

The unrevisability of logic

Thomas Hofweber

May 14, 2021

Abstract

Can it ever be rational to revise one's own logic by one's own lights? In this paper I argue that logic is never rationally revisable, even if one's own logic gives rise to paradoxes and allows one to derive any conclusion whatsoever. Instead of revising logic, we need to revise a certain widely held position in the philosophy of logic, one tied to the standard conception of validity and to the alleged monotonicity of deductive reasoning. I develop the alternative conception of validity and of deductive reasoning, which explains why we are generally entitled to the conclusions of deductive reasoning, even though we rationally do not accept them in certain circumstances tied to paradoxes.

Contents

1	The puzzle	1
1.1	The background of the puzzle	1
1.2	The argument against revision	6
1.3	The argument for revision	8
1.4	What to do?	8
2	Defeating the argument against revision	9
3	Defusing the argument for revision	15
4	Validity and truth-preservation	17
4.1	Generic vs. strict validity	17
4.2	Deductive reasoning and defeasible reasoning	19
4.3	Non-monotonic deductive logic	22
5	Which conception of logic is correct?	26
5.1	The descriptive question	26
5.2	The normative question	29
6	The solution to the puzzle	31
	Bibliography	33

1 The puzzle

The question whether it is ever rational to revise one's own logic gives rise to a real puzzle. There are apparently compelling arguments for both sides: that revision is never rationally possible, and also that revision is not only rationally possible, but in fact rationally required. I will discuss these arguments in detail shortly, but in essence the former is that one can prove with great certainty that one's present rules are valid, and thus the right rules, while the latter is that one's present rules lead to paradoxes, and thus need to be changed. Whatever one says in response to this puzzle has significant ramifications for the epistemology and philosophy of logic. But despite its significance, I dare say that the puzzle has never been properly resolved. In fact, it is arguable that the puzzle has hardly ever been properly addressed. Some would dismiss it outright, since they hold that logic has no connection to reasoning and rationality, and so no issue even arises about how we can rationally revise our own logic. Others see the push for revision not to be connected to revising logic, but only to revising something else, most prominently our "naive" reasoning with truth. And others still dismiss the puzzle quickly as being easily resolved, since they take the arguments against revision to make a simple mistake in epistemology. I will argue that none of these reactions are justified. Instead, the puzzle is real and substantial, and to properly resolve it requires us to step away from what is widely accepted in the philosophy of logic.

In this paper I hope to make progress on resolving the puzzle about the rational revisability of logic by arguing that logic, including the rules for truth, indeed is never rationally revisable. This holds even though rational revision of logic makes sense in principle, but, nonetheless, it is never actually rational to revise one's own logic, even if it leads to paradoxes. The rational reaction to the paradoxes is not to revise logic, nor to revise the rules for truth, but something completely different: giving up a widely held view in the philosophy of logic. The key to the defense of this solution to the puzzle is a certain lesson about the relationship between valid rules of logic and deductive reasoning. This lesson has long been absorbed in our practice of deductive reasoning, but not yet in our philosophy of logic.

1.1 The background of the puzzle

To talk about the revision of logic might sound as strange as talking about the revision of physics. Sure, we can change physics understood as a theory: we can simply formulate and

accept a different theory. But we can't change physics understood as the physical laws of nature. What these laws are has nothing to do with us, and so nothing we can do will make a difference to them. Similarly, one might hold that the only sense of a revision of logic is to be understood as the revision of a theory about what the laws of logic are. We can change our minds about what the right logic is, but we can't change what the right logic is, or so this sentiment. And this seems clearly right for physics, and to an extent also for logic. But nonetheless, there arguably is a sense of revising logic that doesn't carry over to revising physics. And so understood, it does make sense to consider whether it is rational to revise logic. This sense is tied to the traditional view that logic is relevant for good reasoning, in particular that logic is normative for reasoning and thought. Some philosophers reject such a connection, as we will discuss shortly, but even if there is this connection, how precisely logic might be normative for reasoning or thought is unclear and controversial,¹ But if the connection obtains in some form or other, then this will mark a key difference between logic and physics. It would give rise to the possibility that there is something that we can do rationally that changes which norms tied to logic apply to our own thought or reasoning. And this might lead to a sense in which we could rationally revise logic.

Some philosophers, most notably Gilbert Harman,² have objected to tying logic to inference and reasoning, and to holding that logic is normative for reasoning or thought in a way that physics is not. Harman argued that to believe in such a connection is to confuse entailment with inference. For example, if p entails q and you believe p it does not follow that you ought to believe q ; maybe you ought to give up the belief in p instead. This position can be taken in a radical way or in a more modest way. The radical way holds that there is no connection between reasoning and logic at all; they are simply about different things — entailment and inference — and there is no natural way in which these different things are connected. The modest way is to hold that although there is no positive *ought* coming from logic about what you should believe, there still is a connection between logic and reasoning: for example, logic requires that we do not have certain combinations of beliefs, say p and $\neg q$ in the above example, but it leaves open which particular ones of these beliefs we should have. Here there are numerous options on how logic might rationally constrain reasoning or thought that are different than providing a direct guidance on what one ought to believe: logic might prohibit certain combinations of attitudes, or

¹For discussion of some of the options, see [MacFarlane, 2004], [Steinberger, 2019] and [Steinberger, 2017].

²See [Harman, 1986]. A similar view is also defended by David Israel in [Israel, 1980].

permit certain combinations of attitudes, or require something, and so on. It is a little unclear to me whether Harman's official view is the radical or the modest one. In any case, only the radical one would separate logic and reasoning enough so that there likely will be no real issue about rationally revising logic. But for those who maintain that there is some connection between logic and reasoning, even if it is only a constraint on which beliefs are collectively rationally forbidden, there likely will be a coherent issue about whether it is ever rational to revise logic. The easiest way to motivate this is to think of the issue terms of conceptual engineering, that is, thinking about which concepts we should have. In particular, thinking about which logical concepts one should employ in reasoning naturally leads to a coherent sense of revising one's own logic.

Suppose I presently use classical negation when I represent something as not being the case. This, I will assume contrary to the radical option mentioned above, comes with some normative requirements on my reasoning or thinking, at a minimum that certain combinations of beliefs involving that concept of negation are rationally forbidden. Since my concept negation is the classical one, it makes sense to think of classical logic as my own logic, assuming, of course, my other relevant concepts are classical as well. It is the logic that governs my concept of negation, as well as the other logical concepts. Norms tied to classical logic thus apply to me in particular, since my relevant concepts are classical. But now I can ask whether I should not have a different concept of negation instead, say intuitionistic negation. Should I replace one concept of negation in favor of a different one? Should I employ the concept of intuitionistic negation where in the past I have employed my concept of classical negation? And should I consequently put myself under the norms tied to intuitionistic negation, and free myself of the norms tied to classical negation? I might have reasons for and against this. Suppose I have overall better reasons for switching concepts, and that those reasons are better even by my own lights: I can appreciate that intuitionistic negation is on balance of the reasons preferable. If I consequently make the switch, then this would give rise to a sense in which I would rationally revise my own logic: I would switch from one set of requirements on my beliefs to another one, tied to switching from employing one concept to employing another one. And this is a revision of logic not in the sense that I affect the eternal logical truths, whatever they might be, but in the sense that the rational requirements on my own reasoning might rationally change, via a switch of concepts.

This situation illustrates one, but certainly not the only, way in which it might be rational by one's own lights to revise one's own logic. As we saw, not everyone agrees that there is a

coherent sense in which one can revise logic, in particular several philosophers inspired by the radical reading of Harman's position have denied this.³ This paper is not the place to settle the debate about the normative status of logic, and I will thus put aside the radical Harman-inspired view that denies any normative significance of logic for thinking or reasoning. We shall assume instead, as I believe to be the case, that there is such a connection, although we can leave it largely open what that connection is more precisely. And with this we can grant that there is a sense in which a coherent question can arise about the rational revision of logic. That question simply is this: Can it ever be rational by our own lights to change our own logic?

As I announced above, I will argue that logic is never rationally revisable, even in light of the paradoxes. To clarify this we should distinguish two senses of logic being rationally unrevisable: first a strict sense, on which logic is unrevisable because it makes no sense at all to revise logic, and second a loose sense, on which it is coherent to revise logic in general, but it would nonetheless never be rational for us to revise our own logic, for some reason or other. I accept that logic is revisable in the strict sense, i.e. that revision is coherent in principle, but that logic is never rationally revisable in the loose sense: it would always be irrational for us to revise our own logic. Before we can see why, we still need to get clearer on the concepts of *our own logic* and of something being rational *by our own lights*. I shall briefly discuss both of these notions before I can directly address the puzzle about the rational revision of logic, which is the main topic of this paper.

To have a particular logic as *one's own logic* has two parts: first, the norms tied to that logic have an effect on one's reasoning. What these norms are we had to leave largely open, and what effect should be required we can similarly leave largely open, but somehow they have to guide or affect one. If the relevant norms have the required effect we can say that one is reasoning *in accordance with* the deductive rules tied to that logic. But this alone is not enough. The second part is a reflective endorsement of this reasoning. Those two can come apart, and if they do it is not so clear what one's own logic is. I might endorse intuitionistic logic, but in fact generally reason in accordance with rules only licensed by classical logic. Or I might endorse classical logic, but fail to use its power in reasoning, and de facto only reason in accordance with intuitionistic logic. In such cases there is a clear internal tension in one's mind, one that can be

³For discussion of revision and logic in the spirit of Harman's view, see [Priest, 2014], [Russell, 2015] and [Beall, 2015]. For a criticism of Harman's view, see in particular [Field, 2009], which I find persuasive. Harman replies in [Harman, 2009]. See also [Christensen, 2004], which concerns the relationship of logic to the revision of both full belief as well as degrees of belief. I will focus only on the former in this paper.

resolved either by endorsing different forms of reasoning or by reasoning differently. I will leave such mixed cases aside and focus only on cases where one is fully behind a particular logic: one reasons in accordance with the rules of that logic and one reflectively endorses this reasoning as correct, at least for each of the rules taken on their own. To reason in accordance with a logic does not require, of course, that one's reasoning is flawless. One can make mistakes in one's reasoning and thus the application of one's own logic, while still being guided by particular rules and endorsing this guidance.

To revise one's own logic is to adapt and endorse a different one than the one one presently has. To rationally revise one's own logic requires one to have all things considered better reasons to adopt a new logic and abandon one's present logic. I should add that our question here is not whether logic can be rationally revised on empirical grounds, but rather whether it can ever be rationally revised, considering all empirical and non-empirical reasons. In fact, most of the relevant reasons we will consider below are not empirical, but some are. Revision by itself is clearly possible, as one might undergo brain manipulation, brainwashing, or being persuaded by bad arguments to switch from one logic to the next. To rationally revise logic requires one to consider the reasons for and against revision, and then to properly weigh these reasons and conclude that those in favor of switching outweigh those in favor of not switching. In particular, if one concludes that one should switch, then one would have to accept that one should from then on reason in accordance with the new logic, and no longer with the old one. The interesting question here is whether it can ever be rational *by one's own lights* to switch. When one considers all the available reasons from one's present perspective, and evaluates them properly as they should be evaluated according to one's own best standards of doing so, can it ever be that the reasons for switching outweigh the ones for staying? If so, then one's own perspective and current position requires one to move from reasoning one way, tied to the old logic, to a different way, tied to the new logic.

Being rational by one's own lights contrasts with being rational on a purely external standard of rationality. On such an external conception of rationality it is not puzzling how it might be rational to switch. What is rational in the external sense might not be accessible to a thinker, and thus simply be a requirement that is there, externally, with it playing a direct role in one's own deliberation. For example, if in general it is rationally for me to do as God demands, whether or not I know about this, and whether or not I accept God's demand as a good reason for doing what is demanded of me, then I ought to switch if God demands so, and it would then

be rational in the external sense for me to do so. The interesting question is not whether it can be rational to switch on such a conception of what is rational, but rather whether it can ever be rational by one's own lights: where one is in a position to properly appreciate the reasons from one's present position and to properly conclude that they are sufficient reasons for switching.

My main example for a starting point of a logic that is one's own will be classical logic together with the "naive" truth rules.⁴ The argument of this paper does not depend on having this starting point, as it would carry over to many, but not all, other logics one might have at the outset. Still, this starting point seems to me to be a good one, since much of our reasoning does appear to be classical. This holds in particular in mathematics, which exhibits deductive reasoning in its purest form. And our reasoning about truth does seem to be in accordance with the naive truth rules. Both of these claims are clearly debatable. In particular, it is unclear if the classical rules for the conditional in fact capture our reasoning with conditionals. Conditionals in natural language are clearly much more complicated than the material conditional. But we can simply leave this issue aside and focus only on our reasoning with negation, disjunction, conjunction, and truth. For them it is much more plausible that we do in fact reason according to the classical rules. In particular, when I reflect on my own situation, I find myself exactly in this position: I think I am reasoning in accordance with the classical rules and the naive truth rules, and I endorse this reasoning, one rule at a time, upon reflection. The puzzle about revisability arises quite clearly for this starting point, as I will make clear shortly, and it similarly would arise for many other ones. Our starting point is thus that of someone who reasons classically and with the naive truth rules, who approves of this reasoning one rule at a time, and who considers whether or not they should revise these rules and switch to a different logic or different rules for truth, in particular in light of the paradoxes. I take this to be my own situation, and also that of many, although clearly not all, others.

Our main question then is this: assuming a connection between logic and good reasoning, can it ever be rational by one's own lights to switch from one's own logic to a different one?

1.2 The argument against revision

To rationally revise logic is to have better reasons overall by one's own lights to adopt a different logic than the one one presently endorses and reasons in accordance with. Suppose I presently

⁴The "naive" truth rules are simply to infer from 'p' is true to p, and the other way round, from p to 'p' is true. Whether this rule really is naive in a bad way is controversial, hence the quotation marks. I will side below with those who think these are good and valid rules.

reason classically and with the naive truth rules, and endorse this as the proper way to draw logical inferences. And suppose also that my friend Iris suggests, and gives me some reasons for, adopting a different logic instead, say intuitionistic logic as our toy example. Intuitionistic logic differs from classical logic in that intuitionistic logic rejects double negation elimination, i.e. the inference from $\neg\neg p$ to p . Iris, in particular, argues that I should stop inferring p from $\neg\neg p$, and she gives me some reason for why this inference is not valid. I will then have to consider whether I should make the switch: give up classical logic and adopt intuitionistic logic instead. But on reflection I can see that Iris's reasons for the switch, whatever they might be more precisely, ought to be rejected by me. I can reason conclusively that the inference from $\neg\neg p$ to p is valid and thus a proper inference pattern for me to rely on in reasoning. I simply reason as follows: suppose $\neg\neg p$ is true. Then $\neg\neg p$ (using the truth elimination rule), and thus p (using a negation rule: double negation elimination). So p must be true as well (using the truth introduction rule). Thus the truth of $\neg\neg p$ guarantees the truth of p , and therefore is a valid form of inference. In particular, I could establish that double negation elimination is valid by deductive reasoning alone, without relying on any premises. What could be more certain than that? That is at least as certain as results in mathematics, since there I rely on deductive reasoning as well as some mathematical assumptions. The strength of this conclusion will easily outweigh Iris' reasons for switching, and thus overall it would be irrational for me to switch.

Now, I used classical reasoning as well as reasoning about truth to draw this conclusion, and that might seem question-begging and circular. But it isn't question begging in the strict sense where I rely on an answer to the issue at hand — is double negation elimination a valid rule? — as a premise in my reasoning. I didn't rely on any premises at all. All I did is use the rules which at present strike me as the best rules to rely on in reasoning. Those are the rules I at present rely on in evaluating any proposal that will come before me, be it where to park or what logic to adopt. And how else would I evaluate it? What else would I use to evaluate a proposal other than the best rules I presently accept? I use my best tools to assess whether or not the proposal put in front of me is one that I should adopt. And my best tools at this point still include classical logic and the naive truth rules. Using those I can conclude that the rules of classical logic are truth preserving and thus valid. In particular, I can conclude this for double negation elimination. Whatever reasons I was given for the switch, I can rationally reject them with the greatest strength that I accept: something established by deductive reasoning alone from no premises. And the same argument applies to all my deductive rules, including the truth

rules. Thus it can never be rational by my own lights to switch to a different logic.

1.3 The argument for revision

On the contrary, it can be argued that revision of logic, in particular of classical logic and the naive truth rules, is not only rationally permitted, but in fact rationally required. Famously, the rules of classical logic together with the naive rules for truth allow one to derive anything whatsoever. Paradoxes like the Liar Paradox or Curry's Paradox give one simple ways of reasoning towards a contradiction and towards any conclusion whatsoever. This is as bad as it could be for any set of rules of reasoning. We can derive anything we want by simply instantiating those rules, and even by our own lights this is a terrible situation to be in. We must recognize that something has gone wrong with our own reasoning in light of this. And the conclusion is clear: if what I took to be valid rules allow me to derive anything whatsoever then I must have been mistaken about which rules are valid. That they allow me to derive anything makes this obvious. Thus at least one of the rules I presently accept must be invalid, and I can see this by my own lights. I thus need to find the invalid rule or rules, and replace them with an improvement, which is to say: switch logic. And all this is perfectly rational and in fact required of me by my own lights.

Although it might be tempting to immediately put all the blame on the truth rules and leave logic as more traditionally and more narrowly understood alone, this is at first hard to justify. The rules for truth are initially just as compelling as the rules tied to the traditional logical expressions, and they are arguably less problematic than those tied to the conditional. Furthermore, we can deduce that the truth rules are valid in the same way as we deduced that double negation elimination is valid. It is thus at first more natural to consider the truth rules to be in the same ballpark as the other logical rules, and to be rules of logic in a broader sense. So understood, logic itself leads to contradictions, and thus it will be rational by our own lights to revise our own logic. Rational revision is not only possible, but required, assuming we start with classical logic or one of the many other options that lead to paradoxes.

1.4 What to do?

To resolve the puzzle one must take sides: one or the other of the above two arguments must be mistaken. The hard part is to say which one and why. At first it must be more tempting to hold that the argument against revision is mistaken, since it is hard to resist the suspicion that the

argument that our own rules are valid is problematically circular. This does indeed seem to be the most common reaction to this argument for the validity of our own rules, and this reaction is spelled out most explicitly in recent work by Hartry Field.⁵ Tempting as this option is, I will argue in the following that it is mistaken. I will endorse the other side instead: the argument that we need to revise our own logic in light of it giving rise to paradoxes and allowing for the derivation of any conclusion whatsoever is mistaken. The rational reaction to the paradoxes is not to revise logic, but something quite different. Before I can defend this positive proposal we first need to see why a reaction like Field's is mistaken.

2 Defeating the argument against revision

As we saw in section 1.2, we can argue that our own logical rules are valid by using those rules, including the truth rules. The argument needed no premises, and each step was exactly in accordance with one of the logical rules, including the truth rules. Clearly there is something circular in using the rule in question in an argument that this rule is valid. The question is whether this circularity is bad. As is common, we can distinguish a *premise-circular* argument from a *rule-circular* one.⁶ Our argument is clearly not premise-circular, since no premise was relied upon. But it also clearly is rule-circular, since we used the rule to conclude that this very rule is valid. The question is whether this rule-circularity somehow takes away from the argument and its conclusion, in particular in the context of arguing over which logic is the right one. And whatever one says here, it should fit into a larger epistemological account of this situation.

The best way to do this, it seems to me, is to tie it to the issue to epistemic defeat, in particular some version of undercutting defeat. And this is just the line recently taken by Hartry Field in [Field, 2020], where he argues that in light of challenges to one's logic the continued use of the challenged rules gets defeated. I hope to argue in this section that this attempt to respond to the argument is not compelling.

As is common,⁷ we can distinguish rebutting from undercutting defeat. These notions are

⁵See, for example, [Field, 2020].

⁶For a discussion of rule circular reasoning in the justification of particular logical rules, see, for example, [Dummett, 1978] and [Boghossian, 2000] for a more favorable take, and [Dogramaci, 2010], for a more unfavorable one.

⁷See [Pollock, 1974].

primarily motivated by examples like these: Rebutting defeat occurs when Pete tells me that p , but later Sue tells me that not p , and Sue is much more reliable on this matter. Undercutting defeat occurs when Pete tells me that p , but later Sue tells me that Pete is notorious for making things up. As a first approximation, they can be made more precise as follows: Rebutting defeat of (a reason for) a belief that p is a stronger reason to believe not p , whereas undercutting defeat of (a reason for) a belief that p is reason for holding that the reason on which one based one's belief that p isn't a good reason after all.

It is hard to see how the argument for the validity of our own rules can be defeated in the sense of rebutting defeat. The argument for their validity proceeded from no premises by employing only deductive rules, ones we reflectively endorse as the right ones. Deductive proof from no premises is the strongest form of an argument there is, or so I would evaluate the comparative strengths of reasons for believing something. No other consideration can be stronger, and thus no other consideration can defeat, in the rebutting sense, the conclusions I can derive this way. Not even the paradoxical arguments using the liar or Curry paradox for an arbitrary conclusion are this strong. These paradoxical arguments are weaker, since they at least uses the uncontroversial premises that sentence C is the Curry conditional or that sentence L is the liar sentence.⁸ If the argument for the validity of the logical rules is to be defeated, it will first need to lose its strength, that is, it needs to be undercut.

We can distinguish two kinds of undercutting defeat: the strong version, which completely takes away the reason one originally had for the belief, and a weaker version, which doesn't take the reason away, but diminishes the degree to which it supports the belief. The difference corresponds to learning that the person on whose testimony you relied is notoriously making things up, and thus you lose all entitlement to the belief you formed on the basis of their testimony, and to learning that they are sometimes mistaken about the relevant issues, which only diminishes your entitlement, but doesn't completely take it away. In the following I would like to distinguish these two cases and use *undercutting defeat* only for the strong version, and *erosive defeat* for the weak version. This difference will correspond to us being entitled to still rely on the arguments that our own rules are valid, but the arguments will have less force than before after some erosive defeat has occurred, and us not being entitled to rely on these arguments at all any more after undercutting defeat has occurred.

⁸The Curry conditional is the sentence C which says if C is true then Santa exists, the liar sentence is the sentence L which says that L is not true.

In a recent article, [Field, 2020], Hartry Field has argued that reasonable challenges to one's own logic lead to undercutting defeat of the rules in question. He illustrates this with an analogous scenarios of our perceptual beliefs being challenged. Consider the debate about composition: are there ordinary objects like tables, or are there only simples arranged table-wise, but no tables? It is tempting to simply rely on perception to rule out the view that there are no tables. After all, I can see a table right in front of me. But the nihilist challenger, who holds that composition never occurs, will insist that such perceptual beliefs do not entitle you any more to believe in objects after nihilism has been pointed out as an alternative. After all, things would look just the same even if composition did not occur.⁹ In light of the challenge by the nihilist our entitlement to our perceptual beliefs in objects gets defeated. It might have been fine before, but now it is no longer, and thus relying on such beliefs to refute the nihilist is illegitimate, or so the argument. And similarly in the debate about logic. Once intuitionism has been suggested as an alternative logic, my reasoning with double negation elimination has been defeated and thus cannot be relied upon any more in the refutation of intuitionism and my argument that double negation elimination is valid.

What brings about this defeat? There are in principle several options one has to support this, but Field endorses what seems like an especially radical option: assuming the alternative meets some minimal threshold for being reasonable, the proposal of such an alternative alone leads to defeat, even if no evidence or reasons are presented in favor of that alternative. To quote Field: "It needs to be emphasized that the old observational practice isn't defeated *by new evidence*, it is defeated *by the suggestion of a new theory*." ([Field, 2020, 6]. Emphasis in the original).¹⁰ A less radical option would hold that defeat only happens once one has sufficient reasons to take one's present belief formation to be flawed. It is unclear if either the nihilist challenge to our perceptual beliefs about objects nor the intuitionist's challenge to our use of double negation elimination live up to this stronger standard for defeat, but they clear do live up to the weaker standard of being a suggestion of an alternative that is worthy of being considered seriously. And if such a suggestion alone defeats the rules that are being challenged, then I will have to retreat to neutral ground: I will have to reason using only rules that have

⁹Field in [Field, 2020] uses a different, and, I think, somewhat more complex example to illustrate this point: the debate about the heliocentric vs. geocentric worldview and rejecting the former on the basis of seeing objects fall in a straight line. Essentially the same lesson will apply to both examples.

¹⁰See also p.14 for the qualification that not any crazy alternative will work, and that which alternatives should be taken seriously can vary. I take it, though, that all the alternatives under discussion here will meet the threshold of being reasonable in the sense that they are worthy of being taken seriously.

not been challenged. And thus I will have to face the challenge the intuitionist puts forward without relying on double negation elimination and thus without my proof of its validity, which uses it. Those rules would have been defeated in the undercutting sense, or so Field.

Although this might seem like a reasonable position to take at first, I think we can see that it is mistaken. In particular, to hold that we face defeat from the mere suggestion of a (minimally reasonable) alternative strikes me as the wrong reaction. It is true, of course, that the mentioning of an alternative that one had not considered has some legitimate epistemic effect. One should at least consider it, and maybe consider it very seriously. But to take it to simply defeat one's perceptual beliefs or ones conclusions of deductive reasoning is a complete overreaction. To take undercutting defeat to have occurred means that the perceptions or deductive arguments do no longer support these conclusions at all. If the mere proposal of an alternative worth taking seriously were enough for defeat, it would be easy to defeat most of our beliefs, no matter what they are based on. In fact, most would already have been defeated, given all the alternatives that have been mentioned in philosophical discussions: simulations, unreliable memories, hallucinations etc.. What would defeat a particular belief, I take it, is that one has good reason to think that the alternative obtains, and thus that the belief might not be true. That is what happens in the classical example of me relying on the testimony of someone, and then learning from someone else, and thus have good reason to accept, that the first person is a notorious liar. But mentioning an alternative alone, even a reasonable one, is not enough for defeat.¹¹

The case of challenging the logical inferences seems to me to be analogous to someone challenging relying on rational thought more generally. 'Don't think so much, feel what's in your gut', someone might say. Should I do this? Should I rely on my gut more, and not try to figure things out by thinking? Am I still entitled to think about this proposal, and accept the conclusion I draw with this thinking, or do I have to accept a neutral position beyond thinking and feeling and settle the issue that way? Do I have to follow my heart on this issue, as the middle ground between the gut and thought?¹² It seems to me that the answer is 'no'. I am

¹¹Field is not alone in endorsing such a radical picture of defeat in philosophical debates. An endorsement of this epistemological stance in the debate about composition is found, amongst others in [Merricks, 2001], [Rosen and Dorr, 2002], and [Sider, 2013, 260]. In particular, Sider endorses that merely the suggestion of nihilism defeats our entitlement to our perceptual beliefs in objects. I have defended the opposite view, which endorses that our entitlement to our perceptual beliefs in there being tables is not defeated by bringing up the nihilist alternative, and that we do not have to enter neutral ground, in chapter 7 of [Hofweber, 2016] and in [Hofweber, 2019].

¹²Thanks to Kevin Scharp for the tongue-in-cheek suggestion of the heart as neutral arbiter.

still entitled to think about this issue and rightly reject it on the basis of rational thought. My entitlement to thinking is not defeated by this challenge, it is merely challenged and rightly rejected after thinking about it. This, too, has the ring of circularity to it, but it is nonetheless a legitimate way to answer the challenge. If I had reason to think that following my gut is getting me closer to the truth than if I made a decision based on thought, then I would have to accept that I should follow my gut. But I should make this decision on the basis of thinking about the issue. It is not that thinking has been defeated in light of the suggestion of the gut as an alternative, but rather thinking has decided that it is best to follow my gut, at least in certain cases. But whether thinking itself can ever be rationally abandoned in general in favor of the gut or the heart is quite another issue, one where the answer likely is that it can't be. And such an answer, I hope to argue for below, is also the right one when it comes to revising our own logic.

Undercutting defeat, which takes away all entitlement, in light of the suggestion of an alternative is too strong, but maybe erosive defeat, which only takes away some entitlement, is more plausible. And this seems to be right reaction in the case of perception. The nihilist suggestion, even without evidence, that composition might not ever occur, does take away from my confidence in my perceptual beliefs, a little bit at least. Perception does not conclusively settle whether there is a table in front of me, even though I get evidence for it by looking. The evidence is still good, but a little less good than before the challenge. The issue will need to be ultimately settled by bringing in other considerations as well: science, metaphysics, whatever else might work.¹³ In the case of challenges to perception and many other ways in which we form beliefs, erosive defeat does work, at least to an extent, as long as the relevant suggestions of alternatives are at least somewhat reasonable.

But in the case of deductive inferences the situation is not so clear. My deductive proofs from no premises settle an issue conclusively by my own lights. When I weigh the reasons for the conclusion of the proof with those from some challenge to it, it would seem that I always have to side with the proof. After all, the proof is correct by my own lights, and thus settles the issue conclusively, whereas the challenge is just that: a challenge. In the case of perception the reasons for objects are not conclusive, although pretty strong. Here erosive defeat is possible, since even by my own lights I have to accept that I might be wrong in my belief in objects, and

¹³I have argued in [Hofweber, 2016] and [Hofweber, 2019] that overall the evidence is in favor of the existence of ordinary objects and against nihilism.

the attempt at defeat should increase my confidence a little that I am indeed wrong. But when the reasons are conclusive, then I cannot rationally allow for that. If all the reasons available to me conclusively settle that p , then I can't allow that I might be wrong, and thus that a certain consideration should erode my confidence in p .

This might seem like a stalemate and like a chicken and egg situation: who gets to move first, the attempt at erosive defeat, which would result in the reasons not being conclusive any more, or the conclusive reasons, which would result in the attempt at defeat being rationally rejected. But on reflection it seems quite clear that the conclusive reasons come first if we are fully rational. We have those reasons available to us at any time, and thus if we are properly responsive to all the reasons available to us, then we should rationally reject any attempt at erosive defeat. Thus if I am fully rational and any attempt at erosive defeat of my otherwise conclusive reasons comes my way, then I must rationally reject this attempt as being mistaken. If not, then I would not have responded appropriately to all the reasons that I already had available. This applies to conclusions established purely deductively from no premises, and it again nicely contrasts with conclusions reached on the basis of perception. In the case of perception I am entitled to hold that any attempt at defeat is likely mistaken. But this does not neutralize the attempt, it only weakens it. Thus in the case of perception, erosive defeat is possible, since the reasons tied to perception are not conclusive, while in the case of deduction it is not.¹⁴

Thus deductive reasoning from no premises cannot easily be defeated. It cannot be undercut by the suggestion of alternatives, it cannot be eroded, and it cannot be rebutted. If undercutting and eroding defeat does not happen, then the original force of the deductive arguments remains in place. And that force, for the case of deductive proof from no premises, is the strongest one I accept, by my own lights. That means that the argument for the validity of my own deductive rules is stronger than the argument for arbitrary conclusions using the paradoxical reasoning, since the latter relies at least on some trivial premises concerning the identity of the sentences referred to in the liar sentence or Curry conditional. Nonetheless, this conclusion must seem highly problematic. The argument for arbitrary conclusions in the paradoxical reasoning seems airtight, and the trivial premises relied upon are indeed trivial. The natural pull still seems to be that something has gone wrong in the epistemological reasoning engaged in above, and that

¹⁴This issue is related to, but not quite the same as, the “dogmatism puzzle” discussed in [Kripke, 2011], [Harman, 1973], and [Lasonen-Aarnio, 2014].

the rational reaction is to give up one of the classical rules or the “naive” truth rules. But I hope to argue for a different solution to the puzzle in the following. We cannot rationally give up any of the deductive rules we accept by our own lights. And we also cannot rationally accept the arbitrary conclusions we can derive with them. The proper solution to the puzzle lies elsewhere, and we can get there by thinking about what goes wrong with the argument that the paradoxes require a revision of our logical rules.¹⁵

3 Defusing the argument for revision

It might seem almost inevitable that we need to revise classical logic or the naive truth rules in light of the paradoxes. The paradoxes allow us to derive anything whatsoever from the starting point of classical logic and the naive truth rules, and so we must recognize by our own lights that we are in a bad situation that we need to get out of. The rational conclusion to draw from this seems to be to revise our own logic or the naive truth rules. But the situation is in fact more complicated. If my own logic is classical logic including the truth rules, then I have available to me a very compelling reason, by my own lights, to not revise logic. For example, using the reasoning in the Curry paradox with the conditional ‘if this sentence is true then I should not revise my own logic,’ I can conclude that I should not revise my own logic. But I can also see that I can use this same argument for any conclusion whatsoever, including that I should revise my own logic. Thus overall it is not so clear what I should do rationally. Maybe the the proper reaction needs to be that rationally gives out at this point, or maybe there is a better, more appropriate reaction.

The first, and maybe foremost, attempt to resist the revision of classical logic in light of the paradoxes is to give a different status to the more narrowly logical rules and separate them out from the truth rules, and then to put all the blame on truth. This would support the conclusion that the logical rules are fine, but the truth rules need to be revised.¹⁶ I think this is a mistake. The truth rules at first have the same status as the logical rules in that they are forms of inference that I reason in accordance with and reflectively endorse. Furthermore, I can argue for the validity of the truth rules just as I can argue for the validity of any of the other rules. This is not to say that it could not be argued in the end that the truth rules have to go,

¹⁵I am indebted to Ram Neta, Jim Pryor, and Alex Worsnip for discussions of the material in the above section.

¹⁶The history of this move is surely long, but see [Scharp, 2013] for an overview and a recent defense. Scharp’s proposal is especially radical, since he argues that we need to replace the concept of truth itself with two other ones.

but this will need to be worked out, and at first it looks like the truth rules are not rationally revisable just as neither of the rules tied to negation, conjunction or disjunction are. Thus at first we should take them to be on a par.

There is also a second attempt to defuse the argument for revision, which I also want to put aside. It is the attempt to put all the blame on the paradox-inducing sentences: the liar sentences, the Curry conditional, and so on. The claim here might be that they are defective in a way that makes them illegitimate to instantiate in valid rules: they are meaningless, or somehow not well-formed. I think this is also a mistake. Those sentences, it seems to me, are perfectly legitimate instances in the schematic rules, even though they are somewhat unusual sentences. The rules are schematic and need to be instantiated with particular meaningful and grammatical sentences. But the paradox-inducing sentences do meet that condition: they are meaningful and grammatical.

If we put aside these two attempts to defuse the argument for the revision of our own logic, then it seems there is little hope to defuse it otherwise. After all, we are in a position where we accept all of the following:

1. The classical rules, including the naive truth rules, are valid.
2. Legitimate instances of valid rules are truth-preserving.
3. The liar sentence and Curry conditional are legitimate instances.
4. We can derive arbitrary conclusions with rules that are valid and truth preserving and that we accept to be valid.

But clearly, it would be irrational to accept arbitrary conclusions.

All this leads to an apparent trilemma: either we need to give up one of our rules, and thus revise logic, or reject that the paradox-inducing sentences are legitimate instances, or we need to accept anything whatsoever. To defuse the argument for revision seems to require us to pick our poison, as no other options seem to be available. Defusing the argument while saving all the rules and instances seems to be ruled out, unless we accept anything whatsoever.

Nonetheless, I want to argue that we should continue to accept our rules as valid, accept that the paradox-inducing instances are legitimate instances, but we should not accept anything whatsoever. Furthermore, I want to argue that this is the rational reaction in light of the paradoxes, and in particular rationality does not give out when encountering paradoxes. All this might seem impossible, since how could we accept that valid inferences are truth preserving,

that valid inferences allow us to derive anything, but then not accept the conclusions that we ourselves are able to derive, and do so rationally? To see that possibility we need to reconsider a key issue in the philosophy of logic.

4 Validity and truth-preservation

4.1 Generic vs. strict validity

Valid rules are truth preserving, which is to say, instances of those rules preserve truth: if the premises are true, the conclusion has to be true as well. The key to defusing the argument for revision is to see that this can be understood in at least two ways. The sentence

- (1) Instances of valid rules are truth preserving.

has at least two readings. One is the *strict reading*: each and every instance of a valid rule is truth preserving. This is how the statement is generally taken in the philosophy of logic. For a rule to be valid it is required that each and every instance is truth preserving. But (1) also has another reading. This is the *generic reading*, a reading very familiar in other case of bare plural statements like:

- (2) Bears are dangerous.

(2) has a strict and a generic reading as well. The strict reading says that each and every bear is dangerous, and it is false. Some bears are not dangerous: innocent declawed baby bears are not. But the generic reading is true. It can be triggered by restating it differently, although imperfectly, as: in general, bears are dangerous, or normally, bears are dangerous.

We commonly represent the world with generic statements, and they often play a key role in ordinary reasoning. For example, when I hear that there is a bear nearby, I am entitled to conclude that I am in danger, since I know that bears are dangerous. I reasonably and rationally make the inference that I am in danger from the premises that there is a bear nearby and that bears are dangerous. But such inferences with generic statements have the feature that they are non-monotonic: when I learn more I can lose my entitlement to that conclusion. When I learn that the bear nearby is a cute declawed baby bear, then I am no longer entitled to conclude that I am in danger. Note that in this case I still believe all the premises that I believed before: bears are dangerous and there is a bear nearby. But now I know more and that extra information

does not contradict what I knew before, instead it adds to it. With this extra information the inference is no longer a good one, even though it was good before. This non-monotonic feature of reasoning with generics distinguishes them from how deductive reasoning is generally understood. In general it is assumed that deductive reasoning is monotonic: more premises allow you to deduce more things, but never fewer. Whether this is indeed true for deductive reasoning is something we will revisit shortly.¹⁷

We can thus distinguish two conceptions of validity, which correspond to the two readings of (1). Let us call *strict validity* the notion of validity which arises from holding that (1) is true on the strict reading: valid inferences are truth preserving in the sense of that each and every instance is truth preserving. And let us call *generic validity* to notion of validity tied to the generic reading of (1): instances are truth preserving, understood as a generic statement.¹⁸

Normally philosophers of logic take validity to be strict validity. The whole point of logic, they claim, is to uncover patterns that are truth preserving without exception. That logical validity is strict validity is the standard view, but it is a view about logic nonetheless. In principle the alternative view that logic is connected to generic validity should be on the table as well, at least once proposed. Although understanding validity as each and every instance being truth preserving is the standard view, not everyone agrees with it. One exception is Hartry Field, who has argued that validity should not be understood in terms of truth preservation. His argument is related to the unprovability of the consistency of our own assumptions, and what goes wrong with the argument that our assumptions must be consistent, since they are all true and our rules are all truth preserving, thus everything we can derive is true, and so consistent. Field instead holds that validity must be understood as a primitive notion, distinct from truth preservation.¹⁹ Whether that is a good argument is, of course, debatable. For one, the rules we normally accept lead to paradox as well. Is Field proposing a rational revision of our concept of validity, or that it always was a primitive notion, and that philosophers made a mistake in trying to spell it out in terms of truth preservation? There is naturally an issue about whether our notion of validity is rationally revisable, one related to some of the issues that arise

¹⁷As is well known in the literature on generics and their connection to default inferences, not all generic statements are connected to such inferences. For example, that mosquitoes carry malaria does not entitle me to infer that this particular mosquito carries it. Which generics license these inferences is controversial, and I won't aim to settle this here, obviously. I will work with the picture, however, that generic validity tied to our deductive rules is of the kind that licenses them. The justification for why they license them is discussed below. For the basics on generics and their connection to default reasoning, see [Carlson and Pelletier, 1995] or [Pelletier and Asher, 1997].

¹⁸See [Hofweber, 2008] for a discussion of this conception of validity.

¹⁹See, for example, the second half of [Field, 2009], chapter 2 of [Field, 2008], and [Field, 2006].

for revising our own logic. We don't have to settle on Field's proposal here, nor do we have to settle whether he has good reasons for holding that a primitive notion of validity was our notion all along. But we will have to face similar issues when it comes to the question whether we do, or should, have a notion of generic validity as our operant notion of validity. I will address these issues below, where I aim to explore and defend the following view: It is correct to spell out validity in terms of truth preservation, but it was a mistake in the philosophy of logic to do so in terms of strict validity. Instead, validity is generic validity. Not merely should we accept this as the correct account of validity, it is the one we in fact rely on in our actual practice of deductive reasoning. But before we get to this, let us consider what deductive reason would look like when understood as concerning generically valid inference rules.

4.2 Deductive reasoning and defeasible reasoning

It is natural to develop the idea of logical validity being generic validity in the following way. Just as reasoning with certain generic statements entitles us to draw inferences, so we are entitled to draw inference in accordance with the logically valid rules. Just as I was entitled to conclude that I am in danger in the bear example, so I am entitled to draw a conclusion in accordance with modus ponens or any of the other generically valid rule. But such reasoning is non-monotonic in both cases: more information can take the entitlement away, just as it was when we learn that it's a baby bear. Such inferences are thus defeasible in light of further information. The key here is that we are originally entitled to the conclusion drawn in such reasoning, but that entitlement can go away, not because we have reasons to hold the opposite, i.e. that I am safe, since I am in a bear proof room, or reason to reject one of the premises, i.e. when we have reasons that bears are not dangerous after all, but because the inference to the conclusion can be defeated in light of further information. It is uncontroversial, I take it, that reasoning in general is non-monotonic. But it is widely assumed that flawless *deductive* reasoning is monotonic, and thus is a distinctly different sub-part of reasoning. On the natural development of the idea that logical validity is generic validity this is rejected. Instead, deductive reasoning itself, even in the case where one reasons flawlessly in accordance with the simple rules for negation, disjunction, conjunction and truth, is non-monotonic.

This, I want to propose now, and defend below, is how deductive reasoning itself should be understood. When we draw inferences with valid forms of reasoning we are entitled to the conclusions we draw, but that entitlement is defeasible and can go away in light of further

information. Deductive reasoning itself is thus non-monotonic and defeasible, and insofar as deductive logic is tied to deductive reasoning, it, too, should be seen as non-monotonic and defeasible. This is not to be understood the way “default logic” or “non-monotonic logic” is normally understood in mathematical or formal logic. Such logics are based on strictly valid deductive rules together with further rules which are defeasible and allow for exceptions, like rules tied to reasoning about bears.²⁰ I propose not that deductive logic as standardly understood is to be augmented with further rules for default reasoning, but that the deductive rules themselves are default rules. They are valid forms of inference and proper ways to guide one’s reasoning, but even purely deductive reasoning is non-monotonic and defeasible. Our entitlement to deductively valid conclusions can go away in light of further information, even though it originally was there, in all its glory. There are two crucial parts to this: one more positive and one more negative. The positive part is that we indeed are entitled to reason in accordance with the valid rule, and we are entitled to the conclusions we draw this way.²¹ The more negative part is that this entitlement can go away, even when our reasoning is flawless and without us having any reason to reject one of the premises or forms of inference. The positive part is more highlighted by thinking of such reasoning as default reasoning, the negative part more by thinking of it as defeasible reasoning.

And this, I propose, is just what happens in paradoxical reasoning: we are at first fully entitled to draw the conclusions we draw, but we can lose that entitlement, even though we reasoned perfectly. For example, when we argue that Santa exists using the Curry conditional “if this very sentence is true then Santa exists”, or that the liar sentence is both true as well as not true, then we are at first entitled to these conclusions, since we can derive them using only valid inference rules from only uncontroversial premises. But when we reflect on the results we have achieved this way, then that entitlement goes away. Once we realize that this argument would have worked for any conclusion, or that this argument leads to a contradiction, then the entitlement to that conclusion that we originally had gets defeated. The reasoning we engaged in was perfect, just as in the bears example, but the entitlement we get on the basis of perfect reasoning can disappear in light of further information, even in purely deductive reasoning.

How should we understand the defeat of deductive reasoning, in particular in light of our

²⁰See [Reiter, 1980] and [Nute, 1994] for classic treatments of default and defeasible logic.

²¹Of course, such entitlement is only conditional on whatever entitlement I had for the premises relied upon in my reasoning, not absolute.

discussion of defeat in section 2? There are two questions to distinguish here: what kind of defeat is it, and when does it occur. I have argued in section 2 that properly done deductive reasoning is immune from erosive defeat, and that the rules of deductive logic are not defeated by the suggestion of reasonable alternatives. But that does not mean that defeat of a particular deductive argument cannot occur. The proposal here is that a particular deductive argument at first fully entitles us to the conclusion we draw with it, but upon further reflection, the details of which are to be discussed momentarily, that entitlement goes away. Thus this defeat is a version of undercutting defeat: the entitlement is not merely diminished, but disappears completely. But the source of the undercutting defeat is not some general epistemic principle that applies everywhere even outside of deductive reasoning. Instead it is undercutting defeat tied distinctly to deductive reasoning.

But when does this defeat occur? There are two main plausible answers: first, a particular argument gets defeated once you recognize that this argument would have worked not just for the conclusion we in fact drew, but for anything whatsoever. Second, an argument gets defeated if it arrives at both p as well as $\neg p$. Although within classical logic those are closely connected, in the philosophy of logic they are rather different. If p as well as $\neg p$ defeats a deductive argument, then negation plays a double role in logic and is special among all the logical expressions. It first has the usual role in reasoning, just like other logical expressions have their role. But in addition it also has a place in the defeat rule, contrary to the other logical expressions. Here I am taking the defeat rule to say: an argument is defeated once one recognizes that it achieves both p as well as $\neg p$. This is opposed to when it achieves the conjunction $p \wedge \neg p$, in which case conjunction would also play a role in the defeat rule. Although the conjunction can, of course, always be inferred from the conjuncts, I take this to be unnecessary for defeat to occur, and thus I take negation, but not conjunction, to be special among the logical expressions in this case. Alternatively, if we take defeat to occur when one can derive anything whatsoever, then no logical expression plays a role in the defeat rule, and all of them remain on equal standing.

My own opinion is that the defeat rule involving just negation is to be preferred. If we can see that we can derive any conclusions whatsoever with an argument, then at first this should seem as that the argument is extremely powerful. That by itself does not speak against it, arguments that allow one to derive lots of conclusions by themselves are a good thing. But once we recognize that among those conclusions is both p as well as $\neg p$, then we can see that the argument is not powerfully strong, but pathological. The reason why defeat happens in the case

of an argument that allows us to derive any conclusion whatsoever is that it allows us to derive both p as well as $\neg p$. That is the true reason for defeat, and thus the proper defeat rule is tied to negation. And it is tied to only negation, not also conjunction, as mentioned above. Thus I favor taking negation to be special among the logical expressions for philosophical purposes: not only is it tied to valid logical rules, it is the one and only logical expression appearing in the defeat rule for deductive reasoning itself.

To take the option that defeat occurs when we derive both p as well as $\neg p$ leads to the tricky issue of how we should understand proof by contradiction on this conception. Sometimes it seems being able to derive p as well as $\neg p$ does not lead to defeat, but to a proof that one of the assumptions is wrong. But sometimes this is not the lesson we draw, even on a strict conception of validity. For example, no one, to my knowledge, concludes that λ is not identical to the sentence “ λ is not true” from the fact that assuming it allows us to derive p as well as $\neg p$. Instead, the standard reaction within professional philosophy is that one of our deductive rules, including the truth rules, must be mistaken. What the rational reaction is to deriving p as well as $\neg p$ is an issue for both approaches, and clearly one that deserves more attention.

In standard cases of defeasible reasoning like our reasoning about bears, learning that the bear is a baby bear undercuts my entitlement to the conclusion that I am in danger. It does not rebut any of the premises in this reasoning, nor does it undermine the pattern of reasoning that I relied on. But in these cases we can explain why such undercutting is reasonable: I know why baby bears are not dangerous, even though bears are dangerous. In the case of defeasible deductive reasoning this is similar in some ways, but different in others. There, too, what is undercut is the conclusion that is based on a particular argument, not the pattern of reasoning that was relied upon. But contrary to cases of ordinary defeasible reasoning, we do not have an explanation of why this particular argument is undercut other than that it leads to p as well as $\neg p$. That undercutting occurs here is part of deductive reasoning, on the present proposal to be spelled out more shortly, not something that is derivative on any other insight.²²

4.3 Non-monotonic deductive logic

Logic is tied to reasoning, and thus has a normative dimension. But it is also tied to entailment. So far our discussion has mostly been about logic’s connection to reasoning. On the generic conception of validity, deductive reasoning is non-monotonic and defeasible. On the strict con-

²²Thanks to Fabian Klinge for pushing this issue.

ception it is monotonic and not defeasible. This concerns reasoning first and foremost. But what should we think about the entailment relationship? It is a relationship between sentences or propositions. Is it also non-monotonic? David Israel has argued in [Israel, 1980] that logic when it concerns entailments must be monotonic, essentially by definition, which isn't really much of an argument, on reflection. Non-monotonicity only concerns reasoning, he suggests, never logic itself, which is concerned with entailments, but not reasoning. But this seems a bit harsh, and too much tied to a particular conception of what logic is supposed to be. Formally both kinds of relations make sense. We can formulate and consider a non-monotonic entailment relation, as well as a monotonic one. Furthermore, they will look very similar to each other, in particular in cases where defeat will never occur. If the logical principles and assumptions do not entail p as well as $\neg p$, then no defeat of the deductive proofs will ever occur and the two conceptions of logic will agree on what is entailed by what. But when defeat can or does occur, then there will be differences. In these cases, what is entailed by a set of premises can change when you add more premises, which is to say, logic itself will be non-monotonic.

Monotonicity is one of the structural rules of standard deductive logic. And since classical logic is such a standard deductive logic, it is one of the structural rules of classical logic. To give up monotonicity thus seems to give up on classical logic after all. Even if the usual deductive rules of classical logic are valid, it looks like the structural rules won't survive. But here it is important to distinguish two ways in which the structural rules are valid, which correspond to our two ways in which the deductive rules are valid: strict and generic validity.²³ Monotonicity is not a strictly valid structural rule: the example of the paradoxes shows this. But monotonicity is generically valid. We thus only abandon monotonicity in one sense, as a strictly valid structural rule, but not in the other one, as a generically valid rule. In fact, monotonicity is a default structural rule on the present proposal, just as all the deductive rules are default valid rules.

This leads to an important difference from other approaches that deal with the paradoxes by rejecting some structural rule or other. For example, the proposal made by Pablo Cobreros, Paul Égré, David Ripley and Robert van Rooij in [Cobreros et al., 2013] holds that transitivity is to be rejected: it can be that B follows from A, and C follows from B, but nonetheless C

²³To call structural rules valid might seem off, but it can make sense. A structural rule like monotonicity says, informally that if you can derive p from statements in Σ , then you can derive p from statements in Σ together with q . To call such a rule valid is to say that instances (for the schematic letters Σ , p , and q) result in a rule that is guaranteed to preserve truth, or, for our formulation, in a true conditional claim. The notion of validity is sometimes not applied to structural rules directly, but it can carry over in the way outlined.

does not follow from A. The reasoning to paradoxical conclusions inappropriately relies on the transitivity of deductive arguments, they argue, and this points to the flaw in it. Although the rules tied to the connectives were used properly, the reliance on transitivity was the mistake. The problem for this view is to justify the reliance on transitivity in other cases. For example, in mathematics as well as ordinary reasoning we generally rely on transitivity: I use X's Lemma to derive my own theorem, and conclude from this that the theorem holds. But how can I be justified in doing this if transitivity is rejected as a structural rule? Some positive argument must be given by Cobreros et.al. that this is acceptable in mathematics. The likely candidate for such an argument is that in mathematics paradoxes do not arise and thus transitivity can be relied upon. But the problem with this approach is that it seems to require an argument that in mathematics paradoxes do not arise and only after this argument has been given can we rely on transitivity in mathematics. But what is this argument? We know that paradoxes can arise in a number of ways, not merely with the notions of truth, but also satisfaction, denotation, set, etc.. Although few are worried that standard mathematical reasoning can give rise to paradoxes, there is little argument that this is indeed correct. And without such an argument, what is the justification for relying on transitivity? In particular, transitivity cannot be relied upon in the proposed argument that in mathematics paradoxes do not arise, since this argument has to be given first to justify the reliance of transitivity. Thus on this approach the reliance on transitivity in mathematics is not properly justified, and so neither are the results of mathematics.

On the default conception of deductive logic the situation is noticeably different. Monotonicity is only denied on the strict conception of monotonicity. Just as with the ordinary deductive rules, we need to distinguish strict from generic versions of the structural rules. And just as by default the ordinary deductive rules are truth preserving and to be used in proofs, so by default the ordinary structural rules can be relied upon in proofs. The structural rules do have exceptions, in particular monotonicity, and thus they are not correct when understood strictly. But these exceptions are tied to cases that we by default are entitled to put aside, and thus by default the structural rules can be relied upon. In particular, no argument is required at the outset to justify why in this particular case we are not dealing with one of the exceptions. On the default conception we are entitled to this with nothing having to be done at first to earn this entitlement. That's just what it is to be entitled by default. Thus we are entitled to rely on monotonicity in our proofs until we encounter the exceptions and then the entitlement goes

away. Cobreros et.al. have to face the issue what justifies reliance on transitivity in ordinary cases, since they work with what I called a strict conception of deductive logic, although a substructural one. On such a conception the rejected structural rule is invalid in general, and any reliance on it needs to be justified. On the default conception the structural rules is strictly invalid, but generically valid, and thus can be relied up by default, without the need for an argument to justify this particular use of it. This is why there is no threat for mathematics on the default conception, but there is one on the strict, substructural conception of deductive logic.

Overall then we can distinguish two conceptions of logic and their relationship to deductive reasoning. One, the *strict conception of deductive logic* holds that entailment is strictly monotonic and the valid rules are strictly valid. If one holds that there is a connection of valid rules to good reasoning, as one should, then one will naturally hold that a set of rules is jointly refuted if one can instantiate them in a way that leads from truth to falsity. In particular, on this conception of logic the paradoxes refute classical logic including the naive truth rules. On the alternative *default conception of deductive logic* entailment is strictly non-monotonic, but generically monotonic, and valid rules are generically valid. Deductive reasoning is correspondingly defeasible and non-monotonic. The defeat rule will, in what seems to me to be the best version, hold that our entitlement to the conclusion of an argument is defeated if that argument leads to both p as well as $\neg p$. In particular, on the default conception of logic, a set of rules is not refuted if they jointly allow the derivation of a contradiction.

Both conceptions of logic make sense and are coherent options. It might be tempting to object that if we adopt the default conception of deductive logic we are in essence giving up on deductive logic, since, it is claimed, deductive logic concerns inference patterns that are strictly valid. But this is simply adopting a stance in the philosophy of logic, one that is widely held, but still a particular view about logic and what it is supposed to be like. What can be agreed upon by all sides is that logic concerns formally valid patterns of inference or entailment, and maybe even that such patterns are a guide to good reasoning. What is controversial, for present purposes at least, is whether validity is strict or generic, and whether logic so understood is strict or defeasible. This is the topic we need to address next, now that both options are on the table.

5 Which conception of logic is correct?

How then are we to decide which conception of logical validity, and with it of deductive logic, is the correct one? Is validity strict validity or generic validity? Both notions, we can grant, make sense and both conceptions of logic are coherent options. As logics, abstractly understood, strict deductive logic and default deductive logic are both fine. We shouldn't ask which one is the true logic. Considered in isolation, both are fine. But that doesn't address the real issue. The real question is not about abstract logical systems, but about us. It is about which conception of logic and of validity underlies our own deductive reasoning, and which one it should be. Which one, if any, do we reason in accordance with? And which one should it be? So, there are really two questions: a descriptive question about which conception of logic and validity we reason in accordance with, and a normative question about which one we should be reasoning in accordance with. We will consider them in turn.

5.1 The descriptive question

To find out which logic we reason in accordance with, if any, is not simply a matter of our actual reasoning, since we might be misapplying the logic we endorse and otherwise use. Thus even the empirical question cannot simply be answered by looking at how we in fact reason. This is analogous to determining what the syntactic rules for our own language are. We can not simply look at what sentences people in fact utter, since people utter ungrammatical sentences, ones that are ungrammatical even by their own assessment on reflection. We need to look at actual performance as well as implicit competence. Which grammar is correct for our own language does not have to be, and clearly isn't, explicitly available to those who speak the language, but it must manifest itself nonetheless beyond what sentences speakers utter. For example, an ungrammatical sentence can bring with it a certain feeling of something being wrong, even if one can't quite say what it is more precisely. There is thus some felt push towards something different, and the same will apply when we consider the question which logic and which conception of validity is the right one for our reasoning. It is not enough to look at our actual inferences, which are clearly imperfect, but we need to also consider what pressures towards something different we implicitly accept and do not accept. The logic that corresponds to our deductive competence will manifest itself in our reasoning, and in the presence or absence of felt pressure towards something different. All this is still only part of the descriptive question

about our reasoning. The normative question, to be addressed later, adds to that whether we should reason this way.

When it comes to the descriptive question, there are two things to distinguish: First, which logic, if any, is our logic and thus which rules are the valid rules of our logic. Second, which conception of logic, and thus which notion of validity, governs our logic. As mentioned at the outset, I do not hope to settle the first question here, but I take classical logic, including the naive truth rules, to be a plausible and correct answer. Classical logic, in particular, is the logic used in mathematics, and mathematics is the pinnacle of pure deductive reasoning. But the general proposal of this paper does not depend on an agreement with this starting point, and would apply to many, but not all, other starter logics as well. My focus will instead be on the second question: is the logic that we employ in our deductive reasoning one based on strict validity or generic validity. That is the descriptive question we should focus on. And to settle it we need to look not primarily at which inference forms we reason in accordance with, but what rational pressures we feel in our reasoning. This is what pulls the two conceptions of logic apart. And looking at it this way, things look rather good for the generic conception of validity, and rather bad for the strict conception.

On the generic conception of validity, the ordinary reaction to paradox is exactly the right one. When the Curry paradox is presented to an ordinary person, not necessarily a professional philosopher or logician, they accept each and every step of the reasoning, they accept, upon reflection, the forms of reasoning relied upon, but they do not accept the conclusion so reached, even after, and especially upon, reflection. In particular, upon realizing that one could have argued the same way that Santa doesn't exist, the conclusion loses all its support from the argument. Despite this, the standard reaction after being exposed to the paradox is not to change one's reasoning in other contexts, to stop relying on modus ponens or even to have a sense of concern when using modus ponens elsewhere. Instead one shrugs off the paradox with little felt pull towards revision of one's reasoning practice, while one at the same time one does not accept the conclusion reached with an argument one otherwise finds to be flawless. To put it differently: among ordinary reasoners there is little sense of rational pressure to give up one of the rules of inference in light of the paradoxes.

This is quite remarkable, and very significant for our discussion: Even though the paradoxes have been known for at least 2500 years, reliance on the classical rules and the naive truth rules has faced little resistance among ordinary reasoners. Little rational pressure is felt towards

giving them up. The paradoxes are simply put aside, and reasoning goes on as it did before. We continue to reason in accordance with just the rules that figure in the paradoxes, we in general accept the reasoning with those rules, but we simply disregard the paradoxical reasoning and don't accept the conclusions derived with these arguments, once we see what is going on. And we feel no or little rational pressure to change our ways, except, of course, for a dedicated group of philosophers. And they have good reason to look for the bug: which rule is not valid, or why is this instance not a legitimate one? Assuming the strict conception of deductive logic, there must be a bug, and we must find it. But the ordinary reaction to paradoxes, that of ordinary reasoners, points towards this being mistaken. The philosophers reaction would be rational assuming the strict conception of deductive logic, but the ordinary reaction is evidence that this conception is mistaken. Philosophers, after all, are just ordinary reasoners with extra training, and that extra training can introduce mistakes that lead one astray.

The ordinary reaction is evidence for what our competence of deductive reasoning consists in, and with it evidence of generic validity being the notion of validity that guides our deductive reasoning. This is no different than ordinary reactions to utterances of sentences being evidence for what the grammar of our language is. On the default conception of deductive logic, as outlined above, the ordinary reaction to paradox is the rational one, and there indeed is no or little rational pressure to revise the rules in light of the paradoxes. Assuming our competence tracks these pressures, this speaks in favor of the default conception and in favor of generic validity. But on the strict conception of validity, a rule is refuted and shown to be invalid if we find a single counterexample. And a group of rules are jointly refuted if they together can lead us from truth to falsity. We know that on the strict conception of validity, not all the rules we in fact reason with are valid. Nonetheless we still reason with them. Defenders of revision generally like to claim that this is fine, we are entitled to reason with those rules anyways, even though some of them are not valid. But it is unclear what good reasons they have for claiming this. Although the standard pro-revision line is that revision is rational in general, it is nonetheless also rational to reason classically in mathematics, say, and to employ the naive truth rules in ordinary cases. But why is it rational to reason with rules which are not valid? Why am I entitled to the conclusions we draw in mathematics when at the same time I hold that some of the rules I relied on are not valid?

The general pro-revision answer is that the rules are truth preserving in mathematics and thus we are entitled to use them there, even though we are not entitled to use them in general.

In mathematics we are “in a consistent context” and so we may reason classically, or so the proposal. But I don’t see how this solves the problem. If being in a consistent context is a purely externalist condition for being entitled to use generally invalid rules, then it is hard to see how I can remain entitled to these rules once I recognize that they in general are not valid. After all, by my own lights I use rules that are invalid, and even if I am in a consistent context, I might not realize this. Things might change once I learn that I am in a consistent context. But then, isn’t “I am in a consistent context” merely another premise in my reasoning? And if so, how does having a new premise change which rules I am entitled to reason with? It seems rather that I am entitled only to reason with valid rules, and thus not classically, although I can learn more about my own situation and acquire more premises to reason with. But classical reasoning nonetheless is not saved, and with this the results we have achieved in mathematics are in limbo.

This contrasts nicely with deductive reason on the generic conception. One might object that if it were right, then our results in mathematics are also in limbo, since the rules allow for defeat and thus there is the possibility that what has been proven might be defeated later, and thus that we lose our entitlement to the results of mathematics. But it would be a mistake to think that the results are in limbo, since defeat is in principle possible. We are entitled to the results to the degree that deductive reasoning entitles us. That defeat is in principle possible for deductive reasoning does not take away from the entitlement we achieve with it by itself. Only actual defeat takes the entitlement away, and then it takes it all away for a particular argument. Thus nothing is in limbo when we reason properly. We are entitled to accept the relevant conclusions and the results of mathematics to the degree that deductive reasoning in general comes with. Such reasoning is defeasible, but if it is not defeated, then the conclusions are rationally to be accepted, and thus not in limbo.

5.2 The normative question

Even if our logic is in fact based on generic validity, the question remains if we should not better base it on strict validity. Here there are two lines of thought that should be distinguished. The first is that although we might reason in accordance with a generic notion of validity, we are not entitled to this reasoning, and thus we should stop engaging in it. The second is that even if we are entitled to this reasoning, there is a better alternative, namely moving to a strict notion of validity. Let’s look at them in turn.

Although it is popular to try to derive our entitlement to reason deductively in accordance with the rules of logic, in particular to do so by connecting it to the meaning of the logical expressions tied to these rules, this seems to me to be a mistake. My own preferred view is a dogmatist one instead, in the sense of [Pryor, 2000]. We are entitled, by default and defeasibly, to the results of our belief forming mechanisms, whichever ones we find ourselves as having. And thus we are entitled, by default and defeasibly, to reason in accordance with the deductive rules we find to be compelling. No theory of meaning or any philosophical work is required for this entitlement, we simply have it. But it is defeasible, and so it can go away. If the arguments given above in section 2 are correct, then this entitlement has not been defeated, even by the paradoxes, and thus we continue to have it. But the question remains if we should not change our ways nonetheless.

Even if our deductive competence is based on generic validity, strict validity makes sense, and with it the question whether it would not be preferable to reason only in accordance to strictly valid rules. This, too, is a normative question, one about not the revision of our rules directly, but about the notion of validity underlying those rules. Of course, if we switched to strict validity in reasoning, then we would have to change some of our rules, as a consequence. The question is whether this could be rational by our own lights. So, what reasons would we have for this switch?

One natural thought is that reasoning with generically valid rules is risky, since our entitlement might get defeated. It might thus be best to reason only with rules where this is ruled out, and which are thus strictly valid. But this misunderstands the default picture of deductive reasoning. It is not the case that we are entitled to a lesser degree when we reasoning with generically valid rules than with strictly valid ones. Both lead to the full entitlement of deductive reasoning. It isn't the case that once I recognize that deductive reasoning is defeasible that this takes away from or lessens my entitlement to the conclusions reached with deductively valid rules when they are not defeated. I shouldn't conclude that deductive reasoning itself only leads to a lesser degree of entitlement or warrant for its conclusions, instead I should hold that it has the high degree we associate with it when it isn't defeated, and no entitlement at all when it is defeated. Trying to find the strictly valid rules and then guiding one's reasoning only by them would not lead to an improvement in warrant, but to a more limited deductive power. It won't lead to a larger degree of warrant, since the generically valid ones already give you the high

degree of warrant we associate with deductive reasoning. Thinking that strict validity would give a higher degree of warrant, since defeat is ruled out, would be a mistake. Obviously, for the arguments that got defeated, all warrant is lost. But the worry that the possibility of future defeat somehow takes away from our entitlement to the conclusions where defeat does not occur would be just as misguided as the worry that maybe in the future people will accept different rules and thus I should take my present ones less seriously. Searching for strictly valid rules is a fine theoretical enterprise, but switching to strictly valid rules won't improve our deductive reasoning. To the contrary, if we switch to strictly valid rules, then we are forced to accept fewer rules, and thus have less deductive power. Since not all of the classical rules can be strictly valid, we would have to reason in accordance to fewer rules and thus our reasoning would have to be restricted. To switch from generic to strict validity would thus be a mistake.

Overall then, I conclude that we have good reason to think that our actual deductive competence is based on generic validity. Thus defeasible deductive logic is the proper logic that governs our deductive reasoning, not strict deductive logic. The rational pressures we see realized in ordinary reactions to paradoxes support this. In addition, if we adopt a dogmatist stance on the entitlement to our basic belief forming mechanisms, then we are defeasibly entitled to reason in accordance with these generically valid rules. And finally, we have no reason to switch to strictly valid rules instead, but reason to stick with our present practice. Generic validity is thus the correct notion of validity.

6 The solution to the puzzle

The puzzle about the rational revisability of logic was that there are apparently compelling arguments for and against revision. On the one hand there is the argument that our own logic can never be rationally revised by our own lights, since we can prove deductively that all our rules are valid. On the other hand there is the argument that logic must be revised, since the rules we accept allow us to deduce anything whatsoever. I argued in section 2 that attempts to defeat the argument that we can prove that our own rules are valid are not successful. We should instead take the other side and hold that the argument that we must revise our rules in light of the fact that they allow us to derive anything whatsoever should be rejected. This argument would be compelling if our deductive reasoning was based on strict validity: to reason in accordance with a rule, and to endorse this rule, requires us to endorse each and every

instance of the rule. To the contrary, I argued above that our actual deductive competence is in accordance with generic validity, not strict validity. It is this notion of validity that governs our actual deductive reasoning, and we are defeasibly entitled to reason with the rules we find ourselves as having. In light of this we are entitled to conclude that our rules are valid and truth-preserving. But we are not entitled to hold onto arbitrary conclusions, at least once we recognize that the paradoxical arguments allow us to derive both p as well as $\neg p$. Such recognition defeats our original entitlement to these conclusions, as per the proposed defeat rule for deductive reasoning. Thus we are not rationally required to change our deductive rules in light of the paradoxes, nor in light of the fact that we can derive arbitrary conclusions with those rules. The rational reaction instead is to not accept the arbitrary or contradictory conclusions, and to leave the rules in place as they are.

And this, of course, is what in fact has happened all along. We still reason classically and with the naive truth rules, even though we know of the paradoxes. This is how it should be and the argument that we can and need to change logic or the truth rules ultimately comes from two mistakes: one in the philosophy of logic and one in epistemology. The first mistake is to hold that logical validity is strict validity. Once that mistake is given up we can see that the paradoxes do not require a revision of the logical or the truth rules. The second mistake is to think that the rule-circular argument for the validity of our own logical rules can be defeated. Instead the argument that our own rules are valid is a good argument. We can prove that our own rules are valid, from no premises and relying only on deductively valid rules. No philosophical worries can outweigh the strength of these arguments, and thus logic is never rationally revisable. This is the solution to the puzzle I suggest we adopt. Our own rules are valid, even though we can derive anything with them. But we rationally leave these arbitrary conclusions aside, without any threat to our deductive reasoning with just those rules that allows us to derive them.²⁴

²⁴My thanks to Jc Beall, Aaron Cotnoir, Paul Egré, Hartry Field, Fabian Klinge, Matt Kotzen, Graham Priest, and Thomas Sattig for discussing these issues with me. Special thanks to Matti Eklund, Marc Lange, Ram Neta, Jim Pryor, Kevin Scharp, and Alex Worsnip for both discussions and comments on an earlier draft.

References

- [Beall, 2015] Beall, J. (2015). Free of detachment: logic, rationality, and gluts. *Noûs*, 49(2):410–23.
- [Boghossian, 2000] Boghossian, P. (2000). Knowledge of logic. In Boghossian, P. and Peacocke, C., editors, *New Essays on the A Priori*, pages 229–254. Oxford University Press.
- [Carlson and Pelletier, 1995] Carlson, G. N. and Pelletier, F. J., editors (1995). *The Generic Book*. University of Chicago Press.
- [Christensen, 2004] Christensen, D. (2004). *Putting Logic in its Place: formal constraints on rational belief*. Oxford University Press.
- [Cobreros et al., 2013] Cobreros, P., Égré, P., Ripley, D., and van Rooij, R. (2013). Reaching transparent truth. *Mind*, 122(488):841–66.
- [Dogramaci, 2010] Dogramaci, S. (2010). Knowledge of validity. *Noûs*, 44(3):403–432.
- [Dummett, 1978] Dummett, M. (1978). The justification of deduction. In *Truth and other enigmas*, pages 290–318. Harvard University Press.
- [Field, 2006] Field, H. (2006). Truth and the unprovability of consistency. *Mind*, 115:567–605.
- [Field, 2008] Field, H. (2008). *Saving Truth from Paradox*. Oxford University Press.
- [Field, 2009] Field, H. (2009). What is the normative role of logic? *Proceeding of the Aristotelian Society*, LXXXIII:252–68.
- [Field, 2020] Field, H. (2020). A "naturalistic" take on epistemology. *Journal of the American Philosophical Association*.
- [Harman, 1973] Harman, G. (1973). *Thought*. Princeton University Press.
- [Harman, 1986] Harman, G. (1986). *Change in View*. MIT Press.
- [Harman, 2009] Harman, G. (2009). Field on the normative role of logic. *Proceedings of the Aristotelian Society*, 109:333–5.
- [Hofweber, 2008] Hofweber, T. (2008). Validity, paradox, and the ideal of deductive logic. In Beall, J., editor, *Revenge of the Liar: New Essays on the Paradox*, pages 145–158. Oxford University Press.
- [Hofweber, 2016] Hofweber, T. (2016). *Ontology and the Ambitions of Metaphysics*. Oxford University Press, Oxford, UK.

- [Hofweber, 2019] Hofweber, T. (2019). Empirical evidence and the metaphysics of ordinary objects. In Cumpa, J. and Brewer, B., editors, *The nature of ordinary objects*. Cambridge University Press.
- [Israel, 1980] Israel, D. J. (1980). What is wrong with non-monotonic logic? In *Proceedings of the First National Conference on Artificial Intelligence*, pages 99–101. AAAI.
- [Kripke, 2011] Kripke, S. (2011). Two paradoxes of knowledge. In *Philosophical Troubles*, pages 27–51. Oxford University Press.
- [Lasonen-Aarnio, 2014] Lasonen-Aarnio, M. (2014). The dogmatism puzzle. *Australasian Journal of Philosophy*, 92(3):417–32.
- [MacFarlane, 2004] MacFarlane, J. (2004). In what sense (if any) is logic normative for thought? Unpublished manuscript.
- [Merricks, 2001] Merricks, T. (2001). *Objects and Persons*. Oxford University Press.
- [Nute, 1994] Nute, D. (1994). Defeasible logic. In *Handbook of logic in artificial intelligence and logic programming*, volume 3. Oxford University Press.
- [Pelletier and Asher, 1997] Pelletier, F. J. and Asher, N. (1997). Generics and defaults. In van Benthem, J. and ter Meulen, A., editors, *Handbook of Logic and Language*. Elsevier.
- [Pollock, 1974] Pollock, J. (1974). *Knowledge and Justification*. Princeton University Press.
- [Priest, 2014] Priest, G. (2014). Revising logic. In Rush, P., editor, *The metaphysics of logic*, pages 211–223. Cambridge University Press.
- [Pryor, 2000] Pryor, J. (2000). The skeptic and the dogmatist. *Noûs*, 34(4):517–549.
- [Reiter, 1980] Reiter, R. (1980). A logic for default reasoning. *Artificial Intelligence*, 13:81–137.
- [Rosen and Dorr, 2002] Rosen, G. and Dorr, C. (2002). Composition as a fiction. In Gale, R., editor, *The Blackwell guide to metaphysics*, pages 151–174. Blackwell.
- [Russell, 2015] Russell, G. (2015). The justification of the basic laws of logic. *Journal of Philosophical Logic*, 44:793–803.
- [Scharp, 2013] Scharp, K. (2013). *Replacing Truth*. Oxford University Press.
- [Sider, 2013] Sider, T. (2013). Against parthood. In Bennett, K. and Zimmerman, D., editors, *Oxford Studies in Metaphysics*, volume 8, pages 237–293. Oxford University Press.

[Steinberger, 2017] Steinberger, F. (2017). The normative status of logic. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, spring 2017 edition.

[Steinberger, 2019] Steinberger, F. (2019). Three ways in which logic might be normative. *Journal of Philosophy*, 116(1):5–31.