


ARTICLE

A Satisficing Theory of Epistemic Justification

Raimund Pils 

University of Salzburg, Department of Philosophy
Email: raimund.pils@sbg.ac.at

Abstract

There is now a significant body of literature on consequentialist ethics that propose satisficing instead of maximizing accounts. Even though epistemology recently witnessed a widespread discussion of teleological and consequentialist theories, a satisficing account is surprisingly not developed yet. The aim of this paper is to do just that. The rough idea is that epistemic rules are justified if and only if they satisfy the epistemic good, i.e., reach some threshold of epistemic value (which varies with practical context), and believing is justified if and only if it follows said rules.

I argue that this alternative to the implicitly established way of thinking in maximizing terms has significant advantages. First, maximizing epistemic value can be unreasonably demanding; second, a satisficing theory can make finding reasonable rules for belief formation and sustenance much more accessible; and third, a satisficing approach is a better alternative to both general subjectivist and maximizing objectivist attempts to spell out epistemic blame.

Keywords: Satisficing; epistemology; value theory; justification; epistemic normativity; teleological epistemology; epistemic consequentialism

1. Structure

I situate my theory in a framework of teleological epistemology (TE). To develop a satisficing theory of epistemic justification, it will be crucial to specify TE. Since there is still a shortcoming of precise formulations, I will develop general principles of epistemic justification before I specify my satisficing theory. A brief sketch of TE will be given in the second section; a general principle of indirect epistemic justification will be developed in the third section. In [section 4](#), I will develop a satisficing theory on this basis. [Section 5](#) shows how this framework applies to an episode in the history of science, gives three motivations for a satisficing theory, and discusses how to determine the threshold for justified believing. [Sections 6 to 8](#) reply to objections of relativism, permissivism, and arbitrariness.

2. The four pillars of teleological epistemology

The basic idea of TE is that there are one or several central values guiding all our epistemic endeavors and *epistemic justification* is aimed at that which has epistemic value. This view can be divided into four pillars:

1. *The Axiological Pillar: The Good is Identified with Some x*

The first step is to identify the good. Most epistemologists today favor ‘veritism,’ i.e., some variation of the view that the fundamental epistemic good is to believe truths and to avoid believing falsehoods.¹ This paper works within a veritist framework, but the given arguments for a satisficing theory are applicable to various other epistemic goods as well.

2. *The Teleological Pillar: The Good Is the End That We Want to Achieve*

This is the explicit teleological pillar: the good is what we are aiming for or trying to achieve.²

3. *The Deontic Pillar: What Is Right Is Explained via the Good*

We believe generally in the right way (justified) if it promotes the epistemic good (the goal) and we believe generally in the wrong way (unjustified) if it impedes on the epistemic good.³

4. *The Normative Pillar: Norms of Belief Are Explained Entirely via the Right*

Finally, we move from right and wrong believing to obligations and permissions to believe. Note that a variety of epistemologists (e.g., Alston 1988; Plantinga 1988) think that there are no obligations to believe; they skip the fourth pillar, giving merely a theory of justification and not a theory of belief norms. This still qualifies them as TE if justification is aimed at (explained by) the good.

This very broad characterization is all I want to refer to with TE. Note, for instance, that this view is not committed to the claim that justification or norms of belief *only* depend on the promotion of the epistemic good *unrestrictively*. Such restrictions can be not to allow trade-offs over time, or interpersonal trade-offs (*cf.* Foley 1993).⁴ Introducing *some* of those restrictions does not indicate a deontological theory. I will call the limiting case of TE without any restrictions ‘epistemic consequentialism.’ Note, however, that this term is sometimes used more broadly.⁵

3. The third pillar: a theory of epistemic justification

In this section, I will develop basic principles of epistemic justification. Fundamentally, as frequently pointed out before,⁶ one can distinguish between subjective vs. objective, direct vs. indirect principles, and, far less recognized in epistemology,⁷ between maximizing and satisficing ones.

¹The term ‘veritism’ is introduced in Goldman (1999). It is also sometimes called “the Jamesian goal” after James ([1896] 2013), “the twin cognitive good” (Carter, Jarvis, and Rubin 2014), “value t-monism” (Pritchard 2010), or “veritistic value monism” (Ahlstrom-Vij 2013).

²Berker (2013, 344ff.) slices the first two pillars for a different purpose into a theory of final value and a theory of overall value.

³For example, Ronzoni (2010, 455) and Williams (1988, 21).

⁴Ahlstrom-Vij and Dunn (2018) allow even consequentialists to restrict the consequence set to that of a single agent (i.e., restricting social trade-offs). Similarly, I think that some restrictions of the consequence set can be conceptualized as being not genuine restrictions that the right puts on the good (not side constraints) but as being part of the axiology. For example, advocating for the ‘truth-now-goal’ (*cf.* Foley 1993) has a clear teleological structure but avoids the implausible epistemic trade-off of me justifiably believing an obvious falsehood now for gaining more epistemic value later. For the problems of epistemic trade-offs in a teleological framework see, e.g., Firth (1981), Fumerton (1995), and Berker (2013). For a recent argument for not allowing epistemic trade-offs as part of an argument for teleological nonconsequentialism, see Littlejohn (2018, 37–40). Ahlstrom-Vij and Dunn’s (2018) restrictions are an answer to trade-offs for the consequentialist.

⁵In today’s normative ethics, ‘teleology’ is typically used in a broad sense, comparable to mine (*cf.* Portmore 2005, 96 n6; Rawls 1971). In epistemology, for a similar broad use see Littlejohn (2018), Wedgwood (2018), and, to some extent Berker (2013). For using deontology broadly and teleology narrowly instead, see Kagan (1998) and Klausen (2009).

⁶For example, Feldman (1988) and Briesen (2016).

⁷But see Berker (2013) and Ahlstrom-Vij and Dunn (2018).

I start by formulating a principle of direct objective epistemic justification, since it is the most straightforward and all other can be viewed as adding various restrictions or extensions.

(EJ-DO) Principle of Direct Objective Epistemic Justification: For all subjects *S* and propositions *p*: believing that *p* is epistemically justified for *S* if and only if believing that *p* promotes the correct epistemic goal(s).

This is a modified version based on considerations by Briesen (2016, 281), Chisholm ([1966] 1989), Feldman (1988, 248), and Klausen (2009, 163). Without going into too many details, I want to mention three points of departure.

First, to ensure generality in my formulation, I left out proposing a relevance condition or proposing a specific epistemic goal. Second, Feldman's explications have quite a deontological flavor, whereas Klausen's are strictly consequentialist. Klausen goes even as far as allowing for interpersonal trade-offs. EJ-DO (and my other upcoming principles) are teleological but stay mostly uncommitted about the amount of restrictions, except they restrict interpersonal trade-offs since I am concerned with agent-centered epistemology.

Third, Feldman (1988, 248) suggests the following principle: "For any proposition *p* and person *S*, if *S* considers *p* then *S* is epistemically obligated to try his best to bring it about that *S* believes *p* if and only if *p* is true." This principle implies that if *p* is false, then *S* is *not obligated* to try to believe. It does not imply that if *p* is false, then *S* is obligated to avoid believing that *p*. *S* is still permitted to believe since not being obligated does not imply not being permitted. Feldman, of course, does want an *obligation* to try to avoid believing falsehoods. Therefore, his principle, as it stands, is too weak. It permits everyone to believe everything; it merely does not obligate everyone to believe anything. Invoking the twin cognitive good is another way of fixing Feldman's principle, something he infers afterwards without him recognizing that it is not implied by his quoted principle.

I will now build on EJ-DO to develop a principle of indirect justification. I develop my satisficing account in a framework of indirect justification because it is the implicitly presupposed meta-epistemology for most theories of justification and I can, as such, cover the most space.⁸ For instance, process reliabilism is considered a version of rule consequentialism, i.e., indirect justification (e.g., Driver 2018, 114; Firth 1981, 12; Goldman 1986; Kornblith 2018, 70; Resnik 1994). The case for evidentialism is similar. For instance, Feldman argues that justification is entirely explained by evidential support (Feldman 1988, 254) but "one's epistemic goal is to get at the truth" (255). Again, one has this intermediate layer of evidential support that mediates between justified believing and the truth goal.

In ethics, the most prominent representative of indirect justification is 'rule-consequentialism.' Rule-consequentialists state that there is some fundamental rule (or set of rules), and their correctness is defined via the conduciveness to the ethical good. Moreover, the acts of a person are morally right if and only if that person follows the rule(s). Hooker (2016, 6.1)—the main contemporary proponent of rule-consequentialism—explicates actualist (i.e., objective) rule-consequentialism as follows: "An act is wrong if and only if it is forbidden by rules the acceptance of which would *actually* result in the greatest good." Using the same technicalities as in the formulation of EJ-DO, this transfers to epistemology as follows:

(EJ-IO) Principle of Indirect Objective Epistemic Justification: For all exhaustive sets of rules *R*: *R* is epistemically justified if and only if following *R* promotes the epistemic good and for all subjects *S* and propositions *p*: believing that *p* is epistemically justified for *S* if and only if believing that *p* follows *R*.

⁸For theoretical reasons why EJ-DO is insufficient see Alston ([1985] 1989, 98–99), Feldman (1988, 246), and, explicitly from a teleological perspective, David (2001, 161–66), who argues that direct justification might not allow for justified false beliefs.

EJ-IO extends the assignment of deontic properties from beliefs to rules. Restrictions can come in again on the level of the epistemic good, by explicating it for example as ‘believing truths and avoiding error (now, on matters of relevance, for S , ...).’

With ‘exhaustive sets of rules R ’ (henceforth simply R), I refer to a complete rule set of belief evaluation—one that correctly assigns *being justified* or *not being justified* to every belief in question. As a contrast, take an incomplete R_c consisting only of the rule not to believe contradictions. It is reasonable to think that R_c promotes the epistemic good and is thus justified. Now, suppose you believe consistently that the earth is flat based on insufficient evidence, ignoring counterevidence, and not trying to gather further evidence. Since you followed the justified rule R_c , and rule following is sufficient for justification, your belief that the earth is flat would be justified. That cannot be correct. It is not correct because even though you did not violate R_c , you violated other rules which you should have considered as well.

4. A satisficing deontic theory

4.a The basic idea of a satisficing approach

So far, EJ-IO uses the vague notion of ‘promotion.’ But what epistemologists typically suppose is that only those R are justified that *maximize* the epistemic good, or, at least, are better than all alternatives in promoting epistemic value. Going for less than value maximization parallels the discussion of ‘satisficing deontic theories’ in normative ethics and since there is no elaborate account of a satisficing theory in epistemology, it is advisable to take a brief look to normative ethics.

Michael Slote (1984) introduces a form of act consequentialism as ‘satisficing consequentialism’ in contrast to the traditional view of ‘maximizing consequentialism’ proposed by Sidgwick (1874).⁹ Fundamentally, Slote proposes that some acts promote the ethical good sufficiently to be right even if they are not maximizing. Transferred to epistemology, this amounts to the following: reaching some threshold of epistemic value is sufficient for a belief, rule, or method of justification to be justified. As a result, one obtains more relaxed principles of belief formation and sustenance with lower epistemic standards, thereby allowing for a wider variety of rational agents than a maximizing approach does.

Note that this account diverges substantially from the basic idea of Simon (1956) in rational choice theory or Gigerenzer and Goldstein’s (1996) considerations of bounded rationality. Slote’s and my account are not replying to the problem of limited information. Even if one knows that F -ing promotes the ethical or epistemic good better than G -ing, one is justified in G -ing as long as G -ing satisfies value. You are morally permitted to buy a present for your mother even if it does not maximize utility and you know that it does not.

Furthermore, Slote and I attempt to explicate conditions for right actions or right ways of believing respectively and our accounts operate in an explicit teleological argumentative structure, including an associated value theory, which a theory of bounded rationality does not. I would even argue that reframing Gigerenzer and Goldstein’s (1996) account as a naturalized epistemology in TE simply makes for a *maximizing theory*. It attempts to prove model theoretically that violating rules of traditional rationality in favor of simplicity increases inferential speed and predictive accuracy. Given that the epistemic good is to increase inferential speed and predictive accuracy, then, what is really shown is that such an alternative set of rules is aimed at maximizing epistemic value. Something similar can be said in the case of Simon’s account. Since my theory is, however, properly satisficing (in that it is not aiming at any form of value maximization), I claim to offer a genuinely new approach. The precise formulation based on EJ-IO is as follows:

⁹Slote borrows this concept from Simon (1956).

(Satisficing EJ-IO) Satisficing Principle of Indirect Objective Epistemic Justification: For all exhaustive sets of rules R : R is epistemically justified if and only if following R satisfies the epistemic good and for all subjects S and propositions p : believing that p is epistemically justified for S if and only if believing that p follows R .

Note that Satisficing EJ-IO does not propose contradicting deontic properties for beliefs in the case of competing exhaustive sets of rules. Satisficing EJ-IO should not be understood as allowing for beliefs to be absolutely justified and unjustified simultaneously. The correct understanding is as proposing a form of relativism: satisficing EJ-IO allows that some beliefs are justified relative to some R but unjustified relative to some R' .

4.b Two wrong perspectives on a satisficing theory

It is frequently maintained that reliabilism in epistemology is analogous to rule consequentialism in ethics (e.g., Driver 2018, 114; Firth 1981, 12; Goldman 1986; Kornblith 2018, 70; Resnik 1994). Ahlstrom-Vij and Dunn even propose an analogy between reliabilism and satisficing consequentialism. If this were correct, then a satisficing theory was present in epistemology all along without recognition; it is merely reliabilism. They write:

Another view says that right beliefs must instead lead to some threshold level of epistemic value. This is the view of the reliabilist—a process can generate justification while failing to be maximally reliable—and it is in this respect analogous to the satisficing consequentialist in ethics. (2018, 5)

I think that this is not the best analogy from ethics. Consider a very strict process-reliabilist rule, such as the following: *believing that p is justified for S at t if and only if S 's believing that p at t results from a maximally reliable cognitive belief-forming process*. Such a rule will in fact not maximize epistemic value, since the rule is much too strict and excludes far too many true beliefs. On the other hand, take a more relaxed rule: *believing that p is justified for S at t if and only if S 's believing that p at t results from a sufficiently reliable cognitive belief-forming process*. After specifying the threshold of sufficient reliability, such a rule is a much better candidate for *maximizing* epistemic value *without* proposing maximal reliability. This should point to the fact that there is a disconnection between the maximization of epistemic value and being a maximally reliable process. If one identifies the two with each other, then one confuses different layers of TE.

It could be objected that Ahlstrom-Vij and Dunn's point is more basic. A maximally reliable belief-forming process produces only true beliefs but any reasonable version of process reliabilism always falls short of this ideal. As such, any reasonable version of process reliabilism allows for beliefs to be justified but not true. This is what classifies them as satisficing theories. In this case, however, the exact same could be said of *any* version of indirect justification. Reliabilism, evidentialism, and any rule-based deontic theory would be a satisficing theory. However, this would draw upon a flawed analogy from ethics, since this would suggest that rule consequentialism in ethics constitutes a satisficing theory as well but, typically, it is not conceived as such. It is also misleading to pick out process reliabilism as something special by associating it with a satisficing theory, since such a use of the term 'satisficing theory' would not differentiate reliabilism from any other theory of indirect epistemic justification.

Ahlstrom-Vij and Dunn's explication can partly be explained by the following disanalogy between a satisficing account in ethics and a satisficing explication of reliabilism in epistemology: Slote developed satisficing consequentialism explicitly for act consequentialism (i.e., a *direct* form of ethical justification), but reliabilism is structurally a form of indirect justification. As already indicated, to draw an analogy to indirect justification in epistemology, one has to propose that *rules* are epistemically justified as long as they meet some threshold of realizing epistemic value. This is what makes beliefs only *indirectly* aiming at the epistemic good. For reliabilism this reads as

follows: reliabilism is a method of justification which is itself justified by promoting (maximizing or satisficing) epistemic value, and beliefs are justified by being formed according to reliabilism.

Let me move to a second wrong perspective on a satisficing theory. Consider a method of justification, call it R_p , stating that if a cognitive belief-forming process reaches some threshold x of reliability, then the beliefs formed by this process are justified.¹⁰ Suppose further that both believing based on some source A and some source B follows R_p . Furthermore, A is more reliable than B . Then, choosing B still follows R_p and is justified even though there is a more reliable source available. If R_p is truly better than all alternatives in maximizing epistemic value, then one could view this as a maximizing theory on the level of principles, but as a satisficing theory on a lower level.

It has to be objected that R_p by itself cannot truly maximize value. For a maximizing theory there is a need for an additional condition which specifies the following: if multiple belief-acquiring options are above the threshold x , then one has to choose the one with the highest reliability. Not adding this condition makes for a truly satisficing theory since it allows for less than maximizing epistemic value. As such, one version of a satisficing theory is simply one that omits this condition. Both believing based on A or based on B follows R_p . R_p proposes the ideal threshold for reliability, and this might be a good reason why choosing B as a source is justified and the beliefs formed accordingly are as well. If the general idea of this case is convincing, then one might consider it already as a first motivation for a satisficing approach to epistemology.

This motivation for a satisficing approach is especially striking if one formulates the case on the level of processes: if a belief-forming process A is above the ideal threshold for reliability x , then it is reasonable to think that beliefs formed according to A are justified, despite there being another more reliable belief-forming process.

5. Three basic motivations and the threshold

5.a A maximizing approach is too demanding

The aim of this subsection is twofold. My exposition so far is quite abstract because my focus is on explicating a general framework. In this section, I will show how my framework applies to a case in the history of science. Secondly, I will generalize the case as a first motivation for a satisficing approach. The general idea is that value maximization is too demanding.

Han Li gives an interesting case as a motivation for his theory of epistemic supererogation which I want to extend as a motivation for a satisficing theory. He writes:

Take, for example, Albert Einstein's famous theory of general relativity. [...] [It] was an epistemic achievement of the highest order. But it seems that much of the evidence that supported Einstein's theory was well known to physics of the time. Probably every sufficiently well-educated physicist was in position to justifiably believe in the theory of general relativity before it was actually discovered. If all this is correct, should we say, then, that all these scientists were actually irrational in failing to believe in the theory of general relativity? This seems far too harsh a verdict. (2018, 350–51)

Li speaks from a perspective of virtue epistemology and is fundamentally concerned with the explication of the rational person, not with justification. However, Li's case can be adopted as a motivation for satisficing TE, insofar as it shows that maximizing is too demanding.

Let me show first that this very general description of the case is actually not problematic for a maximizing theory. Sure, we would not want to demand of Einstein's colleagues to be epistemically obligated to believe the theory of special relativity (SR) or general relativity (GR) as soon as the evidence was available. However, a maximizing theory can also accommodate this verdict. Usual

¹⁰This is very close to Goldman's (1979, 13) process-reliabilist base clause.

versions of maximizing evidentialism or reliabilism do not suppose that any belief produced by a reliable belief-forming process, or any belief sufficiently backed up by evidence, *ought* to be formed. The reason is, as mentioned earlier, that obligations to believe have to be restricted to questions of relevance, or the truth-goal has to be relativized to propositions under consideration. Before anyone worked out GR with its associated propositions, various of those propositions were not even under serious considerations. Consequently, for many of the propositions associated with GR, there were no obligations to believe, even under the presupposition of a maximizing theory.

The more interesting case, contrary to Li's focus about matters before their discovery, are the false beliefs that were still sustained by Einstein's colleagues, contrary to the evidence (interpreted in the right way). To make this case more specific, let me add some historical details on the development of SR and GR and how they compare to Henri Poincaré's relativity theory. There is some disagreement among historians of physics, as to what extent Einstein's and Poincaré's theories were similar (*cf.* Darrigol 2004) but at least the following is agreed upon: similar to Einstein's SR (1905), Poincaré (1904) built on the mathematics of Lorentzian electrodynamics and recognized the invariance of the Maxwell-Lorentz equation under the Lorentz transformations. He developed a principle of relativity almost identical to Einstein's, understood the relativity of simultaneity, and recognized that the measurements of the velocity of light were identical in different inertial frames. One main difference on the path to GR, however, was their treatment of Euclidean geometry, which I will discuss soon.

What separated them were not merely some details in their proposal but rather a different epistemological framework which resulted in a different outcome.¹¹ Poincaré took up part of the Kantian tradition insofar as he relied on organizing principles for organizing perception and empirical data. This he shares with Einstein. The fundamental difference is that for Einstein those organizing principles were freely created (*cf.* Miller 1984, 40–41), whereas for Poincaré, they were synthetic a priori innate principles (1898; [1902] 1952). This led Einstein to a more flexible framework. It made Poincaré (1898, 41–42) believe that “geometry is not an experimental science; experience forms merely the occasion for our reflecting upon the geometrical ideas which pre-exist in us.” Poincaré took a conventionalist approach to geometry (Poincaré 1898; also see Hagar and Memmo 2013, 363). Consequently, he believed until his death that the choice to add time as a fourth dimension is at best an instrumentally useful convention. This goes even as far as Poincaré thinking that the laws of mechanics cannot be empirically disconfirmed (*cf.* Miller 1984, 41). For Einstein, on the other hand, axiomatic geometry was epistemologically on a par with physics. He wrote: “Geometry thus completed is evidently a natural science; we may in fact regard it as the most ancient branch of physics” ([1921] 1970, 8). As such, he believed that *empirical testing* will decide between competing geometries, something Poincaré denied. In 1908, Poincaré argued for the privileged position of Euclidean geometry, whereas Einstein was driving towards GR with its rejection of Euclidean geometry (Miller 1984, 42).

We can see here two different methods of justification at work. Roughly summarized, one method (Einstein's) had fewer a priori aspects and was thus more flexible in accommodating empirical counterevidence, allowing Einstein to realize the failures of Euclidean geometry, whereas another method (Poincaré's) integrated more a priori aspects and was thus more static, making Poincaré ultimately retain the Euclidean picture. Both methods were, however, excellent in getting the physics of relativity right in a way that was truly exceptional, realizing an amount of epistemic value on the subject matter that clearly superseded their colleagues. Even though, as Darrigol (2004, 619) states, “Einstein provided the version that is now judged better.”¹²

¹¹I will pick here some paradigmatic details; for a more extensive analysis on how Poincaré's and Einstein's epistemological assumptions directly impacted the development of their physical theories, see Miller (1984).

¹²There is some disagreement whether Einstein's development of SR was actually better than his colleagues' at accommodating the evidence available at the time. This complicates the case slightly, but these historical details are not relevant to make my general point.

Analyzing this episode in maximizing terms would result in judging that Poincaré's epistemological method was not justified even considering his extraordinary epistemic results simply because Einstein, by using his method, got some additional parts of the picture right. This seems much too harsh a verdict, which thus speaks against a maximizing theory. A satisficing theory evaluates both Einstein's and Poincaré's epistemological methods as justified considering the exceptional epistemic consequences they yielded, which is the preferable verdict.

To generalize this motivation: the method of justification that is better than all alternatives might turn out to be highly complex and nearly impossible to follow for the average rational person or even the highly trained epistemologist. Going for the highest possible standard for evaluating methods might be too demanding and this can motivate lowering the threshold for picking out a method of justification.

5.b The threshold

The Einstein Case motivates the claim that epistemic value maximization demands too much of us. But there is still the question of how much one should be allowed to lower the *threshold*. I will argue that this varies with practical context, and furthermore depends on the amount of epistemic good needed, one's ability to do epistemically good, and the position one is in.

I want to start by putting forward a typical case of pragmatic encroachment on epistemic standards (*cf.* Stanley 2005).

Egg Case: Suppose there is an exceptional method of belief formation M_1 that is 0.95 reliable¹³ in some domain D . Suppose for a subject S , such M_1 is reading the grocery list she wrote yesterday to determine whether there are eggs in her fridge. Suppose there is another method M_2 . M_2 is 0.99 reliable for S in D —trusting the testimony of her daughter's direct perception via a phone call. Suppose further that there is a method M_3 . M_3 is 0.999 reliable for S in D . S could go back home to check herself, eliminating potential errors (buying a detector that checks for fake eggs, etc.). Finally, method M_0 has 0.9 reliability for S in D : trusting her recent memory from yesterday based on her sense perception. M_0 – M_3 are all methods available.

What should S believe? We need to know more about the threshold. In normative ethics, McKay (2021) recently argued that a satisficing theory must determine the threshold via three factors: amount of good needed, one's ability to do good, and uniqueness of one's position. The reasons to choose these three factors are quite analogous in epistemology and can be motivated as follows:

- (a) **Amount of good needed:** If there were a specific amount of good needed in the world, then agents acting as satisficers will be just good enough to get to this ideal state of affairs.

This is the notorious trickiest part for any satisficing theory in ethics. For many standard ethical theories (e.g., classical utilitarianism), there is no maximally good state of the world. There is no natural upper limit. As McKay (2021) admits, this makes a satisficing theory in ethics challenging. Fortunately, epistemologists are in a better position. James ([1896] 2013) already noticed that there is some variability to how much we (ought to) value believing truths compared to avoiding error and practical considerations inform us as such. Some contexts require higher reliability, stronger evidential support, etc., than others. This brings me to the Egg Case. I suggest that S should not form beliefs based on M_0 , but not because 0.9 reliability is bad. Rather, switching from M_0 to M_1 gives S an increase of reliability by 0.05 just by looking at the grocery list. It is thus reasonable to choose M_1 over M_0 (you can adjust the percentages, if you have different intuitions). It also seems perfectly reasonable not to call the daughter, since S 's memory is sufficiently reliable given the situation, but it

¹³A similar story can be told for strong evidence or high confidence.

seems also perfectly reasonable to call her just to make sure. Both M_1 and M_2 are reasonable. Only M_3 seems to be unreasonable and excessive. Hence, in the Egg Case for S , the threshold for a method to be justified is 0.95 reliability. But if the importance of having eggs increases (maybe S needs to prepare a dinner for important dinner guests), M_1 would start to be unreasonable because the value of avoiding error increases. This is the practically informed context relativity of the threshold.

Notice that the lesson here is *not* that only M_3 maximizes value because it has the highest reliability, using M_3 is wrong, and thus we should not be maximizers. This would be the mischaracterization of a satisficing theory warned for in section 4.b. M_3 is actually *not* maximizing value because it values avoiding error too highly and believing truths not highly enough. If we were to demand such high epistemic standards for our methods of justification in such contexts, we would end up hardly believing anything and miss out on many relevant truths. M_3 is not maximizing value.

An additional advantage of this picture with a variable threshold is that it solves the notorious tricky problem of lottery beliefs for veritists. Many commentators (*cf.* e.g., Littlejohn 2012, 79) have raised the following worry: the belief that l (you lose in a large enough fair lottery) is intuitively unjustified, but since the chance of winning can be set arbitrarily low, the threshold for believing would need to be arbitrarily high for this belief to come out as unjustified. Following a rule with such a high threshold would make almost all of our beliefs come out as unjustified. That must be wrong.

My response is that lottery beliefs too have practical implications and believing that l would make you throw away your ticket. This, however, is unreasonable¹⁴ because there is no practical advantage to throwing away your ticket but there is a disadvantage if you win. As such, the value of accurately believing that l if it is true approaches 0 and the value of accurately believing that l is false if it is false approaches infinity. Thus, you are not justified believing that l .¹⁵

In sum, the amount of good needed explains the range of the threshold by the value we put on believing truths versus avoiding error, which is informed by the context of our practical demands. My explication of a satisficing theory and some maximizing theories can give such a range.¹⁶ However, additionally satisficing theories propose that this range is relative to the subject based on two factors:

(b) **One's ability to do good:** ought implies can, but this grades off.

In the Egg Case, consider a person S_2 who forms his beliefs based on his shopping lists with only 0.9 reliability. Furthermore, S_2 could also use M_2 (calling his daughter with 0.99 reliability), and these are all methods available. Now 0.9 (i.e., M_0 —just using her memory) was not good enough for S_1 because the jump to 0.95 was too easy. Thus, we set the threshold at 0.95. For S_1 , however, getting over 0.95 would mean applying M_2 , which is significantly more costly. We conclude that for S_2 , 0.9 is sufficient. (Again, you can modify the numbers if you have different intuitions.) S_2 is justified to believe based on 0.9 reliability while S_1 is not, and this is explained by their different abilities to do epistemically good.¹⁷

We get the same subject relativity in the Einstein Case. Einstein, Poincaré, and their scientific peers all had different abilities to achieve epistemically good outcomes. Einstein and Poincaré were

¹⁴This part is quite similar to Littlejohn's solution.

¹⁵Of course, your credences in l still ought to be very high. However, in some contexts and for some propositions high credences will not imply full belief. See Dorst (2017, 186–92) for such a proposal.

¹⁶Recently, Dorst (2017, 188) argues quite analogously in a Lockean framework that the magnitudes of the values of true and false belief determine the threshold, and this threshold varies with context and proposition in question. The threshold varies, but Dorst's theory is still a maximizing one (as he recognizes) since for every proposition p in a specified context c there is only one correct threshold, the one that maximizes value.

¹⁷This strategy is quite similar to Rogers's (2010) and Chappell's (2019) strategies for their satisficing theory. They only allow nonmaximization if the maximizing strategy is too costly or an undue burden.

extraordinary individuals, so we demand epistemically more from them. We lower the threshold for the average scientist, but we still expect to change their mind to the high standards of the community and the expertise of their peers after the results are more accessible.

(c) **Unique position:** Some are simply in a better position to do good than others.

In the epistemic realm, this applies especially to scientists and public science communicators. We expect scientists to have extraordinary responsibility for accuracy. Einstein's approach, even though supererogatory, is still not beyond expectancy for a scientist because we expect epistemic supererogation from scientists. If we think that science is one of our best knowledge-generating source, then it is reasonable to think that in a mature progressive science at least the median scientist fulfills their epistemic demands. This is of course open to further evaluation.¹⁸ Poincaré clearly fulfilled such demands, and thus his epistemic methods were justified.

This version of a satisficing theory can also answer Bradley's (2006) well-known objection for ethical satisficing consequentialism. Bradley's criticism transferred to (ego-centered) epistemology is as follows: if one's total epistemic value is above the threshold n , one could permissibly believe a random falsehood as long as one's total epistemic value does not drop below n . The presented theory does not fall prey to this objection. We cannot simply choose a method of justification with lower (expected) epistemic value if this choice is not motivated by its higher practicability on the background of one's ability to do good and one's unique position.

5.c Simplifying rule finding

As in ethics so in epistemology, it is contentious what counts as a rule. This goes together with the most iconic criticism of rule consequentialism: it just collapses into act consequentialism. The idea is as follows. Sometimes an act promotes the good but violates a general good-promoting rule such as "Don't steal." Thus, adding an exception clause gets a rule that is better in promoting value. Then, adding exception clause after exception clause in the same fashion will ultimately make act prescriptions of rule consequentialism and act consequentialism to be coextensive. The standard reply is (*cf.* Hooker 2016, 8) that adding that many exception clauses simply makes consequentialism impractical; it plausibly even contributes to a wrong application of exception clauses due to complexity. The lesson: rules should be simple! Interestingly, we find similar reasoning in epistemology. Williamson (2002, 223) argues against indirect *subjective* principles of epistemic justification by pointing out that a rule such as "Add salt when the water boils" is far superior to "Do what appears to you to be adding salt when what appears to you to be water appears to you to boil." It is not that advocates for the simpler rule deny that sometimes we mistake salt for a similar looking ingredient, it is rather that motivating one to check whether it is really salt is simply done by the rule "add salt!" Williamson then transfers this to epistemology:

Just as we can follow the rule 'Add salt when the water boils', so we can follow the rule 'Proportion your belief in a proposition to its probability on your evidence'. Although we are sometimes reasonably mistaken or uncertain as to what our evidence is and how probable a proposition is on it, we often enough know enough about both to be able to follow the rule. (2002, 223)

This not only is an argument against subjective principles of epistemic justification: it also motivates that it is valuable to have rules that are easy to follow.

One advantage of a satisficing theory is that it is easier to follow than a maximizing one. A satisficing theory can do away with the task of evaluating which exhaustive set of rules is better than

¹⁸For example, degenerate research programs (Lakatos 1974).

all alternatives in promoting epistemic value. Finding some *R* that satisfies epistemic value is already sufficient. This makes it much easier to find an *R* to follow since as soon as some *R* is good enough you do not have to look any further for better alternatives.

5.d Epistemic blameworthiness

Two related objections to the motivation of section 5.c—i.e., simplifying rule finding—arise. First, one might worry that it mixes up objective principles of epistemic justification with the pragmatics of rule application or rule selection. Second, one might worry that, *epistemic* rules in particular are not something (rational) agents explicitly, consciously do follow, and thus considerations about the practicality of rule following are beside the point.

I reply that at least for those philosophers who think that the practicality of rules has some bearing on whether a rule is right or wrong (see the discussion of Hooker and Williamson in the last section), the motivation has force. If, however, one is from the opposing camp, then I will concede this one motivation. Still, even if one thinks that an idealized theory of epistemic justification should be stripped away of all subjectivist, relativist aspects, and of all practical considerations of rule finding, then, I will argue, a satisficing theory is still valuable to spell out the concept of ‘epistemic blameworthiness.’

Differentiating wrongness from blameworthiness is widespread in ethics (cf. Hooker 2016, 6.1).¹⁹ Recently, Driver (2018, 118) argued that our critical practice warrants such a separation in epistemology as well. Kvanvig ([2005] 2014, 361) argues that it is reasonable to differentiate Epistemic Blameworthiness with its subjectivist aspects from purely objective Justification. Furthermore, Singer (2018) argues that adopting this differentiation in epistemology is an important lesson one should draw from ethical consequentialists to avoid common objections. Thus, one might reformulate the proposed satisficing deontic theory as an explication of the notion of Epistemic Blameworthiness but stick with a maximizing deontic theory when explicating Justification. This preserves a completely nonrelativist handling of justification as an ideal epistemic theory, but also preserves the motivation from section 5.a that a maximizing theory is too demanding, such that rational agents are at least not epistemically blameworthy if they are value satisficer.

A satisficing explication of Epistemic Blameworthiness has also advantages compared to a completely subjective explication. Merely believing that your beliefs maximize the epistemic good cannot be sufficient to not be epistemically blameworthy. In the Flat Earth Case, suppose you believed that you did everything epistemically right by believing that the earth is flat. You still should be epistemically blamed for holding such a belief since you made too many mistakes.²⁰ But as soon as one restricts subjectivity to something like “you are not epistemically blameworthy if and only if it is justified to believe that your beliefs maximize the epistemic good,” one is just back to equating blameworthiness with justification. Consequently, for spelling out epistemic blameworthiness, mere believing is not good enough, but justified believing goes too far. Now, there might be a middle ground, but that is very hard to spell out correctly. A satisficing explication of Epistemic Blameworthiness does this in a very natural way. The epistemic standards are lower than the standards for justified believing but higher than mere believing.

¹⁹Note, however, that if one differentiates blameworthiness from wrongness in ethics, then blameworthiness is, typically, explicated as a form of expectabilist consequentialism (cf. Hooker 2016, 6.1), i.e., a subjective form of wrongness, or it is explicated as objective wrongness, but one adds some additional conditions, especially control and knowledge conditions. What I am suggesting is quite different. It ties blameworthiness to *not satisficing value* and wrongness to *not maximizing value*. I am not aware of any theory in ethics that does that.

²⁰For the implausibility of purely subjectivist principles of epistemic justification, also cf. Alston ([1985] 1989, 88–89).

6. Objection 1: Contradictory instructions

I want to turn now to three expected objections. What is ultimately expected from a complete normative theory in epistemology is to give clear instructions about obligations to believe or at least justified believing. There is a worry that a satisficing theory cannot live up to this demand. If justified R justifies believing that p and justified R' does not, how would an agent decide whether to believe that p ? What cannot be wanted is that a complete epistemic normative theory gives contradictory instructions for belief formation and sustenance. I call this the Contradictory Instructions Objection:

Contradictory Instructions Objection: For all subjects S and propositions p : the correct overall normative theory of belief formation and sustenance should neither imply that S is obligated to believe that p at t and obligated not to believe that p at t nor that S is permitted to believe that p at t and is not permitted to believe that p at t .

To address this objection, a detour to the intersection of the third and the fourth pillar of TE is necessary, i.e., it has to be addressed how to get from justifications to obligations to believe. The most straightforward principle is as follows:

(BN-So) Straightforward Norm of Belief Formation and Sustenance, Obligation: For all subjects S and propositions p : S is epistemically obligated to believe that p if and only if S is justified (according to *satisficing* EJ-IO) in believing that p . S is epistemically obligated not to believe that p if and only if S is not justified in believing that p .

I formulated this principle with *obligations* because in epistemology it is not sufficient to merely speak about permissions. If there were no obligations to believe, then you would always be permitted to withhold beliefs and would not get to your goal of believing truths. As such, it would defy the whole teleological motivation. This, however, puts some restrictions on the correct epistemic goal. If the goal were merely believing truths and avoiding error, then it seems that one would be obligated to believe all kind of propositions that one did not even consider and were not of relevance because it increases epistemic value. Thus, there needs to be some kind of relevance condition,²¹ or some kind of restriction to propositions under consideration.²² Furthermore, one cannot omit the additional condition for *obligations not to believe*. Without that condition, S would be merely not obligated to believe that p but still allowed to believe that p (which Feldman did not recognize, as mentioned earlier) if there were no justification not to believe that p . This would be too weak.

BN-So is as elegant as it is intuitive. What speaks for it is that it draws the most straightforward connection from the epistemic goal to belief obligations. However, it could be troublesome in connection with a satisficing deontic theory since it seems to run into the Contradictory Instructions Objection: for conflicting exhaustive sets of rules which satisfy epistemic value, there will be some p which S is obligated to believe according to R and will be obligated not to believe according to R' . This outcome can be circumvented if a person simply chooses one and only one R at any given time. After all, the basic idea of a satisficing theory is to give the subject the freedom to choose between satisficing R and R' . Thus, by choosing precisely one R at t , no contradictory instructions arise. The question arises now, however, if we can live with the resulting permissivism and relativism, which will be the topic of the next two sections.

²¹For the relevance truth goal, see Haack (1993, 199), Harman (1986), Briesen (2016), and Khalifa (2020). For famous counter-cases, see Grimm (2008, 742).

²²See David ([2005] 2014, 365–56).

7. Objection 2: Permissivism and arbitrariness

The relativist consequences of a satisficing theory imply ‘Permissivism,’ i.e., a violation of what Feldman ([2006] 2011, 148) calls the ‘Uniqueness’ thesis that “a body of evidence justifies at most one proposition out of a competing set of propositions [...] and that it justifies at most one attitude toward any particular proposition.”

Suppose there are two evidentialist methods that satisfy epistemic value but justify contradicting sets of beliefs (see section 5.1), then Uniqueness is violated because the same body of evidence (i) justifies competing propositions relative to different R and (ii) justifies more than one attitude toward some propositions because it can obligate one to believe that p relative to R and obligate one not to believe that p relative to R' , and thus justify two doxastic states towards p , i.e., believing and withholding belief.

First, if one is already convinced by a Permissivist picture, then this further empowers a satisficing framework. As Li (2018, 351) recognizes: “Many philosophers have found the uniqueness thesis to be intuitively implausible, but there are relatively few fully developed epistemic theories that can explain why it is false.” A satisficing theory of epistemic justification does just that: since there will be closely related evidentialist rules that satisfy value but justify different sets of beliefs, a satisficing theory of epistemic justification can explain in a very natural way why Uniqueness is false.²³

But what if one thinks violating Uniqueness is wrong? While this is not the place to solve the Permissivism debate, I want to argue why the Permissivism of a satisficing theory can be quite reasonable. It is useful to differentiate here two versions of violating Uniqueness: Inter-Personal Permissivism—i.e., violating Uniqueness across persons—and Intra-Personal Permissivism—i.e., violating Uniqueness within one single person (cf. Kelly 2014). Most Permissivists defend Inter-Personal Permissivism. For instance, Schoenfield (2014) argues that it is reasonable for epistemic standards to vary between individuals (also cf. Kelly 2014; Podgorski 2016; Simpson 2017). Already simply involving a veritist framework with a twin cognitive goal can make a form of Permissivism plausible since, arguably, the balancing of the value of believing truths and avoiding error is not exhausted by epistemic reasons.

Intra-Personal Permissivism, on the other hand, is much less defended.²⁴ Allowing one single rational agent to have varying epistemic standards seems unintuitive. The presented satisficing theory sanctions such Permissivism. It can be put most bluntly with what I call the ‘Arbitrary Switching Objection.’ Consider there are two competing R that satisfy value for S . What would prevent S to follow R at t and R' at t' in the exact same epistemic situation arbitrarily. Nothing, it seems, since both satisfy value. Maybe S believes p based on R on weekdays and $\neg p$ based on R' on weekends. The objection is: if Satisficing EJ-IO sanctions *this*, then there must be something wrong with it.

White (2005) challenges Permissivism by arguing precisely that it leads to some unacceptable arbitrariness of one’s doxastic attitudes (also cf. Kolodny 2007, 248). In response to White’s arguments, Permissivists typically respond by trying to avoid arbitrariness at least of epistemic standards²⁵ but it seems that Satisficing EJ-IO has arbitrariness built into it. It is also resilient to Jackson’s (2021) solution for the Intra-Personal Permissivist. She argues with cases of supererogation for permissive switching. If supererogative reflection on your beliefs suggests revising your belief, then such change is permitted but not required because it is supererogative. Thus, revising and not revising is permissible. Whatever we think about this solution, it does not work for the present Arbitrary Switching Objection because switching appears to be based on some epistemic procedure, such as rational reflection.

²³Li (2018) makes the structurally equivalent argument for his theory of epistemic supererogation.

²⁴but see Jackson (2021); Ru Ye (2019).

²⁵See, e.g., Kelly (2014); Meacham (2016); Schoenfield (2014); Simpson (2017); also cf. Ye (2019).

What is then the solution? Ye (2019) pushes Permissivism further than Jackson and concludes that arbitrariness is simply fine. His answer is, in short, that there is nothing unreasonable about choosing an action between two permissible actions arbitrarily, and analogously the same holds for belief. This solution works for a satisficing theory as well, but I do not even have to go as far as Ye.

First, contrary to Ye, my view is not committed to Permissivism about credences. For example, in lottery beliefs your credences that your ticket will lose can be incredibly high without you believing that it loses. Since practical considerations impede on the balancing of the truth goal but not necessarily on credences, there is no arbitrariness of credences. As noted earlier, following Dorst (2017), this still can preserve Lockeanism since the threshold is variable. As such, the arbitrariness of choosing epistemic standards in my framework is already weakened, since uniquely tailoring credences to evidence is preserved.

Second, what about the arbitrariness of choosing rules that satisfy epistemic value? Similar to current responses of arbitrariness (*cf.* Schoenfield 2014, 199; Meacham 2016, 472–73), I respond that any *R* still singles out one specific belief; believing is not arbitrary (also see section 6.1). Ye (forthcoming; also *cf.* White 2005, 452; Feldman 2007, 205–6) objects to this move that this just pushes the arbitrariness to the epistemic standards or rule choice. Note, however, since I, contrary to Ye, explain the choice of epistemic standards via practical considerations, the choice of epistemic standards is *not* arbitrary. It is not arbitrary to believe whether there are eggs in my fridge based on a higher effort-/high reliability method or a lower effort-/slightly lower reliability method. There is a practical trade-off, and this trade-off can explain the *R* (with its implied epistemic standards) that one chooses. So going beyond Ye, there is not simply an analogy between permissive action and permissive belief. Permissive actions straightforwardly lead to permissive beliefs, and since there is nothing strange about cases of permissive action, there is nothing strange about permissive believing either. If I buy eggs on weekends by consulting my shopping list and on weekdays by calling my daughter, I am not irrational because both are reasonable things to do. If I form beliefs about the eggs on weekdays by consulting my shopping list and on weekends by calling my daughter, I am not irrational either, as long as both methods satisfy value.²⁶ Note further, that the presented theory actually blocks *completely* arbitrary switching, it just does not block *epistemically* arbitrary switching. If I consulted my shopping list, and thus believe that I have eggs in my fridge, afterwards call my daughter (higher reliability, stronger evidence) and she is saying that I do not, then I cannot simply switch back to looking at my shopping list believing that I have eggs in my fridge after all. If I have the result of applying various methods of justification available, then I am not allowed to base my belief on an inferior one because this would practically *not* be motivated. There is no practical advantage in trusting the lower reliability method.

Now one might worry that such belief switching is much stranger in complex interdependent belief systems. I do not think so, and this is independent of a satisficing theory. Consider the following case from personal experience: by thinking about the question of scientific realism (at least for some scientific theories), some days I found myself to be more of a realist based on abductive reasoning, and other days I was more of an antirealist based on avoiding inflationary metaphysics. Nothing changed in the evidence I had but my methods of justification were different. I still think both solutions are quite reasonable, and I would not view myself as being unreasonable to prefer one over the other.²⁷ If cases such as this are reasonable epistemic practice, then they motivate that there is generally nothing wrong with switching.

²⁶Note that if there were situations where switching would lead to problems in your action guidance, this would be a practical reason to avoid switching in such cases. But for cases, as in the egg case, where switching seems practically fine, there is no reason not to allow switching.

²⁷For versions of voluntarism in the realism debate see e.g., Chakravartty (2018) and van Fraassen (2002).

8. Objection 3: Relativism

As a last objection, I want to reply to more general relativist concerns. I will show that even a maximizing theory plausibly runs into relativism and violates Uniqueness, and thus the worries about a satisficing theory should be reconsidered.

A maximizing principle of epistemic justification proposes that R is justified if and only if it maximizes value or at least promotes epistemic value better than all alternatives R' . Such a principle might still justify conflicting R in the following cases:

- (i) Two (or more) conflicting R promote the correct epistemic goal equally well but better than all alternatives.
- (ii) Two (or more) conflicting R promote the correct epistemic goal better than all alternatives but are incommensurable or incomparable in their promotion of the correct epistemic goal between themselves.

Maximizing EJ-IO would still justify those conflicting R in both cases, and one gets all the consequences of violating Uniqueness. Some conditions need to be fulfilled, however.

Merely not knowing which one of two conflicting R is better than all alternatives in promoting the epistemic good is not yet problematic. The only thing that matters is whether those R are *actually* better or not, since we are arguing in a framework of objective justification. Thus, the statement in (i) can only mean that R justifies a subject to believe one set of propositions and R' justifies a subject to believe a different set of propositions and it is exactly equally valuable to believe one set as it is to believe the other. If, for example, the correct explication of veritism is to improve the ratio of true over false beliefs for a reasonably sized set of beliefs, then both sets would have to have the exact same ratio. In real-life cases that might not happen too often.

The same line of reasoning applies to (ii). It is not an objection in a framework of objective justification that two R are merely incommensurable *given what S knows*. But what could objective incommensurability mean? Since it is clearly defined how rules derive their value—i.e., only by maximizing the epistemic good—the most plausible way for this to arise is some variability of correct epistemic goods. For instance, there might not be a justification for how much weight one should put on believing truths compared to avoiding error, or at least there might be some permissible spectrum. Then, some R could be better in promoting one epistemic good and some R' could be better in promoting another. If those goods are incommensurable, then both R would be objectively incommensurable.

If cases of the kind (i) or (ii) truly exist, then even a maximizing theory would violate Uniqueness, and has relativist consequences. These consequences would be limited in scope compared to a satisficing theory, but still, if one had reservations against a satisficing theory because one wants to preserve a *completely* nonrelativistic objectivist theory of justification, then such reservations have to be reconsidered if a maximizing theory cannot deliver that either. Then the difference between a maximizing and a satisficing theory of epistemic justification is just a matter of degree and not a matter of kind regarding their relativist consequences.

9. Conclusion

I put forward a satisficing theory of indirect epistemic justification in a framework of teleological epistemology stating that rules or methods of epistemic justification are justified if and only if they satisfy the epistemic good—i.e., reach some threshold of epistemic value (which varies with practical context)—and believing is justified if and only if it follows said methods or rules.

I argued that by drawing the correct analogy from normative ethics, a genuine satisficing approach has to be understood as putting forward a form of subjective relativism and Permissivism. There is some leeway for rationality where it is up to the subject to choose between different

methods or rules of epistemic justification as long as epistemic value is satisfied. The threshold varies with practical context, and furthermore depends on the amount of epistemic good needed, one's ability to do epistemically good, and the position one is in.

I gave three motivations: (i) a maximizing approach is too demanding (Einstein Case, Egg Case), whereas a satisficing theory can give the right verdict; (ii) a satisficing theory can make finding reasonable rules for belief formation and sustenance more accessible; and (iii) a satisficing approach has major advantages for spelling out the concept of Epistemic Blameworthiness, since, contrary to the maximizing objectivist, it can preserve the intuition that epistemic standards for epistemic blameworthiness are lower than those for justified believing, but contrary to the general subjectivist, it does not lead to implausibly low epistemic standards.

I argued that the framework implies violating Uniqueness, which Permissivists will regard as a strength because a satisficing theory can naturally explain this violation. For opponents of Permissivism, I argued that the resulting Intra-Personal Permissivism is weaker than expected because it does not imply Credence Permissivism, and is, furthermore, a direct consequence from a plausible kind of Permissivism about action. Finally, I argued that a maximizing alternative is most likely not able to avoid all relativist consequences either and will violate Uniqueness as well, so reservations against a satisficing theory should be reconsidered.

Acknowledgments. I want to thank Philipp Schoenegger and two anonymous reviewers of the *Canadian Journal of Philosophy* for their very helpful comments on an earlier version.

Raimund Pils is a university assistant at the University of Salzburg, Austria. His primary research interests are in philosophy and epistemology. Currently, he is working on transferring various insight from epistemic value theory, such as consequentialist theories, to the scientific realism debate.

References

- Ahlstrom-Vij, Kristoffer. 2013. "In Defense of Veritistic Value Monism." *Pacific Philosophical Quarterly* 94 (1): 19–40.
- Ahlstrom-Vij, Kristoffer H., and Jeffrey Dunn. 2018. "Introduction: Epistemic Consequentialism." In *Epistemic Consequentialism*, edited by Kristoffer Ahlstrom-Vij and Jeffrey Dunn, 1–22. Oxford: Oxford University Press.
- Alston, William P. 1988. "The Deontological Conception of Epistemic Justification." *Philosophical Perspectives* 2: 257–99.
- Alston, William P. [1985] 1989. "Concepts of Epistemic Justification." *The Monist* 68 (2); re-issued in *Epistemic Justification: Essays in the Theory of Knowledge*. Ithaca, NY: Cornell University Press: 81–114.
- Berker, Selim. 2013. "Epistemic Teleology and the Separateness of Propositions." *Philosophical Review* 122 (3): 337–93.
- Bradley, Ben. 2006. "Against Satisficing Consequentialism." *Utilitas* 18: 97–108.
- Briesen, Jochen. 2016. "Epistemic Consequentialism: Its Relation to Ethical Consequentialism and the Truth-Indication Principle." In *Epistemic Reasons, Norms, and Goals*, edited by Pedro Schmechtig and M. Grajner, 277–306. Boston: De Gruyter.
- Carter, Adam J., Benjamin W. Jarvis, and Katherine Rubin. 2014. "Varieties of Cognitive Achievement." *Philosophical Studies* 172 (6): 1603–23.
- Chakravartty, Anjan. 2018. "Realism, Antirealism, Epistemic Stances, and Voluntarism." In *The Routledge Handbook of Scientific Realism*, edited by Juha Saatsi. New York: Routledge.
- Chappell, Richard Y. 2019. "Willpower Satisficing." *Nous* 53 (2): 251–65.
- Chisholm, Roderick M. [1966] 1989. *Theory of Knowledge*. 3rd ed. Englewood Cliffs, NJ: Prentice Hall.
- Darrigol, Oliver. 2004. "The Mystery of the Einstein–Poincaré Connection." *Isis* 95 (4): 614–26.
- David, Marian. 2001. "Truth as the Epistemic Goal." In *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*, edited by Matthias Steup, 151–69. Oxford: Oxford University Press.
- David, Marian. [2005] 2014. "Truth as the Primary Epistemic Goal: A Working Hypothesis." In *Contemporary Debates in Epistemology*, 2nd ed., edited by Matthias Steup, John Turri, and Ernest Sosa, 363–77. Chichester: Wiley Blackwell.
- Dorst, Kevin. 2017. "Lockeans Maximize Expected Accuracy." *Mind* 128 (509): 175–211.
- Driver, Julia. 2018. "The 'Consequentialism' in 'Epistemic Consequentialism.'" In *Epistemic Consequentialism*, edited by Kristoffer Ahlstrom-Vij and Jeffrey Dunn, 113–22. Oxford: Oxford University Press.
- Einstein, Albert. 1905. "Zur Elektrodynamik bewegter Körper." *Annalen der Physik* 322 (10): 891–921.
- Einstein, Albert. [1921] 1970. *The Meaning of Relativity: Four Lectures Delivered at Princeton University*. 5th ed., translated by Edwin P. Adams. Princeton, NJ: Princeton University Press.
- Feldman, Richard. 1988. "Epistemic Obligations." *Philosophical Perspectives* Vol. 2, *Epistemology*: 235–56.

- Feldman, Richard. [2006] 2011. "Reasonable Religious Disagreements." In *Social Epistemology: Essential Readings*, edited by Alvin Goldman and Dennis Whitcomb, 137–57. Oxford: Oxford University Press.
- Feldman Richard. 2007. "Reasonable Religious Disagreements." In *Philosophers without Gods*, edited by Louise Antony, 194–214. Oxford: Oxford University Press.
- Firth, Roderick. 1981. "Epistemic Merit, Intrinsic and Instrumental." *Proceedings and Addresses of the American Philosophical Association* 55 (1): 5–23.
- Foley, Richard. 1993. *Working without a Net. A Study of Egocentric Epistemology*. New York: Oxford University Press.
- Fumerton, Richard. 1995. *Metaepistemology and Skepticism*. Lanham, MD: Rowman & Littlefield.
- Gigerenzer, Gerd, and Daniel G Goldstein. 1996. "Reasoning the Fast and Frugal Way: Models of Bounded Rationality." *Psychological Review* 103 (4): 650–69.
- Goldman, Alvin I. 1979. "What Is Justified Belief." In *Justification and Knowledge*, edited by G. Pappas, 1–25. Boston: Reidel.
- Goldman, Alvin I. 1986. *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.
- Goldman, Alvin I. 1999. *Knowledge in a Social World*. Oxford: Oxford University Press.
- Grimm, Stephen R. 2008. "Epistemic Goals and Epistemic Values." *Philosophy and Phenomenological Research* 77: 725–44.
- Haack, Susan. 1993. *Evidence and Inquiry. Towards Reconstruction in Epistemology*. Oxford: Blackwell.
- Hagar, Amit, and Meir Memmo. 2013. "The Primacy of Geometry." *Studies in History and Philosophy of Modern Physics* 44: 357–64.
- Harman, Gilbert. 1986. *Change in View: Principles of Reasoning*. Cambridge, MA: MIT Press.
- Hooker, Brad. 2016. "Rule Consequentialism." In *The Stanford Encyclopedia of Philosophy* (Winter), edited by Edward N. Zalta. <https://plato.stanford.edu/archives/win2016/entries/consequentialism-rule>.
- Jackson, Elizabeth. 2021. "A Defense of Intrapersonal Belief Permissivism." *Episteme* 18 (2): 313–27.
- James, William. [1896] 2013. "The Will to Believe." In *Reasons and Responsibility. Readings in Some Basic Problems of Philosophy*, 15th ed., edited by John Feinberg and Russ Shafer-Landau, 129–37. Australia: Wadsworth.
- Kagan, Shelly. 1998. *Normative Ethics*. Boulder, CO: Westview Press.
- Kelly, Thomas. 2014. "Evidence Can Be Permissive." In *Contemporary Debates in Epistemology*, 2nd edition, edited by Matthias Steup, John Turri, and Ernest Sosa, 298–312. Chichester: Wiley Blackwell.
- Khalifa, Kareem. 2020. "Understanding, Truth, and Epistemic Goals." *Philosophy of Science* 87 (5): 944–56.
- Klausen, Søren H. 2009. "Two Notions of Epistemic Normativity." *Theoria* 75: 161–78.
- Kolodny, Niko. 2007. "IX—How Does Coherence Matter?" *Proceedings the Aristotelian Society* 107 (1): 229–63.
- Kornblith, Hilary. 2018. "The Naturalistic Origins of Epistemic Consequentialism." In *Epistemic Consequentialism*, edited by Kristoffer Ahlstrom-Vij and Jeffrey Dunn, 70–84. Oxford: Oxford University Press.
- Kvanvig, Jonathan L. [2005] 2014. "Truth Is Not the Primary Epistemic Goal." In *Contemporary Debates in Epistemology*, 2nd ed., edited by Matthias Steup, John Turri, Ernest Sosa. Chichester: Wiley Blackwell: 352–62.
- Lakatos, Imre. 1974. "Science and Pseudoscience." In *Philosophy in the Open*, edited by Godfrey Vesey. Milton Keynes: Open University Press.
- Li, Han. 2018. "A Theory of Epistemic Supererogation." *Erkenntnis* 83 (2): 349–67.
- Littlejohn, Clayton. 2012. *Justification and the Truth-Connection*. Cambridge: Cambridge University Press.
- Littlejohn, Clayton. 2018. "The Right in the Good. A Defense of Teleological Non-Consequentialism." In *Epistemic Consequentialism*, edited by Kristoffer Ahlstrom-Vij and Jeffrey Dunn, 23–47. Oxford: Oxford University Press.
- McKay, Daniel. 2021. "Solving Satisficing Consequentialism." *Philosophia* 50: 149–57.
- Meacham, Christopher J. G. 2016. "Ur-priors, Conditionalization, and Ur-prior Conditionalization." *Ergo: An Open Access Journal of Philosophy* 3 (17). <https://doi.org/10.3998/ergo.12405314.0003.017>.
- Miller, Arthur I. 1984. "Poincaré and Einstein." *Imagery in Scientific Thought Creating 20th-Century Physics*. New York: Springer.
- Plantinga, Alvin. 1988. "Chisholmian Internalism." In *Philosophical Analysis. Philosophical Studies Series*, vol 39, edited by D.F. Austin, 127–51. Springer: Dordrecht.
- Podgorski, Abelard. 2016. "Dynamic Permissivism." *Philosophical Studies* 173 (7): 1923–39.
- Poincaré, Henri. 1898. "On the Foundations of Geometry." Translated by Thomas J. McCormack. *Monist* 9 (1): 1–43. <https://www.jstor.org/stable/27899007>.
- Poincaré, Henri. [1902] 1952. *La Science et l'Hypothese*. In *Science and Hypothesis*, translator unknown. New York: Dover.
- Poincaré, Henri. 1904. *L'état actuel et l'avenir de la Physique mathématique*. Lecture delivered on 24 September 1904 to the International Congress of Arts and Science, Saint Louis, Missouri. *Bulletin of Mathematical Sciences* 28: 302–24.
- Portmore, Douglas W. 2005. "Combining Teleological Ethics with Evaluator Relativism: A Promising Result." *Pacific Philosophical Quarterly* 86: 95–113.
- Pritchard, Duncan. 2010. "Knowledge and Understanding." *The Nature and Value of Knowledge: Three Investigations*, edited by Duncan Pritchard, Alan Millar, Adrian Haddock, 3–88. Oxford: Oxford University Press.
- Rawls, John. 1971. *A Theory of Justice*. Revised ed. Cambridge: Harvard University Press.
- Resnik, David. 1994. "Epistemic Value: Truth or Explanation?" *Metaphilosophy* 25 (4): 348–61.
- Rogers, Jason. 2010. "In Defense of a Version of Satisficing Consequentialism." *Utilitas* 22 (2): 198–221.

- Ronzoni, Miriam. 2010. "Teleology, Deontology, and the Priority of the Right: On Some Unappreciated Distinctions." *Ethical Theory and Moral Practice* 13 (4): 453–72.
- Schoenfield, Miriam. 2014. "Permission to Believe: Why Permissivism Is True and What It Tells Us about Irrelevant Influences on Belief." *Noûs* 48 (2): 193–218.
- Sidgwick, Henry. 1874. *The Methods of Ethics*. London: Macmillan.
- Simon, Herbert A. 1956. "Rational Choice and the Structure of the Environment." *Psychological Review* 63 (2): 129–38.
- Simpson, Robert M. 2017. "Permissivism and the Arbitrariness Objection." *Episteme* 14 (4): 519–38.
- Singer, Daniel J. 2018. "How to Be an Epistemic Consequentialist." *The Philosophical Quarterly* 68 (272): 580–602.
- Slote, Michael. 1984. "Satisficing Consequentialism." *Proceedings of the Aristotelian Society* 58 (supp. vol.): 139–63.
- Stanley, Jason. 2005. *Knowledge and Practical Interests*. New York: Oxford University Press.
- Van Fraassen, Bas C. 2002. *The Empirical Stance*. New Haven, CT: Yale University Press.
- Wedgwood, Ralph. 2018. "Epistemic Teleology. Synchronic and Diachronic." In *Epistemic Consequentialism*, edited by Kristoffer Ahlstrom-Vij and Jeffrey Dunn, 85–112. Oxford: Oxford University Press.
- White, Roger. 2005. "Epistemic Permissiveness." *Philosophical Perspectives* 19: 445–59.
- Williams, Bernard. 1988. "Consequentialism and Integrity." In *Consequentialism and Its Critics*, edited by Samuel Scheffler, 20–50. Oxford: Oxford University Press.
- Williamson, Timothy. 2002. *Knowledge and Its Limits*. Oxford: Oxford University Press.
- Ye, Ru. 2019. "The Arbitrariness Objection against Permissivism." *Episteme*: 1–20.