

Swiss Skeptics Discussion Paper Series
Volume 1, Issue 1

A generalized definition of critical thinking

Marko Kovic*

June 2016

*marko.kovic@skeptiker.ch / +41 76 335 06 17

Swiss [Skeptics]

association for critical thinking

About the Swiss Skeptics

Swiss Skeptics – Association for Critical Thinking is an independent nonprofit organization based in Zurich, Switzerland. The Swiss Skeptics promote a scientific worldview, and, more specifically, critical thinking as a tool for making rational inferences about the world.

About the author

Marko Kovic is a founding member and, at the time of the publication of this paper, the sitting president of the Swiss Skeptics. By training, he is a social scientist and holds a PhD in communication science. He is also president of ZIPAR, the Zurich Institute of Public Affairs research, a nonprofit thinktank that conducts research in the areas of politics, communication, and science.

Recommended citation

Kovic, Marko (2016): “A generalized definition of critical thinking.” *Swiss Skeptics Discussion Paper Series* 1(1): 1-31.

Copyright

This document is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International license. You can and are encouraged to freely distribute this document and reference it in other works. For any questions about copyright, feel free to contact the author.

Abstract

The concept of critical thinking enjoys a near-universal positive connotation. Existing definitions of critical thinking, however, tend to be rather vague, and, as a consequence, they provide neither an accurate nor a precise understanding of critical thinking. In this paper, I propose to understand critical thinking as a *metacognitive skill applicable to the evaluation of truth claims*. Critical thinking as a metacognitive skill consists of three components: *Minimization of logical fallacies*, *minimization of cognitive biases*, and a *probabilistic epistemology*. Understood in this manner, critical thinking can improve the quality of our inferences about the world.

Contents

1	Introduction: The fuzzy concept of critical thinking	5
2	The problems with current definitions of critical thinking	6
2.1	Descriptive definitions	6
2.2	Proceduralist definitions	9
2.3	Outcome-oriented definitions	12
2.4	Domain-specific definitions	14
2.5	Summarized: Existing definitions describe a lot and define little	15
3	A generalized definition of critical thinking	16
3.1	Metacognition and truth claims	16
3.2	Minimization of logical fallacies	17
3.3	Minimization of cognitive biases	21
3.4	Probabilistic epistemology	22
4	Some implications	24
4.1	What's critical about critical thinking?	24
4.2	What is it good for?	24
4.3	How, exactly, does it work?	25
4.4	Critical thinking and skepticism	26
5	Conclusion	26

1 Introduction: The fuzzy concept of critical thinking

Are you a critical thinker? Chances are, your answer is Yes. Critical thinking is a concept that enjoys a strong and, quite possibly, universal positive connotation. We intuitively associate critical thinking with the idea of not simply believing things others tell us to believe. Instead, a critical thinker is one who uses his or her own faculties in order to find out the truth, autonomously and without succumbing to any naïveté in the process. Critical thinking, then, means a form of highly desirable thinking; critical thinking is how we *should* be thinking.

Even though the notion of critical thinking probably enjoys universal support, the specific nature of critical thinking is not very clear. There is some baseline agreement that critical thinking as «good» thinking means thinking that is autonomous, inquisitive and thorough, whereas regular, uncritical thinking means thinking that is automatic, unreflected and error-prone. Such a baseline agreement, however, is little more than an *ex negativo* definition¹, and it leaves much to be desired. With such a vague understanding of critical thinking, we might think that we *sort of* know what is meant, but really, we don't.

The fact that we deem critical thinking to be normatively desirable but struggle to properly define it gives it the air of an *essentially contested concept* [1]: A concept for which there is some basic agreement over its definition, but for which there is disagreement over its *complete* definition. A prototypical example of an essentially contested concept is the concept of democracy. Most people probably agree that democracy is a good thing², but it is very difficult to provide a clear and universally acceptable definition of what democracy should and should not entail.

It's very much possible that even such a vague, essentially contested understanding of critical thinking is of non-trivial practical use. By merely pointing out that not every kind of thinking is good and that, consequently, we should strive to reduce the amount of our uncritical thinking, we will probably induce some positive change to our thinking routines. Or, put more simply: By being told to think critically, we are being motivated to think about our thinking. Thinking about our thinking, in turn, results in more deliberate and careful thinking, which is desirable. This means that by the mere act of thinking about how we think, we enter a different mode of thinking: For the most part, thinking simply happens, but when we think about thinking, we are being more aware of our thinking.

¹Critical thinking is essentially understood as ¬uncritical («not uncritical») thinking. That statement is logically true, but definitorially empty if neither the properties of critical thinking nor the properties of uncritical thinking are explicitly defined.

²Or, put more precisely: A majority of people who live in democratic polities agree that democracy is a good way for a polity to be organized.

If a vague, fuzzy idea of critical thinking as «good thinking» is likely to have positive impact, should we, then, conclude that critical thinking is an essentially contested concept and simply regard the concept's vagueness to be its defining property? After all, we have just concluded that even such a vague and fuzzy notion of critical thinking is beneficial. Such a conclusion would be misguided. Critical thinking is *not* an essentially contested concept, but it is actually a concept that so far has not received enough attention. This means that critical thinking is a fuzzy concept not because opposing definitions have reached a state of conceptual equilibrium, but because there are actually only relatively few attempts to define critical thinking, and the existing attempts are, for the most part, preliminary and incomplete.

The goal of the present paper is, *first*, to present a brief critique of common definitions of critical thinking in section 2. Of course, I am not the first person to argue that there is something amiss with how we think of critical thinking [2]. However, just pointing out that there is a problem is not enough – that is why, *second*, I introduce a generalized definition of critical thinking in section 3. This newly introduced definition aims to overcome the shortcomings of the definitions that have been proposed so far. *Finally*, in section 4 I address some implications that arise given the generalized definition introduced in this paper.

2 The problems with current definitions of critical thinking

There are a number of attempts to define critical thinking, and those attempts do not stem from a single academic discipline or conceptual tradition; there is no one canon or canalization of ideas about critical thinking. This section aims to provide a rough typology of critical thinking definitions, so as to outline the main components of the current definitorial trajectories, and with them, the main shortcomings of the current definitorial trajectories.

2.1 Descriptive definitions

The dominant type of critical thinking definition is the *descriptive* one. This type of definition is easily recognizable because it consists of a description of some attributes of critical thinking instead of a definition proper. Or, to put it differently, they are merely adjectives that describe the adjective «critical». A widely used definition of critical thinking that was introduced by Ennis [3] is precisely such a descriptive one:

[R]easonable reflective thinking that is focused on deciding what to

think and do.

This definition introduces the descriptors «reasonable» and «reflective» in lieu of «critical», but in doing so, no *actual* definition is provided. The terms «reasonable» and «reflective» are themselves rather vague and thus in need of additional conceptual clarification.

An influential conceptual work by Lipman [4] proposes to understand critical thinking as a form of thinking that has three properties (*italics in the original*):

(1) it is *self-corrective* thinking; (2) it is thinking *with criteria*; and (3) it is thinking that is *sensitive to context*.

Criticizing this definition as a descriptive one might seem unfair, because the properties of self-correctiveness, criteria and context sensitivity might very well be elaborated in greater detail. And they are, but only, again, in a descriptive manner and not really a definitorial one. For example, the author describes the property of self-correctiveness as follows:

What has come to be known as scientific method is a distillation of the exploratory and self-corrective procedures employed by ordinary persons in everyday life. These same, self-corrective procedures are responsible for the emergence of logic. In turn, science and logic provide us with models that we can attempt to internalize and emulate in our thinking.

This clarification amounts to little more than additional descriptions. In addition, this specific argument is circular in nature: The source of the «scientific method» and logic are «self-corrective procedures» that people employ in everyday life. But science and logic, in turn, inform our thinking and provide us with «models» that we can internalize and «emulate» in our thinking. This reads a whole lot like our thinking gives rise to our thinking, which is, first, a circular way of thinking about the chronological, causal order of affairs, and it is, second, a purely descriptive way of thinking about critical thinking, and thus, it is no definition at all.

Another well-known definition of critical thinking is proposed by Paul [5]:

Critical thinking is disciplined, self-directed thinking that exemplifies the perfections of thinking appropriate to a particular mode or domain of thought.

In the above descriptive definition, two vague attributes are introduced, «disciplined» and «self-directed». The main definitorial component, however, is the

notion that critical thinking exemplifies the «perfections of thinking». Unfortunately, the author describes these «perfections» only in a fairly vague manner. «Perfections» of thinking are qualities such as clarity, precision, logicalness, depth, and «imperfections» of thinking are the opposites of the perfections, such as unclarity, imprecision, illogicalness, and superficiality. Of course, intuitively, we deem these «perfections» as more desirable than the «imperfections», but in definitorial terms, such fuzzy and emotionally loaded notions are of very little use. To make matters worse, in the above definition, the author is suggesting that critical thinking should only «exemplify» the desirable qualities that he calls «perfections». Ultimately, the above definition, then, proposes that critical thinking should be sort of something that is like some things that we think are good. There's a double vagueness in that definition that defeats the whole purpose of proposing a definition.

Another definition that has considerable impact on how critical thinking is understood is proposed by Facione [6]:

We understand critical thinking to be purposeful, self-regulatory judgment which results in interpretation, analysis, evaluation, and inference, as well as explanation of the evidential, conceptual, methodological, criteriological, or contextual considerations upon which that judgment is based. CT is essential as a tool of inquiry.

[...]

The ideal critical thinker is habitually inquisitive, well-informed, trustful of reason, open-minded, flexible, fair-minded in evaluation, honest in facing personal biases, prudent in making judgments, willing to reconsider, clear about issues, orderly in complex matters, diligent in seeking relevant information, reasonable in the selection of criteria, focused in inquiry, and persistent in seeking results which are as precise as the subject and the circumstances of inquiry permit.

This definition is the result of the deliberation of a range of experts, and it is a very vivid description of critical thinking. But it is, essentially, just that: A description that is a (fairly large) collection of attributes. The definitorial component of the definition *in toto* (which is even longer than the excerpt above) is that critical thinking is «judgement»; a statement that carries precious little information.

To summarize, the problem with descriptive definitions of critical thinking is that they, essentially, only introduce *attributes* of critical thinking, but they fail to provide an *actual* definition. Descriptive definitions of critical thinking fail to propose what critical thinking is, and instead, they describe what it looks like.

You can think of the problem with descriptive definitions with an analogy. Say you asked me «What is a chicken?», and I was to tell you «Chicken are usually not taller than 50cm and weigh between one and four kilograms; they have wings and feathers, but they are unable to fly». You would have learned some properties of chicken, but you'd be none the wiser about what chicken actually are – I described a chicken, I did not define it.

2.2 Proceduralist definitions

The «thinking» part of the term critical thinking is quite suggestive of the fact that critical thinking is related to some sort of *process*. Indeed, this can almost be regarded as a truism: Thinking does not describe a state, but a *change* of states; to think means to process information in some way. It is only logical, of course, that the process character of thinking applies to the critical kind as well.

The relevance of the process dimension of critical thinking features prominently in the *proceduralist* understanding of critical thinking. In the proceduralist understanding, critical thinking is described primarily in terms of the specific steps that are to be performed during the thinking process. This means that in the proceduralist understanding, critical thinking is defined in terms of ritualistic actions. I am not using the term «ritualistic» in a condescending way; I don't mean to imply that proceduralist understandings of critical thinking simply propose rituals, similar to, say, rituals in religious contexts. The term «ritualistic action» simply summarizes the main problem with proceduralist understandings of critical thinking: Those understandings are focused on describing the things one should do in order to think critically, but they do not actually define what critical thinking is.

One prominent streak of proceduralist definitions is the idea that critical thinking consists of «asking the right questions». This idea is prominently on display in a definition by Browne and Keeley [7, p. 2]:

Consequently, critical thinking, as we will use the term, refers to the following:

1. awareness of a set of interrelated critical questions;
2. ability to ask and answer critical questions at appropriate times; and the
3. desire to actively use the critical questions.

In the above definition, the ritualistic nature of the proceduralist understanding of critical thinking is on clear display: Knowing the right questions, asking

(and answering) the right questions, and desiring to use these right questions simply means performing a set of instructions over and over again. As I mentioned in the introduction, any mechanism that makes us reflect upon our thinking and thus makes our thinking more deliberate is likely to have a positive impact on the quality of our thinking. In this sense, proceduralist definitions such as the above probably do have some positive impact. However, since they are instructions on *how to perform* critical thinking, proceduralist definitions fail to actually define critical thinking. Obviously, defining critical thinking as an instruction for how to perform critical thinking is not a valid definition, but rather a somewhat recursive circular argument.

Another proceduralist definition of critical thinking that bears great similarity to «asking the right questions» is proposed by Willingham [8]:

[C]ritical thinking consists of seeing both sides of an issue, being open to new evidence that disconfirms your ideas, reasoning dispassionately, demanding that claims be backed by evidence, deducing and inferring conclusions from available facts, solving problems, and so forth.

Everything in that definition sounds very plausible, and it can be boiled down to something like «try to be objective». That is, of course, good advice in and of itself, but it is not sufficient as a definition. To be fair, the author of the above definition does actually present it as a definition in «layperson terms». However, it's not exactly clear what a «layperson» means when it comes to critical thinking, and, more importantly, the problem with the above definition is not its terminology, but its latent proceduralist premise. The «expert» definition that the author offers in contrast to the layperson one is the following (italics in the original):

From the cognitive scientists point of view, the mental activities that are typically called critical thinking are actually a subset of three types of thinking: reasoning, making judgments and decisions, and problem solving.

[...]

Critical reasoning, decision making, and problem solving—which, for brevity's sake, I will refer to as critical thinking—have three key features: effectiveness, novelty, and self-direction.

This definitions starts out as yet another proceduralist one – «reasoning», «making judgements and decisions», and «problem solving» – , and the proceduralist definition is then defined in descriptive terms, with a set of attributes

of critical thinking. This cannot serve as a proper definition of critical thinking, because it is merely a marriage of a proceduralist with a descriptive notion.

Another very popular definition of critical thinking is proposed by Scriven and Paul [9]:

Critical thinking is the intellectually disciplined process of actively and skillfully conceptualizing, applying, analyzing, synthesizing, and/or evaluating information gathered from, or generated by, observation, experience, reflection, reasoning, or communication, as a guide to belief and action.

This definition is, at first, appealing, because it encompasses a range of actions that make intuitive sense. But there are at least three problems with this definition. First, and obviously, it is a strongly proceduralist one, given that it is, basically, a list of actions that are to be performed in order to think critically. Second, the definition contains an «and / or» and an «or». This means that we don't know which of the activities are necessary conditions, and which are sufficient conditions. Third, this definition operates with the vague descriptor «intellectually disciplined» which does not add any definitorial clarity, but rather obscures the proposed definition.

Another definition with notable prevalence in the literature is proposed by Paul and Elder [10, p. 4]:

Critical thinking is the art of analyzing and evaluating thinking with a view to improving it.

This definition begins with the statement that critical thinking is an «art». What is that supposed to mean, exactly? Oftentimes, when something is called an art, it is juxtaposed with something that is clearly defined and measurable; «it's not a science, it's an art». By stating that critical thinking is an art, the authors of the above definition are preemptively defeating the purpose of their definitorial efforts – by calling critical thinking an art, critical thinking is described as a fuzzy, vague concept. The rest of the definition is rather standard proceduralist fare: Critical thinking is the «art» of analyzing and evaluating thinking.

In summary, proceduralist definitions of critical thinking are definitions that are *ritualistic* in nature: Proceduralist definitions describe different actions that, supposedly, lead to critical thinking by means of taking these actions. Proceduralist understandings of critical thinking, therefore, take the form of a set of *instructions*, which in and of themselves might very well have some positive impact. However, as definitions, they ultimately fail, since they are only an extension of descriptive definitions. Instead of proposing a set of vague attributes that

describe critical thinking, proceduralist definitions propose a set of actions that, together, lead to the process of thinking critically. If we go back to the example of the definition of a chicken, the problem of proceduralist definitions becomes obvious. If I defined a chicken as «An animal that is constantly walking around, but not flying, and it lays an egg about once a day», that definition would be very lacking. Basically, it's little more than a descriptive definition, for which the adjectives have been replaced with verbs. All the vagueness and uncertainty of the descriptive definitions remain.

2.3 Outcome-oriented definitions

While descriptive and proceduralist notions predominate in the existing attempts to define critical thinking, some definitions (additionally) define critical thinking in terms of the outcome of critical thinking. One such definition is proposed by Halpern [11, p. 8]:

Critical thinking is the use of those cognitive skills or strategies that increase the probability of a desirable outcome. It is used to describe thinking that is purposeful, reasoned, and goal directed – the kind of thinking involved in solving problems, formulating inferences, calculating likelihoods, and making decisions, when the thinker is using skills that are thoughtful and effective for the particular context and type of thinking task.

Large parts of the above definition are descriptive, but there is an outcome-oriented element mentioned twice. First, the author states that critical thinking is supposed to «increase the probability of a desirable outcome», and then, later on, that critical thinking means «using skills that are thoughtful and effective for the particular context and type of thinking task». Such an understanding of critical thinking is intuitively appealing. Of course, there is a latent expectation that critical thinking will have some sort of outcome – usually, when we engage in any sort of thinking, something results from it, be it a specific decision, or more generally, some sort of insight or inference. However, turning this factoid around and expecting that the outcome of critical thinking can be used to define critical thinking is a fallacy.

Suppose, for the sake of argument, that critical thinking produces one specific, particular type of «good» outcome, outcome X. Furthermore, let there be one specific, particular type of «bad» outcome, outcome Y. Furthermore, let us assume that critical thinking deterministically always produces outcome X. Would the presence of outcome X be enough to infer the presence of critical thinking? No, of course not – as a matter of fact, such an inference is always logically false. This

kind of fallacy is called affirming the consequent. Depicted as the relationship between premises and conclusion, the fallacy becomes obvious:

1. *critical thinking* \rightarrow *X*
2. *X*
3. *critical thinking*

If you are wondering why this fallacy is a fallacy, consider that constellation in another context. Let critical thinking be replaced by «to rain», and the outcome X by «wet street». Then you arrive at something like this:

1. *rain* \rightarrow *wet street*
2. *wet street*
3. *rain*

Now, of course, we know that when the street is wet, it's possible that it has been raining. But logically, we cannot state that this is always the case. Maybe someone has been washing their car on the street and that's what made it wet, or some kids in the neighborhood have been cooling off with a garden hose, or an underground pipe burst, and so forth.

An outcome-oriented definition of critical thinking is well-intentioned and intuitively sensible, but it is logically fallacious. That doesn't mean that we should not care about outcomes at all. Of course, it's very important to assess whether critical thinking produces outcomes that are superior to outcomes generated by uncritical modes of thinking. I would go even as far as saying that one of the main reasons why we even care about critical thinking in the first place is precisely the assumption that critical thinking as a way of «throughput» produces superior «output». Even though this question of outcomes is very important, it is not suitable for definitorial purposes. If you don't care about the above point about logical fallacies, there's another angle to the problem of outcome-oriented definitions. Defining critical thinking in terms of its outcomes means to engage in a peculiar form of circular reasoning. Critical thinking, so the assumption goes, produces a set of outcomes X. Therefore, critical thinking is defined as X. However, if we then try to define X, we define it as critical thinking. This means that critical thinking is defined in terms of X, and X is defined in terms of critical thinking. Obviously, there is information missing here.

2.4 Domain-specific definitions

The final group of definitions of critical thinking is what I call domain-specific definitions. This type of understanding is not very prevalent, but it warrants attention. A domain-specific understanding of critical thinking might not, in and of itself, propose a full definition of critical thinking, but the component of domain-specificity is far-reaching enough so as to have a fundamental impact on any kind of definition of critical thinking. Essentially, the argument of domain-specificity is that critical thinking is not a universally applicable skill, but rather a domain-specific skill: Applying critical thinking in context A and for task 1 is not the same as applying critical thinking in context B and for task 2.

In an influential paper, Willingham [8] argues that critical thinking is inextricably linked to the specific domains in which it is to be applied:

The processes of thinking are intertwined with the content of thought (that is, domain knowledge). Thus, if you remind a student to “look at an issue from multiple perspectives” often enough, he will learn that he ought to do so, but if he doesn’t know much about an issue, he can’t think about it from multiple perspectives. You can teach students maxims about how they ought to think, but without background knowledge and practice, they probably will not be able to implement the advice they memorize.

This idea is not only conceptual in nature, but seems to be supported by some experimental evidence as well [12]. What, then, are we to make of this? Is critical thinking not a skill that can be used independent of specific contexts? I think the argument of domain-specificity is somewhat of a misunderstanding. In the definition quoted above, critical thinking is implicitly understood in the outcome-oriented sense. When students³ have to think about some issue that they are familiar with, they can, so the argument goes, apply critical thinking, because they know the subject matter well enough so that they can think about it in non-routine ways. When, on the other hand, it comes to issues that they are unfamiliar with, students cannot apply critical thinking, because they don’t really know what to apply it to. This idea implicitly states that critical thinking means obtaining an outcome that fits some unspecified criteria. That notion is wrong. Of course, it is very well possible that people in general, not just students, will arrive at different outcomes, depending on how well they know some area or topic. For example, when presented with an engineering problem, an engineer will almost certainly arrive at different outcomes than a non-engineer. However, the outcome really doesn’t matter. When applying critical thinking, a non-engineer

³I’m talking about students here simply to stick with the example from the above quote.

will, ideally, come to the conclusion that he or she lacks the necessary knowledge in order to come to a useful or usefully certain solution to the problem at hand – critical thinking will help evaluate one's own thinking, and it will help quantify the uncertainty of one's own thinking⁴.

As for the empirical study cited above that, apparently, supports the domain-specific notion of critical thinking: The study doesn't actually measure critical thinking, but the susceptibility to the so-called "belief bias" [13]. The belief bias refers to the fact that we are more likely to accept conclusions as true that seem subjectively right to us, regardless of the quality of the argument that leads to those conclusions. The study cited above essentially states that belief bias is heightened in domains where test subjects were relatively less knowledgeable. Of course, that doesn't mean that critical thinking is domain specific – but rather that critical thinking is especially necessary when making inferences in low-information settings. In that sense, domain-specific definitions of critical thinking are welcome, because they, inadvertently, demonstrate *why* critical thinking is necessary in the first place – the belief bias is one of many reasons why our routinized way of thinking can lead us astray.

2.5 Summarized: Existing definitions describe a lot and define little

In the preceding sections, I have argued at some length that and how existing definitions of critical thinking fall short of being actual definitions. Does that mean we should completely discard everything the existing definitions propose? No: *ceteris paribus*, the existing definitions help understand critical thinking better than a total absence of definitions. This means that I don't believe that the existing definitions point in the wrong direction; the net amount of understanding critical thinking is higher than zero with the existing definitions. The problem with the definitions that have been proposed so far, then, is not that they are completely wrong *per se*, but rather that they avoid the crucial step of precisely and accurately defining critical thinking, and instead leapfrog to various types of descriptions of critical thinking.

Describing critical thinking in different ways – the purely descriptive, the proceduralist and the outcome-oriented way – is very welcome. But if we do not have a definitorial foundation upon which these descriptions are based on, they are bound to remain vague and, ultimately, of very limited use. For that reason, I introduce a generalized definition of critical thinking.

⁴Don't take this as a definition of critical thinking! The definition is following in section 3

3 A generalized definition of critical thinking

I propose the following as a *generalized definition of critical thinking*:

Critical thinking is a metacognitive skill applicable to the evaluation of truth claims. Critical thinking consists of three components: Minimization of logical fallacies, minimization of cognitive biases, and a probabilistic epistemology.

As far as definitions go, this one is not really ideal, because it consists of two parts, the actual definition, and an additional explanation for the definition. But I believe this definition is still concise enough so as to be of some generalized use. There are five elements in this definition that warrant some explanation: Metacognition; truth claims; logic; cognitive biases; probabilistic epistemology. I delve into these elements in the following subsections.

First, however, let me address another aspect of this definition: What does it mean for the above definition to be «generalized»? The descriptor «generalized» is there to make explicit the fact that the definition I introduce is intended to be *universal*: Critical thinking as introduced here and elaborated in the subsections below is a metacognitive skill that transcends empirical contexts and actors. I propose critical thinking as a generalized skill because, first, I regard human cognition as essentially universal. Of course, no two people think exactly alike, but the basic facts of human biology, to borrow a term from Searle [14], are much the same for everyone. Second, the principle of generalization also expresses a firm opposition to the misguided notion of critical thinking as a domain-specific skill (cf. subsection 2.4 for a critique of that notion): Critical thinking is a skill that can be applied in all manners of domains, independent of contextual factors.

3.1 Metacognition and truth claims

The term *metacognition* has been around for close to forty years [15], even though what it describes has been around for as long as humans have been engaging in philosophy. Metacognition simply means directing one's thinking at one's thinking – when we think about the way we think, we engage in a metacognitive process. Of course, most of us are not professional philosophers, but metacognition comes to us fairly naturally, mostly without us actively seeking to engage in metacognition. For example, if you have ever told someone something along the lines of «What were you thinking?», you were engaging in a metacognitive process, and you tried to elicit the same in someone else⁵.

⁵This understanding of «What were you thinking?» is probably a little bit too charitable. Usually, when we address someone like this, we are not genuinely curious about why someone

Metacognition, then, is not something special, nor is it especially rare. Critical thinking as a metacognitive skill, however, is a skill intended for one primary purpose: The evaluation of *truth claims*. Truth claims are propositional utterances that posit that the validity claim in question is true. More specifically, the truth claims that I am referring to are sometimes described as representative [16] or constative [17] speech acts: The claim that an uttered proposition about the objective world is true. Some people tend to recoil a bit when there is talk about things like «objective world» and «truth», simply because not everyone shares the same ontology, the understanding of the nature of being. There are two positions one can take when it comes to potential ontological worries. The first one is mostly practical: You don't have to care about ontology at all, because what matters is simply that someone makes a truth claim. What we are subsequently interested in is not the claimants general beliefs about the world, but rather the degree to which the specific posited truth claim can be regarded as true or as false. But, of course, this is an *ad hoc* approach, and I advise against it. In order to use critical thinking as a metacognitive skill in a consistent manner, you have to have some ontological foundation. A foundation that is appropriate here is ontological *realism*: The physical world exists independently of our representations of it [18, 19], and the degree to which a truth claim can be regarded as true or false depends on how well the truth claim corresponds to reality. If you reject realism, now is probably a good time to stop reading this paper⁶.

Critical thinking, as defined above, is a metacognitive skill applicable to the evaluation of truth claims. That does *not* mean, however, that critical thinking will directly tell you whether a truth claim is true or not. Critical thinking is a *metacognitive* skill, which means that it is used to think about the way you are evaluating truth claims. Critical thinking, then, is supposed to make the way you are trying to evaluate whether a truth claim is true or not less error-prone and, consequently, more reliable and more valid. In other words: Critical thinking is not going to give you answers, but it will increase the probability that you arrive at correct answers.

3.2 Minimization of logical fallacies

Being logical is, probably, seen as a positive quality in a similarly quasi-universal way as is critical thinking. Thankfully, the concept of logic is much less vague

thinks the way they do, but rather, we are making an accusation, or expressing frustration and anger, and so forth.

⁶Just to prevent potential confusion: Realism does *not* mean that we have perfect access to reality, nor that we can perfectly evaluate whether a truth claim corresponds to reality. Critical thinking, as I argue in subsection 3.4, is precisely intended to make these epistemological uncertainties explicit and quantifiable.

and fuzzy than the concept of critical thinking, so it's much easier to specify what is meant by being logical.

In essence, a logical statement means an argument for which the conclusion follows from the premises. Such as this one:

1. $day \rightarrow \neg night$
2. day
3. $\neg night$

This is a trivial argument which everyone will immediately understand to be true, without any scrutiny of logic: If it's day, it's not night. It's day; therefore, it's not night. This argument is logically true, because the conclusion follows from the premises. But maybe you're wondering about the empirical nature of that argument. Sure, when it's daytime at location A, it's not nighttime at location A. But there are places in the world, and on many other planets in the known universe, where it is nighttime while it's daytime at location A. You are perfectly correct to interject in that manner – but you are referring to the empirical content of the argument, and not to its *formal logic*. Formally, the argument is true. Another argument that is formally true but empirically questionable is the following one:

1. $love \rightarrow God$
2. $love$
3. God

If there is love in the world, there is a God. There is love in the world, and therefore, there is a God. This is, in a nutshell, the reasoning behind so-called «proofs» for the existence of a God, such as the famous one by Gödel [20]. The above argument is formally true, but you might wonder whether its informal, empirical content is true as well; you might question the premises.

The two above examples show that there are, generally, two types of logical consistency that matter, formal and informal. Consequently, we are prone to formal as well as informal *logical fallacies* [21, 22, 23, 24, 25]. Critical thinking entails arguing in a logically consistent manner both formally and informally. In doing so, we are minimizing logical fallacies.

There are many possible types of logical fallacies; far too many to list them all here. Instead of exhaustively discussing all logical fallacies, it's more important to keep in mind some general properties of logical fallacies. *First*, formal fallacies supersede informal fallacies with regard to their impact on the conclusion of the

Table 1: Impact of formal and informal truth on overall argument truth.

formal truth	informal truth	overall argument truth
false	false	false
true	false	false
true	true	true

argument. Every argument contains both formal and informal logic. If there is a formal fallacy present in the argument, the argument is always false, regardless of the quality of its informal component. If, however, an argument is formally true, it can either be informally true or informally false. The relationship between formal and informal logic is summarized in Table 1.

Back in subsection 2.3, we have already encountered an argument that is formally fallacious. It was this one:

1. *rain* \rightarrow *wet street*

2. *wet street*

3. *rain*

This argument is formally false, because you are affirming the consequent. Because this argument is formally false, its informal, empirical content is automatically false as well. Informally, meaning empirically, we know that rain can make the street wet; this informal knowledge is what we express formally with *rain* \rightarrow *wet street*. But we also know that many other things can make the street wet – this additional piece of informal knowledge is expressed by the fact that we *don't* posit the additional premise *wet street* \rightarrow *rain*.

Above, the «God proof» argument has also been introduced. That argument is formally true, but informally false⁷. There are many more examples of purely informal fallacies. It is difficult to quantify which informal fallacies are most prominent and have the greatest impact on the quality of our evaluation of truth claims. Some informal fallacies that can be encountered frequently are the *ad*

⁷Or, at the very least, the informal status of the premise *love* \rightarrow *God* is uncertain; there is, as far as I can tell, no empirically sound reason to accept that premise as true.

hominem fallacy⁸, the argument from popularity⁹, the appeal to tradition¹⁰, the argument to moderation¹¹, or the appeal to authority¹².

The *second* important aspect of logical fallacies is that, for the purpose of critical thinking, an argument can be regarded as true only when it is true *both* formally and informally, as is depicted in the bottom row in Table 1. As I argue in subsection 3.1, critical thinking applies to the evaluation of truth claims, and truth claims are claims about reality. That is why informal fallacies matter as much as formal fallacies when it comes to critical thinking: We do not only care whether conclusions follow properly from the premises on a formal level, but also, additionally, whether the content of the premises is true, meaning, whether the content of the premises itself corresponds to reality.

The *third* important aspect of logical fallacies is that they are insidious. Of course, logical fallacies are not actors who act insidiously; fallacies are not insidious in a literal way. They are insidious in a figurative way: It's very easy to unambiguously detect and describe logical fallacies in theory, but it's not easy to do so in practice, when we are evaluating truth claims. Real-world truth claims are usually not isolated as arguments with clearly demarcated premises and conclusions, but rather, they are embedded in a greater narrative, with a lot of verbiage and «storytelling» that serves the dual purpose of explaining one's arguments in greater detail as well as persuading others of the truth of one's arguments.

⁸By attacking some trait of a person, you try to prove the person's argument as wrong. Example: «Steve says that vaccines do not cause autism. However, Steve is wrong, because he is arrogant.» The informally false premise here is the assertion that arrogance makes someone wrong on the issue of vaccines and autism.

⁹You try to prove the truth of a claim by pointing out its popularity. Example: «Homeopathy works beyond the placebo, because millions of people use homeopathy regularly.» The premise that the mere popularity of therapeutic procedures is indicative of their efficacy and effectiveness is false.

¹⁰You try to prove a claim as true by arguing that something has been done for a long time. Example: «Witches exist, because we have been burning witches for hundreds of years». The premise that the existence of a tradition is proof of the truth of the belief of that tradition is false.

¹¹If there are two opposing opinions, you argue that the best solution lies in the middle. Example: «Marc says that astrology works. Jennifer says that astrology doesn't work. The truth is that astrology works sometimes, but not always.» The premise that the middle ground is always the best solution is false.

¹²You try to prove a claim by pointing out that some person or entity of authority shares your opinion. Example: «The earth is a flat disc. John says so, and John has a PhD.» The premise that some sort of authority in the form of, for example, expertise in a scientific field automatically proves a claim is false. Even if, in this example, John was an expert in a field relevant to the claim at hand, simply referring to his status as an expert without evaluating his concrete arguments is fallacious.

3.3 Minimization of cognitive biases

When we are making inferences about the world, we are wholly dependent on our cognitive faculties to do so—our brain is where our thinking takes place, and the amount and quality of our thinking is enabled as well as limited by our biology. The course of our evolutionary path happens to have been one where we ended up with considerable brain power; after all, no known animal in the universe comes even close to our level of reasoning about the world (and, as a consequence, to our capability for influencing the world). But we humans are just that, products of nature, who have evolved not to be perfect, but to be good enough.

«Good enough» is a very apt description of our cognitive abilities. We are usually able to handle very complex thoughts in a consistent and systematic manner, and we arrive at good enough inferences about the world most of the time. We are, in a nutshell, able to make inferences that correspond to reality well. However, at the same time, much, if not most of our thinking is not very deliberate and reflected, but rather automated. Every day, we make a myriad of *ad hoc* inferences without even actively thinking about them. Instead of carefully and explicitly thinking through many of our inferences, we rely on certain shortcuts, so-called *cognitive heuristics*. These heuristics are rough inferential rules of thumb that greatly reduce the cognitive cost of making all the inferences that we make while delivering results that are usually good enough. Our thinking, thus, can be described in terms of a dual system [26, 27, 28]: One system, or mode, of thinking is slow, deliberate, and careful, while the other is fast, automated and efficient.

Our cognitive heuristics serve us well, for the most part. They allow us to reliably maneuver through many situations of routinized inference making, which saves cognitive resources for high-stakes situations where we really want to take our time for more careful and more explicit thinking. Cognitive heuristics, then, are a very crucial trait of human cognition. However, even though cognitive heuristics are very useful in general, they can also be detrimental to the quality of our inferences in situations where a good enough inference can lead to conclusions that are substantially, if not categorically false. That is why cognitive heuristics are sometimes labeled as cognitive *biases* [29].

There are far too many cognitive biases to list them all here. Furthermore, similar to the situation with logical fallacies, it's difficult to quantify which biases have the greatest negative impact on our inferences. Some well-known cognitive biases are the confirmation bias [30], anchoring [31], the availability heuristic [32], overconfidence [33], *status quo* bias [34], the ingroup bias [35], loss aversion [36], the halo effect [37], and the fundamental attribution error [38]. But these biases comprise only a small part of all the ways in which our cognition falls short. Cognitive biases are essentially «baked into» our cognitive hardware, our brains; susceptibility to cognitive biases is a universal trait of human cognition.

But, fortunately, our cognitive biases are not just completely random, but rather, they are systematic patterns of flawed cognition. This systematic nature of our cognitive biases is sometimes expressed with the adage that we humans are *predictably irrational* [39]: Precisely because the heuristics of our cognition are a universal trait, we can get a grasp on how and in which situations cognitive biases can have an impact on our inferences. Furthermore, by having a grasp of the systematic nature of cognitive biases, we can engage in what is sometimes referred to as *debiasing* [40, 41]: By way of learning about cognitive biases, we can reduce the probability of falling prey to them. In the context of critical thinking, the minimization of cognitive biases can be understood as a constant exercise in debiasing. When evaluating truth claims, it is necessary to evaluate our inferences in light of cognitive biases. It is, therefore, necessary to make sure that the inferences we make are not only logically sound, but that they are also truly objective and not influenced by cognitive biases.

3.4 Probabilistic epistemology

In subsection 3.1, the issue of ontology was briefly touched upon: The whole idea of critical thinking makes sense when and only when one accepts realism as a basic model of reality. Within a realist framework, so the argument above goes, the degree to which a truth claim can be regarded as true or false depends on the degree to which the truth claim in question corresponds to reality. In that brief description, an implicit, but very crucial *epistemological* statement was made: Our beliefs about the validity of truth claims should not be deterministic, but *probabilistic* in nature.

The idea of a probabilistic epistemology goes against our intuitive reasoning about the world. For the most part, we are seeking to think about our inferences in deterministic terms, absent of any uncertainty – we are just very fond of having clear, unambiguous answers. The desire for determinism does not mean categorically faulty reasoning, because most people in most contexts actually do have an implicitly probabilistic epistemic *modus operandi*, but they only fail to make it explicit. An area where such an implicit probabilism in a situation of ostensible determinism can be observed is scientific research. Epistemologically, many scientists will identify as *falsificationists* in the sense of Popper [42]. The falsificationist view of science is important, because it addresses the so-called demarcation problem [43], which is the question of how to distinguish science from non-science. In the falsificationist view, an important trait of scientific reasoning is that, epistemically, truth claims – often called «hypotheses» – cannot actually be verified. Instead, it's only possible to *falsify* truth claims. Falsify here means to demonstrate that the truth claim in question is false; in everyday language, «to falsify» has a very different meaning. The falsificationist epistemology is, in prin-

ciple, very sound, and the idea of falsificationism still plays a prominent role in how many scientists understand the rules of their craft. However, falsificationism paints a rather deterministic picture. Say, for example, that you performed ten independent experiments in order to test a truth claim (a hypothesis). The first nine experiments yielded results that support the truth claim. The tenth experiment, however, yielded results that falsify the hypothesis. In a very strict falsificationist manner, the truth claim would have been falsified with the tenth experiment, thus making all previous experiments obsolete. Of course, in reality, that is not how science works: If nine out of ten experiments yielded results in favor of a truth claim, researchers would interpret the overall results probabilistically as substantive support for the truth claim.

The idea of an explicitly probabilistic epistemology instead of a deterministic one is sometimes referred to as Bayesian epistemology [44, 45]. Thomas Bayes was a statistician (and minister) who became famous for his work on conditional probability. Today, Bayesianism is used as an umbrella term for an epistemology that is explicitly probabilistic, and where one is supposed to explicitly state so-called prior and posterior beliefs. In the above example with ten experiments, the posterior belief after nine experiments that supported the truth claim would have been that, given the data, the probability that the truth claim was true was around 100%¹³. After the tenth experiment that yielded results that didn't support the truth claim, the prior belief (the results of the first nine experiments) would have been updated with the new information, and the new posterior belief would be that the probability that the truth claim is indeed true was around 90% percent. Of course, this is an extremely simplified example¹⁴. But the principle of probabilistic epistemology has a number of concrete applications, one of them systematic reviews and meta-analyses [46]. The goal of systematic reviews is to assess the overall probability of a truth claim given a number of studies about that truth claim. It is usually clear from the very outset that some studies do not support the truth claim in question, but that does not mean that the truth claim is outright falsified. Rather, all available evidence is assessed so as to arrive at a probabilistic assessment of the truth of the truth claim.

¹³Note that, in probabilistic terms, you can assign a probability of 100% without claiming absolute truth. A probability of 100% merely means that, given the information at hand, the probability for a truth claim being true is 100%. However, if you received additional information that did not support the truth claim, you would still act rationally and adjust the probability accordingly. Even though I don't believe that there are too many real-life scenarios where one can assign a probability of 100%, even if one did, that would mean no declaration of absolute and unquestionable truth, but rather something like a temporary lack of uncertainty.

¹⁴For example, for every one of those experiments, one would have to apply critical thinking in order to get some sense about the experiment's validity and reliability.

4 Some implications

In section 3, I have introduced the definition of critical thinking as a metacognitive skill applicable to the evaluation of truth claims, and subsequently described its components in some detail. In this sections, I address some implications of that understanding of critical thinking. Not *all* implications of the understanding of critical thinking as it was introduced above are addressed, but, I believe, some of the more relevant ones.

4.1 What's critical about critical thinking?

Up to this point, I have glossed over the meaning of «critical» in critical thinking. In our everyday understanding, to be critical of something means to have a negative attitude towards something, being against something. If that meaning of critical extended to critical thinking, then to think critically would amount to distrusting something, being against something, and looking for ways how the object of one's negative attitude might be wrong. That is *not* what the critical in critical thinking means.

The meaning of critical in critical thinking is better understood as *critique* than as criticism. Understood in that sense, critical thinking is thinking that represents a thorough and justifiable assessment, and not simply an expression of disagreement. Of course, the term critical has more than one meaning and more than one connotation, and the understanding proposed here is not necessarily the only true meaning of critical thinking, but it is the appropriate one when it comes to critical thinking as defined in this paper.

4.2 What is it good for?

The term critical thinking has one connotation that I wholeheartedly agree with: It's critical in the sense that it is necessary and important. But how important is critical thinking, really? That is a question that is surprisingly difficult to answer in a quantifiable way.

If critical thinking has the goal of improving the quality of our inferences through the minimization of logical fallacies, the minimization of cognitive biases, and a probabilistic epistemology, then, in a sense, almost all instances of faulty inferences are actually inferences that were made with a lack of critical thinking. And since we make inferences mostly not just for the sake of making inferences, but in order to base real-world actions on them, bad inferences can lead to bad outcomes. These bad outcomes as results of bad inferences span, in principle, the whole repertoire of human decision-making.

Such a statement might very well read as a sweeping generalization, and there is indeed some danger of confirmation bias here. If one believes that critical thinking is very important, one might be tempted to ascribe all of the world's ills to a lack of critical thinking. That is not exactly what I believe, but the overall point stands: It is desirable to try and quantify the impact of critical thinking, or the lack thereof, in the real world. That project is, obviously, beyond the scope of the current paper.

4.3 How, exactly, does it work?

If you think that critical thinking as defined in this paper makes conceptual sense, then you might wonder how exactly to translate this theory of critical thinking into actual critical thinking. In subsection 2.2, I have criticized proceduralist definitions of critical thinking as being definitorially insufficient. However, proceduralist notions of critical thinking are valuable because they try to tackle the very important aspect of how to perform the action of critical thinking.

There are two basic ways of translating the concept of critical thinking into real-world inference making. The *first* approach is to learn about critical thinking in a conceptual manner, such as by reading the paper at present. This might, *prima facie*, sound counterproductive: Is not the goal to actually apply critical thinking and not just to talk about it theoretically? Isn't conceptual thinking about critical thinking a waste of time, when we should be applying critical thinking? As a matter of fact, I believe that learning and thinking about critical thinking from a conceptual, theoretical perspective is very beneficial; by learning about critical thinking, we increase the probability that we will apply critical thinking. If we are (continually) exposed to arguments about how and why critical thinking matters, then the idea of critical thinking will gradually gain salience in our minds. At some point, this amounts to an effect of conditioning: By being aware of the importance of critical thinking *in principle*, that awareness will spill over into our inference making *in practice*.

The *second* approach to applying critical thinking to real-world inference making is more paternalistic in nature: We can increase the probability that other people will apply critical thinking by exploiting their cognitive biases. This is sometimes referred to as nudging [47, 48]. Of course, nudging people into critical thinking is not a universal solution, but in some contexts, it might be more effective than exposing people to conceptual information about critical thinking. One obvious example is education: Young students in schools might have trouble understanding the conceptual minutiae of critical thinking, but they might be receptive to nudges that motivate them to apply critical thinking.

These two approaches are outlined in a very superficial manner in the present paper. A more thorough example of implementation strategies for different

audiences in different contexts is a project that is beyond the scope of the present paper.

4.4 Critical thinking and skepticism

Philosophical skepticism, in its traditional sense, describes a range of epistemological positions that stress the fundamental uncertainty of knowledge and of knowing [49, 50]. Philosophical skepticism roughly holds that absolute knowledge is impossible, because the correspondence between our belief states and the states of the outer world can never be conclusively ascertained.

The everyday meaning of skepticism is very different from this philosophical position. When we say that we are skeptical of something or someone, we are usually expressing concern and disbelief. In our everyday use of the term, then, we are basically saying that something doesn't feel right to us, but we can't quite put the finger on it. Or, to put it differently: We opt not to believe something, without necessarily having rational reasons to do so.

Skepticism in the philosophical sense is preferable to skepticism in the *ad hoc* sense. However, even philosophical skepticism is in and of itself of only limited use. Questioning the nature of knowledge is, of course, a fine exercise, but it does not necessarily help us with the challenge at hand – how do we make our inferences more reliable and valid? This is precisely where critical thinking comes into play. If we understand skepticism as the application of critical thinking, then skepticism transcends both its misguided *ad hoc* meaning and its almost exclusively theoretical original philosophical meaning. In doing so, skepticism does not lose its basic philosophical idea of the uncertainty of knowledge. The opposite is the case: The uncertainty of knowledge is actually made explicit as the component of a probabilistic epistemology in critical thinking. Let it be proposed, then, to henceforth understand skepticism as the application of critical thinking.

5 Conclusion

The main goal of the present paper is to introduce a definition of critical thinking that overcomes the shortcomings of current definitions. To that end, I have proposed the following definition:

Critical thinking is a metacognitive skill applicable to the evaluation of truth claims. Critical thinking consists of three components: Minimization of logical fallacies, minimization of cognitive biases, and a probabilistic epistemology.

This definition has a number of advantages. It explains, first, what critical thinking is and where it is to be applied, and second, it explains what exactly critical thinking consists of. Critical thinking understood in this manner can help to make our inferences about the world more reliable and more valid.

The arguments put forward in the present paper are, for the most part, *conceptual* in nature. Therefore, the present paper can only be regarded as a foundation or a blueprint for the necessary steps beyond the conceptual discussion: The challenge of translating critical thinking as a concept into critical thinking as a practice. I have briefly touched upon two strategies for increasing the practical use of critical thinking, disseminating and discussing the concept of critical thinking, and «nudging» people into actually thinking critically. Both of these strategies could lead to an increase in the use of critical thinking, by increasing awareness of critical thinking as well as by inducing people to think more critically by introducing targeted incentives into people's choice architecture. Specific measures for realizing those general strategies, however, need to be fleshed out in future research¹⁵.

¹⁵When I say research, I mean something like applied research. Of course, measures for increasing the use of critical thinking need to be devised in a problem-oriented manner. But doing so absent of systematic conceptualization, execution and evaluation bears the risk that those measures slip away from an evidence-based perspective and into a feel-good exercise.

References

- [1] Walter B. Gallie. “Essentially Contested Concepts”. In: *Proceedings of the Aristotelian Society* 56 (1955), pp. 167–198 (cit. on p. 5).
- [2] Jennifer Wilson Mulnix. “Thinking Critically about Critical Thinking”. In: *Educational Philosophy and Theory* 44.5 (2012), pp. 464–479. DOI: 10.1111/j.1469-5812.2010.00673.x (cit. on p. 6).
- [3] Robert H. Ennis. “A taxonomy of critical thinking dispositions and abilities”. In: *Teaching thinking skills: Theory and practice*. Ed. by J. B. Baron and R. J. Sternberg. Series of books in psychology. New York, NY, US: W H Freeman/Times Books/ Henry Holt & Co, 1987, pp. 9–26. ISBN: 0-7167-1789-1 (Hardcover); 0-7167-1791-3 (Paperback) (cit. on p. 6).
- [4] Matthew Lipman. “Critical Thinking: What can it be?” In: *Analytic Teaching* 8.1 (1987) (cit. on p. 7).
- [5] Richard Paul. “Critical thinking: What, why, and how”. In: *New Directions for Community Colleges* 1992.77 (1992), pp. 3–24. DOI: 10.1002/cc.36819927703 (cit. on p. 7).
- [6] Peter A. Facione. “Critical Thinking: A Statement of Expert Consensus for Purposes of Educational Assessment and Instruction. Research Findings and Recommendations.” In: (1990) (cit. on p. 8).
- [7] M. Neil Browne and Stuart M. Keeley. *Asking the Right Questions: A Guide to Critical Thinking*. Pearson Education, 2014. ISBN: 978-0-321-90795-0 (cit. on p. 9).
- [8] Daniel T. Willingham. “Critical Thinking: Why Is It So Hard to Teach?” In: *Arts Education Policy Review* 109.4 (2008), pp. 21–32. DOI: 10.3200/AEPR.109.4.21-32 (cit. on pp. 10, 14).
- [9] Michael Scriven and Richard Paul. “Critical thinking as defined by the National Council for Excellence in Critical Thinking”. In: Rohnert Park, Ca, 1987 (cit. on p. 11).
- [10] Richard Paul and Linda Elder. *The Miniature Guide to Critical Thinking: Concepts & Tools*. Foundation Critical Thinking, 2001. ISBN: 978-0-944583-10-4 (cit. on p. 11).
- [11] Diane F. Halpern. *Thought and Knowledge: An Introduction to Critical Thinking*. Psychology Press, 2013. ISBN: 978-1-134-63786-7 (cit. on p. 12).
- [12] Walter C. Sá, Richard F. West, and Keith E. Stanovich. “The domain specificity and generality of belief bias: Searching for a generalizable critical thinking skill”. In: *Journal of Educational Psychology* 91.3 (1999), pp. 497–510. DOI: 10.1037/0022-0663.91.3.497 (cit. on p. 14).

- [13] Jonathan St B. T. Evans and Jodie Curtis-Holmes. “Rapid responding increases belief bias: Evidence for the dual-process theory of reasoning”. In: *Thinking & Reasoning* 11.4 (2005), pp. 382–389. DOI: 10.1080/13546780542000005 (cit. on p. 15).
- [14] John R. Searle. *Making the social world*. Oxford: Oxford University Press, 2010. ISBN: 978-0-19-957691-3 (cit. on p. 16).
- [15] John H. Flavell. “Metacognition and cognitive monitoring: A new area of cognitive–developmental inquiry”. In: *American Psychologist* 34.10 (1979), pp. 906–911. DOI: 10.1037/0003-066X.34.10.906 (cit. on p. 16).
- [16] John R. Searle. “A classification of illocutionary acts”. In: *Language in Society* 5.1 (1976), pp. 1–23 (cit. on p. 17).
- [17] Jürgen Habermas. “Some Distinctions in Universal Pragmatics: A Working Paper”. In: *Theory and Society* 3.2 (1976), pp. 155–167 (cit. on p. 17).
- [18] John R. Searle. *The Construction of Social Reality*. Free Press, 1997. ISBN: 978-0-684-83179-4 (cit. on p. 17).
- [19] John Searle. *Seeing Things as They are: A Theory of Perception*. Oxford ; New York: Oxford University Press, 2015. ISBN: 978-0-19-938515-7 (cit. on p. 17).
- [20] C. Anthony Anderson and Michael Gettings. *Gödel’s ontological proof revisited*. Association for Symbolic Logic, 1996. ISBN: 978-3-540-61434-0 (cit. on p. 18).
- [21] Ralph H. Johnson. “The Relation between Formal and Informal Logic”. In: *Argumentation* 13.3 (1999), pp. 265–274. DOI: 10.1023/A:1007789101256 (cit. on p. 18).
- [22] Madsen Pirie. *How to Win Every Argument: The Use and Abuse of Logic*. Continuum, 2006. ISBN: 978-0-8264-9006-3 (cit. on p. 18).
- [23] Douglas Walton. “Why Fallacies Appear to be Better Arguments Than They Are”. In: *Informal Logic* 30.2 (2010), pp. 159–184 (cit. on p. 18).
- [24] S. Morris Engel. *With Good Reason: An Introduction to Informal Fallacies*. Bedford/St. Martin’s, 1999. ISBN: 978-0-312-15758-6 (cit. on p. 18).
- [25] T. Edward Damer. *Attacking Faulty Reasoning: A Practical Guide to Fallacy-Free Arguments*. Cengage Learning, 2008. ISBN: 978-0-495-09506-4 (cit. on p. 18).
- [26] Keith Frankish. “Dual-Process and Dual-System Theories of Reasoning”. In: *Philosophy Compass* 5.10 (2010), pp. 914–926. DOI: 10.1111/j.1747-9991.2010.00330.x (cit. on p. 21).

- [27] Jonathan St B. T. Evans. “Dual-Processing Accounts of Reasoning, Judgment, and Social Cognition”. In: *Annual Review of Psychology* 59.1 (2008), pp. 255–278. DOI: 10.1146/annurev.psych.59.103006.093629 (cit. on p. 21).
- [28] Daniel Kahneman. *Thinking, Fast and Slow*. Penguin, 2012. ISBN: 978-0-14-103357-0 (cit. on p. 21).
- [29] Amos Tversky and Daniel Kahneman. “Judgment under Uncertainty: Heuristics and Biases”. In: *Science* 185.4157 (1974), pp. 1124–1131. DOI: 10.1126/science.185.4157.1124 (cit. on p. 21).
- [30] Raymond S. Nickerson. “Confirmation bias: A ubiquitous phenomenon in many guises”. In: *Review of General Psychology* 2.2 (1998), pp. 175–220. DOI: 10.1037/1089-2680.2.2.175 (cit. on p. 21).
- [31] Adrian Furnham and Hua Chu Boo. “A literature review of the anchoring effect”. In: *The Journal of Socio-Economics* 40.1 (2011), pp. 35–42. DOI: 10.1016/j.socec.2010.10.008 (cit. on p. 21).
- [32] Amos Tversky and Daniel Kahneman. “Availability: A heuristic for judging frequency and probability”. In: *Cognitive Psychology* 5.2 (1973), pp. 207–232. DOI: 10.1016/0010-0285(73)90033-9 (cit. on p. 21).
- [33] Don A. Moore and Paul J. Healy. “The trouble with overconfidence”. In: *Psychological Review* 115.2 (2008), pp. 502–517. DOI: 10.1037/0033-295X.115.2.502 (cit. on p. 21).
- [34] William Samuelson and Richard Zeckhauser. “Status quo bias in decision making”. In: *Journal of Risk and Uncertainty* 1.1 (1988), pp. 7–59. DOI: 10.1007/BF00055564 (cit. on p. 21).
- [35] Miles Hewstone, Mark Rubin, and Hazel Willis. “Intergroup Bias”. In: *Annual Review of Psychology* 53.1 (2002), pp. 575–604. DOI: 10.1146/annurev.psych.53.100901.135109 (cit. on p. 21).
- [36] Richard H. Thaler, Amos Tversky, Daniel Kahneman, and Alan Schwartz. “The Effect of Myopia and Loss Aversion on Risk Taking: An Experimental Test”. In: *The Quarterly Journal of Economics* 112.2 (1997), pp. 647–661 (cit. on p. 21).
- [37] Richard E. Nisbett and Timothy D. Wilson. “The halo effect: Evidence for unconscious alteration of judgments”. In: *Journal of Personality and Social Psychology* 35.4 (1977), pp. 250–256. DOI: 10.1037/0022-3514.35.4.250 (cit. on p. 21).

- [38] Ross Lee. “The intuitive psychologist and his shortcomings: Distortions in the attribution process”. In: *Advances in experimental social psychology* 10 (1977), pp. 173–220 (cit. on p. 21).
- [39] Dan Ariely. *Predictably Irrational, Revised and Expanded Edition: The Hidden Forces That Shape Our Decisions*. Harper Perennial, 2010. ISBN: 978-0-06-135324-6 (cit. on p. 22).
- [40] Scott O. Lilienfeld, Rachel Ammirati, and Kristin Landfield. “Giving Debiasing Away: Can Psychological Research on Correcting Cognitive Errors Promote Human Welfare?” In: *Perspectives on Psychological Science* 4.4 (2009), pp. 390–398. DOI: 10.1111/j.1745-6924.2009.01144.x (cit. on p. 22).
- [41] Pat Croskerry, Geeta Singhal, and Sílvia Mamede. “Cognitive debiasing 1: origins of bias and theory of debiasing”. In: *BMJ Quality & Safety* 22.Suppl 2 (2013), pp. ii58–ii64. DOI: 10.1136/bmjqs-2012-001712 (cit. on p. 22).
- [42] Karl Popper. *The Logic of Scientific Discovery*. Routledge, 2002. ISBN: 978-0-415-27844-7 (cit. on p. 22).
- [43] Karl R. Popper. “The Problem of Demarcation”. In: *Popper Selections*. Ed. by David Miller. Princeton, 1985, pp. 118–130 (cit. on p. 22).
- [44] Luc Bovens and Stephan Hartmann. *Bayesian Epistemology*. Oxford : Oxford ; New York: Oxford University Press, 2004. ISBN: 978-0-19-927040-8 (cit. on p. 23).
- [45] William Talbott. “Bayesian Epistemology”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. 2015 (cit. on p. 23).
- [46] Lindsay S. Uman. “Systematic Reviews and Meta-Analyses”. In: *Journal of the Canadian Academy of Child and Adolescent Psychiatry* 20.1 (2011), pp. 57–59 (cit. on p. 23).
- [47] Richard H. Thaler and Cass R. Sunstein. *Nudge: Improving Decisions About Health, Wealth, and Happiness*. New York: Penguin Books, 2009. ISBN: 978-0-14-311526-7 (cit. on p. 25).
- [48] Richard H. Thaler and Cass R. Sunstein. “Libertarian Paternalism”. In: *The American Economic Review* 93.2 (2003), pp. 175–179 (cit. on p. 25).
- [49] Barry Stroud. *The Significance of Philosophical Scepticism*. OUP Oxford, 1984. ISBN: 978-0-19-824761-6 (cit. on p. 26).
- [50] Peter Unger. *Ignorance: A Case for Scepticism*. OUP USA, 1978. ISBN: 978-0-19-151969-7 (cit. on p. 26).