

How tolerant can you be? Carnap on rationality*

Florian Steinberger

Penultimate Draft

Forthcoming in *Philosophy and Phenomenological Research*

October 7, 2014

Abstract: In this paper I examine a neglected question concerning the centerpiece of Carnap's philosophy: the principle of tolerance. The principle of tolerance states that we are free to devise and adopt any well-defined form of language or linguistic framework we please. A linguistic framework defines framework-internal standards of correct reasoning that guide us in our first-order scientific pursuits. The choice of a linguistic framework, on the other hand, is an 'external' question to be settled on pragmatic grounds and so not itself constrained by these (framework-internal) standards. However, even if choosing a framework is a practical matter, we would nevertheless expect the process of framework selection to be subject to rational norms. But which norms might those be? And where do they come from? I begin by showing that these questions are crucial to the success of Carnap's entire philosophical project. I then offer a response on behalf of the Carnapian which guarantees the rationality of the process of framework selection, while remaining true to Carnap's firm commitment to tolerance.

*I am grateful to Steve Awodey, Patricia Blanchette, André Carus, Richard Creath, Branden Fitelson, Hannes Leitgeb, Erich Reck, Georg Schiemer, Marta Sznajder, Neil Tennant and Clinton Tolley for helpful conversations and comments. I am especially thankful to Michael Friedman who provided invaluable feedback on earlier drafts of this paper, and for the extremely insightful comments of an anonymous referee. Their input deepened my understanding of the topic and led to significant improvements. Needless to say, I am solely responsible for any remaining errors. Work on this paper was funded by the Alexander von Humboldt Foundation whose support I gratefully acknowledge.

1 Introduction

Carnap's 'most original and fundamental philosophical move' (Friedman 1999a, p. 169), consists in abandoning the 'absolutist' conception of logic—the notion that there is but one true *Logic*. Unlike his teacher Frege, for whom the logical laws were beholden to an objective reality, Carnap jettisons the notion that they should be answerable to an external standard against which they could be judged to be true or false, correct or incorrect. Against Wittgenstein's universalist conception of a single logic setting the inescapable bounds of the thinkable in the *Tractatus*, Carnap advances a voluntarist pluralism about logic enshrined in his celebrated *principle of tolerance*, according to which¹

everyone is at liberty to build his own logic, i.e. his own form of language, as he wishes. All that is required of him is that, if he wishes to discuss it, he must state his methods clearly, and give syntactical rules instead of philosophical arguments (Carnap 1937, §17)

Carnap's picture is this. Logics (and rational principles more generally) come as part of a linguistic framework. We are 'entirely free' (*völlig frei*, xiv) to devise any framework we please. None of the countless possible logico-linguistic frameworks has any exclusive claim to being correct. Indeed, the question of correctness cannot be so much as meaningfully posed in the case of linguistic frameworks.

But if we cannot choose linguistic frameworks on account of their correctness, on what grounds do we decide which framework to adopt?² Carnap's reply is that the choice of a linguistic framework is a matter of theoretical expedience: in the absence of external standards of correctness for logical frameworks, our theoretical needs dictate which framework ought to be selected. The process of deliberation over which framework to select becomes a 'practical not a theoretical question' to be settled on pragmatic grounds (in terms of simplicity, fruitfulness, convenience, etc.)

This much, I take it, is clear. But here is the rub. According to Carnap's voluntarism, rational norms are self-imposed. We impose such norms upon ourselves

¹For a helpful account of Carnap's intellectual development leading up to his espousal of the principle of tolerance and of the role that Wittgenstein's universalism occupies therein see Awodey and Carus (2009) and Friedman (1999b).

²As an anonymous referee has rightly pointed out, it only becomes necessary to be explicit and precise about the linguistic framework one uses when one is engaged in an exchange with other rational agents. Accordingly, the business of constructing and choosing a linguistic framework is typically a collaborative effort. All the same, for simplicity's sake I often write as if the choice of a linguistic framework was the product of a deliberative process of an individual agent. Presenting matters in individualistic terms does not affect my arguments below.

by adopting a linguistic framework that gives rise to them. But on what basis do we choose a linguistic framework? Surely, our choice, even if it is of a pragmatic nature, can nevertheless not be extrarational. It may be appropriate for you to criticize me on rational grounds, if not, perhaps, for choosing the theoretical ends I do, then at least for choosing a linguistic framework that is ill-suited for the ends I have set myself. That is to say, the process of framework selection ought itself to be guided by rational norms of some sort. But if rational norms are always relative to a linguistic framework for Carnap, it becomes difficult to see how the process of framework choice can itself be rational. This problem strikes me as fundamental to Carnap's entire tolerance-based philosophy. Curiously, to my knowledge neither Carnap himself nor his readers (including recent Carnap scholars) have considered the question in writing. My goal here is to clearly articulate the problem of the rationality of framework choice and to offer a Carnap-friendly solution to it.

The plan is as follows. Sections §2 and §3, set the scene by offering brief explanations of Carnap's conception of linguistic frameworks and of his distinction between the theoretical and the practical domains. §4 identifies the threat of relativism faced by Carnap's position and proposes a solution which, I argue, is both compatible with Carnap's mature position and independently plausible. In §5 I raise a potential regress problem for my account. I then consider and eventually reject a 'foundationalist' response. Finally, §6 advances a position I dub 'soft framework-foundationalism' that dodges the regress worries.

2 Linguistic frameworks

A linguistic framework, technically, is simply a vocabulary along with a set of formation rules and a set of transformation rules. It is the job of formation rules to determine which grammatical constructions are permissible. It is the job of the transformation rules to specify which types of inferential moves are permissible. Transformation rules fall into logico-mathematical *L-rules*³ and (possibly) extra-logical *P(hysical)-rules*, which embody principles involving *descriptive* (as well as *logical*

³It should be noted that Carnap's conception of logic is rather more encompassing than what most contemporary authors would be comfortable with. At the time Carnap was writing the *The logical syntax of language* it was still more or less commonplace to include set theory under the heading of logic. See e.g. (Quine 1994, p. 347) Nowadays, most would shirk from lumping together inference rules for the standard connectives and the axioms of induction and of choice under the heading of 'L(ogical)-rules', as Carnap does (Carnap 1937, §30). That is not to suggest, of course, that there is currently anything like a consensus about what falls under the remit of logic.

vocabulary).⁴ The L-rules are central. For they generate a relation of logical consequence and so set out the ground rules that

define standards for the acceptance and the rejection of sentences and theories formulated within the calculus and [so define the] standards for a language-relative notion of cognitive correctness (Ricketts 2007, p. 206).

How does this work? Linguistic frameworks are evaluated on the basis of their potential to serve as suitable (local or global) languages of science. In order for a linguistic framework to perform this function, it must be ‘coordinated’ with our sensible experience by identifying (or adding, if need be) a descriptive vocabulary that includes observation predicates (O-predicates). The descriptive vocabulary receives its meaning from the P-rules that govern it and from its logical connections with the O-predicates. Rather than assuming, as Frege did, a pre-existent realm of senses and thoughts that one’s vocabulary would then have to be mapped onto by means of necessarily imprecise verbal ‘elucidations’, Carnap starts with a purely syntactic specification of the language, which receives meaning through the said process of coordination with our theoretical principles and our observations.⁵ It is thus by forging connections between P-rules, the mathematical machinery and our empirical data that logic contributes to endowing the descriptive vocabulary with meaning. As Alan Richardson aptly puts it:⁶

logic provides the formal conditions for sense-making. Suppose we wish to know the reason the sky is blue. The object for which we want the reason is the sentence ‘the sky is blue’; a theoretical reason is then another sentence within the same language from which our target reason logically follows. The very notion of a theoretical reason, therefore, makes sense only internal to a logical framework. Thus there is no realm of theoretical reasons that can be appealed to in advance of the adoption of a logical system. [...] The adoption of some logical system is necessary for there to

⁴I am deliberately skirting a number of intricacies that are inessential to my point here. See Bonnay (2009), Friedman (1999a), Koellner (Unpublished) and Ricketts (2007) for helpful discussions of some of the issues involved.

⁵I am describing Carnap’s view during his syntactic period, relying mainly on Carnap (1936 and 1937) and Carnap (1937). After Carnap’s shift to semantics it becomes possible to specify the relations between the appropriate syntactic items and their *designata* in the metalanguage explicitly. I take it that the transition to his semantic period does not change anything about the constitutive role of logic in determining meaning. See Carnap (1939).

⁶See Creath (2007), Friedman (1999a), Ricketts (2007) for accounts that similarly attribute a constitutive role to linguistic frameworks in scientific inquiry.

be a notion of evidence or theoretical reason in the first place (Richardson 2007, p. 300–301).

Thus, a logical framework is necessary for there to be theoretical reasons and hence the possibility of making and evaluating meaningful claims at all. It is in this sense that logic ‘provides the formal conditions of sense-making’. It is only once we have specified what follows from what, which statements are logically incompatible with which other ones, and so on, that we determine the truth-conditions and hence the meanings of the sentences in our language.

Thus, the norms set forth by one’s adopted logic (or the L-rules of one’s preferred linguistic framework, if you prefer) have a genuinely constitutive function. They create the very possibility of truth-evaluable and so of cognitive discourse. For only a system of linguistic representation with *some* such logical structure can be said to give rise to meaningful theoretical questions, questions regarding the truth or falsity of claims.

We thus arrive at a partial answer to our question as to why linguistic frameworks are not in the business of being correct or incorrect: notions of correctness, of truth, and so on can only be made precise *within* a framework that provides the necessary normative framework. However, when we choose between competing frameworks, we seem to stand ‘outside’ of a framework. It is for this reason that the choice of a linguistic framework is not itself a theoretical, but rather a practical question. To fully understand this, we must therefore get clearer on Carnap’s distinction between the theoretical and practical. It is to this question I now turn.

3 The theoretical and the practical realms

What, then, is a theoretical question and how does it distinguish itself from a practical question? And what does this mean for the issue of framework choice? Now, theoretical questions concern subject matters admitting of truth and falsity and so are settled via either empirical or logico-mathematical inquiry or, typically, some combination of the two. Inasmuch, as we have just seen, as questions of correctness, of truth and falsity, can only be meaningfully posed *within* the context of a linguistic framework in which the relevant concepts have been clearly defined (whence, as we have seen, the constitutively normative role of linguistic frameworks), a theoretical question will always be a question that can be settled in a suitable framework. In other words, a theoretical question is an *internal* question with respect to some

framework.⁷ By contrast, practical decisions conclude deliberations over what to *do*. They are concerned not with determining what is true, but with deciding how to act. According to Carnap, the labels of ‘truth’ and ‘falsity’ are misapplied here.⁸ For instance, no amount of theoretical knowledge about the world can *answer* the question whether I should go to law school. Whether I should go depends on what my goals are; for instance, if I have decided to devote my professional life to fighting injustices in the criminal justice system; or if I wish to ensure I have a sizable income; or again if I have my mother’s *naches* in mind, then perhaps I should go to law school. On Carnap’s view the adoption of goals is ultimately an exercise of the will that is not determined by theoretical reasons.⁹

This is not to deny, of course, that theoretical knowledge typically has an extremely important role to play in our quotidian practical deliberations. After all, it only makes sense for me to go to law school given my goal of leading a financially comfortable life, if I have good reasons for thinking that obtaining a law degree will increase my chances of earning a respectable income. However, for Carnap, theoretical knowledge never determines the ends but only ever the means to those ends (Carnap 1934).¹⁰

Given Carnap’s sharp distinction between practical and theoretical questions, we arrive at a picture in which linguistic frameworks define the rules that first make it possible to formulate theoretical questions (as internal questions), which can be settled (at least in principle) by appeal to the norms set by the framework as well as, perhaps, through empirical investigation. However, the choice of the rules themselves poses a fundamentally different kind of question, to be settled by appeal to different kinds of considerations. Ricketts usefully compares the opposition between theoretical (and hence internal) and practical (and hence external) questions to that

⁷Carnap famously explicitly introduces the distinction between ‘internal’ and ‘external’ questions in his (Carnap 1950). However, as Friedman (2001, p. 18, fn. 18) notes, it already, albeit implicitly, plays a central role in Carnap’s discussion of debates in the foundations of mathematics in *The logical syntax of language*.

⁸In cases in which practical decisions have moral valence, Carnap’s view famously amounts to a strong form of non-cognitivism about moral statements; e.g. the indicative statement: ‘I ought to save the drowning man’ is not susceptible of truth and falsity.

⁹See Kaplan (1963) and Carnap (1963b).

¹⁰That being said, Carnap would presumably agree that knowledge of the necessary means to one’s ends can lead one to adopt different ends in some cases. For instance, upon realizing that all the professions that promise the kind of income level I am after prove to be too onerous or vapid for me, I may well decide to modify my goals in light of this realization. André Carus (2014) offers a discussion of this ‘feedback relation’ between theoretical knowledge and the ‘realm of practical decisions’ in Carnap. His reading, which he attributes to Howard Stein, strikes me as germane to the view advanced here.

between questions that can be settled by appeal to the rules of a game and the question of the choice of the rules themselves:

The choice of a formal language as the language of science thus resembles the choice of rules for a game we intend to play. The rules define what is permitted in the course of the game. But no such question of legitimacy applies to the choice of these rules themselves. Rather, we construct the rules for the game by weighing various considerations in order to frame rules that define, say, an enjoyable, engaging, competitive activity whose course depends on a mixture of physical skill, quick wits, strategy, and luck. We may, in the light of our experience playing the game, decide to modify the rule, perhaps in far-reaching ways, and so to play a different game. Similarly, the option to change languages is always open as regards the choice of a language of science (Ricketts 2007, p. 207).

So, linguistic frameworks define *internal* standards of acceptability for our modes of inference; however, our choice of the framework itself is an *external* matter that is not subject to those standards or indeed to any ‘theoretical’ standards at all. Rather we are guided by our theoretical ends; we seek the optimal or at least a satisfactory linguistic framework given those ends and perhaps other theoretical virtues (simplicity, conservativeness, etc.).¹¹

4 From framework-relativity to relativism?

That, I take it, is Carnap’s position. In abandoning the logical absolutism of his forefathers, Carnap makes rational norms framework-relative. Each framework comes with its distinct set of norms. We impose these norms on ourselves by adopting the framework. And we adopt the framework through a process of selection which is not itself regulated by the framework-internal norms, since they are applicable only within a framework. But—and herein lies the worry—if standards of rationality have no bearing on external, practical questions, then it starts to look as if the process of framework selection itself is not subject to any norms of rationality. And this seems like a worrying conclusion. Yet, it is a conclusion that Ricketts appears happy to ascribe to Carnap when he claims that

¹¹The definite article preceding ‘optimal’ is to be taken with a grain of salt. As in the case of systems inductive logic, it is plausible that there is no one framework that best negotiates these demands. It may well be the case that many frameworks perform equally well.

there can be no method or set of rules which dictates to us a choice of a framework. A linguistic framework is chosen on extrarational pragmatic grounds on whose application there need be no antecedent agreement: convenience lies in the eyes of the beholder (Ricketts 1982, p. 120)

But is such a radical framework-relativism really acceptable to Carnap? Is it independently plausible? So far I have simply assumed an affirmative answer to both questions. The time has come to address them. Begin with the former question.¹² On the face of it, saying that framework selection is practical and external is one thing; saying that it is not guided by any norms of reasoning and hence ‘extrarational’ is quite another. After all, even if frameworks cannot be said to be correct or incorrect, it seems that one framework choice may be better than another given our theoretical aims. Why would framework choice be any different from other types of practical decisions? And since ordinary practical decisions are subject to norms of instrumental rationality, so should decisions over which frameworks to select. Would Carnap be on board with this? I cannot see why not. We already noted that he conceives of framework choices as practical decisions made with a view to promoting a certain end. Less obviously, Carnap also accepts that there are norms of instrumental rationality. Indeed, a good deal of his work during his ‘probabilistic period’ is devoted to the analysis of decision problems, in which one examines which choice rules best capture what it would be rational for an agent to do (in light of her goals and given what she has reason to believe).¹³ Moreover, in an unpublished fragment entitled ‘Wertebegriffe’ of 1958, which was originally intended to be part of Carnap’s reply to Kaplan in the Schilpp Volume, Carnap in fact explicitly defines instrumental rationality—which he calls ‘relative rationality’—in decision-theoretic terms. In modern terminology, actions are ‘relatively rational’ to the extent that they maximize expected ‘value’ or utility (Carnap 1958, p. 3).¹⁴ Thus, Carnap’s position seems not only perfectly compatible with the idea that framework choice, like other types of practical decisions, should be bound by norms of instrumental rationality, it looks as if Carnap is in fact committed to it. As will emerge in section 6, however, the Carnapian notion of instrumental rationality at play here turns out to be a rather thin one.

Is framework-relativism independently plausible? I think not. Even if we grant Carnap’s non-cognitivism to the effect that our ultimate goals and values are not dictated to us by theoretical reasons, it seems that we may nevertheless be rightfully criticized on rational grounds for selecting a linguistic framework that is ill-suited for

¹²I am grateful to an anonymous referee for pressing me on this point.

¹³See e.g. (Carnap 1962, §50-51).

¹⁴See (Carus 2014, p. 13) for further discussion of this fragment.

the end being pursued.¹⁵ Take Carnap's example from *The logical syntax of language*. If my goal is the formulation of a physical theory, then (bracketing foundational issues of consistency and constructivity) I would be ill-advised to avail myself of Carnap's language I (a form of primitive recursive arithmetic) for that purpose. It would indeed be meaningless to call the framework 'incorrect': the framework is not incorrect; it is merely inconvenient because it would weaken and needlessly complicate our theory. But it seems equally clear that my choice is sub-optimal. If we agree on the end being pursued, you would be right to call me out for my mistake, just as you would be right to point out to me that I am making a poor decision if, in planning a trip, I choose the longer of two travel itineraries (assuming that we both know that my goal is a swift journey and that there are no other relevant considerations). Denying that such choices can be subject to rational criticism is not only an affront to our intuitions about such cases, in the context of Carnap's account it also encourages an unpalatable relativism about scientific discourse in general.¹⁶ If no one were potentially subject to rational criticism for their framework choices, the only motivation for different researchers to agree on a linguistic framework would presumably be extrarational (perhaps sociological), which seems troubling. I submit, therefore, that the process of framework selection, though an external and a practical question, must nevertheless be responsible to norms of rationality. We are now finally in a position to formulate what I described as the central question about Carnapian tolerance above:

What is the status and the nature of the meta-norms of rationality that guide our choice of linguistic frameworks?

One possible response would be to maintain that the process of framework selection involves various heuristics that defy formalization and indeed any form of rational reconstruction. Practitioners simply have the ability to make the right kind of judgments about the overall suitability of linguistic frameworks: they have a capacity to weigh and balance the relevant *desiderata* such as 'convenience, fruitfulness, simplicity', etc. (Carnap 1963a, p. 66). Scientists—to the extent that scientists can be said to operate on the basis of anything like linguistic frameworks in Carnap's sense at all—have the capacity to make such calls, and that's that. There simply is no way to spell out the principles and norms that implicitly guide the exercise of this capacity. This response seems problematic for two reasons: one philosophical, the other interpretational.

¹⁵Richard Creath (2007, 2009) makes this point very clearly.

¹⁶Cf. (Friedman 2001, p. 57).

Philosophically, there just is something deeply unsatisfactory about claiming that it is not possible to offer a rational reconstruction of the process by which scientists choose linguistic frameworks without offering any explanation as to why that should be. If framework choice really is central to scientific practice, it would be all the more worrisome if it were an unanalyzable know-how comparable to chicken sexing. If, on the other hand, Carnap's conception of linguistic frameworks is itself better thought of as an idealized reconstruction of scientific practice, it seems equally problematic that no account of how frameworks ought to be selected can be offered.

As far as the second, interpretational issue is concerned, the position that the norms guiding framework choice should somehow resist formalization and so be essentially informal sits very uneasily with both Carnap's philosophical methodology and his overarching philosophical aims. The urgency with which Carnap calls for greater rigor in philosophy—in particular, his demand that we 'explicate' fuzzy terms and render them precise enough to make associated questions scientifically tractable—makes it clear that it is a central methodological commitment of Carnap's philosophy. Our reading of Carnap should not, therefore, place an irremediably informal notion at such a central point in Carnap's system. It is similarly at odds with Carnap's philosophical projects. He was convinced early on that philosophical problems are not 'first-order' problems 'about the world', but rather concern our systems of linguistic representations thereof. This insight is at the heart of Carnap's revolutionary methodological reorientation away from traditional philosophizing in favor of his *logic of science*, a meta-theoretical logical analysis of the language(s) of science. Much of Carnap's work from the early thirties onward is concerned with laying the groundwork for a 'general theory of linguistic forms', which essentially consists in making

metalanguage[s] more precise so that an exact conceptual system for meta-logic could be constructed in it (Carnap 1963a, p. 54).¹⁷

It would be incongruous, to say the least, that after expending such considerable efforts on the development of metalinguistic tools allowing us to make our theoretical proposals precise in the form of rigorously defined linguistic frameworks, the choice between linguistic forms would boil down to a mysterious and brute practical ability. After all, the demand for precision in formulating our framework stems precisely from the need to make accurate predictions and assessments of the consequences of adopting such a framework. Nevertheless, as things stand, we are completely in the dark as to how this rich body of detailed information about the proposed framework should filter into and ultimately help determine our choice of a framework.

¹⁷See also Carnap (1937, p. xiii)

The upshot of these considerations is that framework choice must be governed by some set of rational norms that it must be possible to identify. Very good. But what, then, are these norms? What shape do they have? One thought might be this. Based on what we have seen, the Carnapian picture does not appear to leave room for norms that float free of any linguistic framework—norms are always framework-relative. Therefore, if the process of framework selection is subject to norms of rationality, those norms must have their source in some linguistic framework. We thus arrive at the chiasmic slogan that *framework selection happens in a selection framework*. The *selection framework* is the linguistic framework that articulates the standards against which framework choices are to be appraised.

But this line of thought seems to run up against a major obstacle. The objection is that the very idea of a selection framework is fundamentally at odds with Carnap's account. For if there were a selection framework within which to settle the question of which framework to choose, then that would *ipso facto* transform the question into an internal question. But, by definition, the question of framework choice is not an internal question but an external, practical question. Does this not show that the notion of a selection framework cannot be made intelligible within a Carnapian account?

I think not. And here is how I think the Carnapian should stave off the apparent contradiction. The Carnapian needs to explain how the availability of a selection framework is compatible with both the practical and the external character of the question of framework selection. The reason why an appeal to a selection framework does not undermine the practical nature of the question is that selection frameworks *do not* settle the practical question, 'Which form of language ought I to accept?' They settle the *theoretical* question, 'Which linguistic framework is most fruitful given my goal G (where G is fixed antecedently)?' In other words, a selection framework does not dictate which ends I ought to pursue; it merely specifies rules governing my means-ends reasoning *once* I have decided to pursue that goal. The problem can thus be dissolved by divorcing the properly practical component of deciding which goals to pursue, from the theoretically tractable question as to which means best serve to bring about these ends. And this manoeuvre of splitting the question of framework selection into a properly practical and a quasi-theoretical part is perfectly compatible with Carnap's mature position. In his 'Reply to Kaplan', Carnap explicitly states that his non-cognitivism about evaluative language does not extend to 'statements on means-ends relationships' (or statements 'on the utility of a possible event [...] for a person'). Statements such as these, though 'connected with values or valuations[,] are clearly *factual*' and hence truth-apt.¹⁸ But if questions of means-ends reasoning

¹⁸(Carnap 1963b, p. 999)

are factual and so theoretical, they must be internal questions with respect to some framework. What kind of framework? A selection framework, of course!¹⁹

We can take a further step to dispel the sense of contradiction by distinguishing different levels among linguistic frameworks. First-level frameworks are proposals for the language of science (or sublanguages thereof); second-level frameworks are proposals for the sets of rules that ought to govern our choice of first-order frameworks; third-level frameworks are proposals for the sets of rules that ought to govern the choice of second-order frameworks, and so on up for higher-level frameworks.²⁰ The means-ends portion of the question of choosing a (first-order) linguistic framework (though not the determination of the end itself) could then still be viewed as an external question *with respect to* first-order frameworks (or, more generally, lower-level frameworks). However, it would be an *internal* question with respect to the appropriate second-order framework (i.e. higher-order framework).

In this way we can make sense of the fact that there are norms that guide our decisions about which first-level frameworks to choose. And we retain the central Carnapian tenet that norms are framework-relative. This gives us (the sketch) of a model of how we should think about meta-norms of framework selection. So far, perhaps, so good. But is there anything more we can say about what these norms amount to? Again, Carnap barely offers so much as a clue, with the exception of the remark that the question of framework selection is not ‘a question simply of yes or no, but a matter of degree’ (Carnap 1950, p. 39). The norms of our selection framework, therefore, must reflect this and allow the factors that enter into our evaluation of a framework to admit of degrees: frameworks can be more or less simple, convenient, powerful, etc.

How might we model this procedure? The candidate that comes to mind is decision theory. Not only does decision theory naturally lend itself to the task at hand, it also is the focus of a great deal of Carnap’s work in probability theory (his lengthy treatment in (Carnap 1962) is a good example of this). Late in his career he actually defines a notion of practical rationality on a decision-theoretic basis in an unpublished manuscript (Carnap 1958) (originally intended as part of (Carnap 1963b)).²¹ Let us pursue the idea of applying decision theory as a model for framework selection.

¹⁹The analysis is also consonant with our discussion of Carnap’s distinction between theoretical questions and practical decisions in section 6.

²⁰Of course no intrinsic feature of a framework destines it to being a framework of a particular level. It may be that a framework devised for ground-level applications lends itself to higher-level selection framework duties and vice versa. What matters is how the frameworks are applied.

²¹That being said, to my knowledge Carnap himself nowhere proposes to model framework selection by decision-theoretic means.

Among other machinery, our selection framework would then comprise the following familiar elements: a set of worlds W , a probability measure pr over W , a set of actions A (for a given subject) and, finally, a utility function u , which specifies a particular outcome for every action relative to a state. On the basis of these components we can then calculate the expected utility (EU) of an action $a \in A$ in the customary way:

$$EU(a) = \sum_{w \in W} pr(a)u(a, w)$$

As for our choice principle, let us simply say that an agent should choose to act in such a way as to maximize expected utility.

How, now, must we reinterpret these decision-theoretic notions in order to be able to co-opt them for the purpose of constructing a normative framework for choosing linguistic frameworks? Well, the ‘action’ in question is that of choosing a linguistic framework. Frameworks are appraised along various (and variable, depending on the ends pursued) dimensions, e.g. for their simplicity, convenience, etc. We may assume that all of these (presumably weighted) criteria are agglomerated in our utility function. A framework’s utility is thus an aggregate of the ‘scores’ it achieves with respect to each of these constraints. It may be very flexible in terms of which constraints are considered and how they are to be weighted. We might model a framework’s score as follows.²² Let $C = \{c_1, \dots, c_n\}$ be our set of criteria. For each of our criteria, there is a function μ_{c_i} , which measures how well a framework performs with respect to criterion c_i relative to a fixed goal. The scale is largely conventional. We may simply suppose our measure functions to take their values from the unit interval. We have also said that our criteria need not be equally important; they can be differently weighted. Let g_{c_i} be a weight given to criterion c_i and assume that our weights also take real values between 0 and 1. The pragmatic component of our utility function can then be thought of as the weighted sum of its scores for each of the criteria. For the linguistic framework L we thus get:

$$v(L) = \sum_{c_i \in C} w_{c_i} \mu_{c_i}$$

But there is another important aspect to utility functions: they specify payoffs *relative to how the world is* (in relevant respects). In our case, our utility function will determine the outcome of selecting a particular linguistic framework relative to a world. How is this to be understood? Well, in general a framework’s pragmatic aspects—simplicity, convenience, etc., which are captured in our function v —are not

²²I am adapting a simple model for theory selection due to Graham Priest (2006, Ch. 8) here.

all that matters for evaluating a linguistic framework. As Carnap puts it, we must also look at the *consequences* of adopting a framework, its fit with the available data and its predictive success. The framework's success on this front will obviously depend on how the world is. A linguistic framework that allows us to talk about material objects is useful in a world like ours; it would be less useful were everything in our world in constant flux, as it is on the surface of the sun. And this brings us to our probability function: this function tells us how probable it is (or is deemed to be by the agent) that the world is thus and so; for example, that it is improbable that the world should, everywhere and always, be like the surface of the sun.²³ An agent then ought to select the linguistic framework that, in sum, has the greatest utility across all likelihood-weighted worlds. In some cases neither the probability function nor the relevant set of worlds will enter the picture. This will be so whenever purely analytic frameworks are under consideration.

5 Regress or not to regress

Let us run with the idea of a decision theory-based selection framework. Given this account of the meta-norms guiding our choice of framework, we now face the obvious follow-up question: 'Does the principle of tolerance apply to selection frameworks in just the same way it applies to first-order frameworks?' In other words, is it the case that our decision theory-based selection framework is just one among a multiplicity of (second-level) frameworks with nothing but pragmatic reasons to choose between them?

At first blush, it seems that this is an assumption we should make if we want to stay true to the Carnapian spirit. The trouble with this assumption, however, is that it appears to lead to an infinite regress. We arrive at it via the following reasoning. Framework selections must themselves be carried out in accordance with principles laid down by a selection framework. Is there but one selection framework? No, the selection framework, we are assuming, is subject to the principle of tolerance. Thus the selection framework is itself selected among a number of candidates on pragmatic grounds. But surely, the process of selecting our selection framework, too, must be subject to certain norms of rationality. Where do *these* norms come from? They must in turn be provided by a linguistic framework, a selection framework for selecting selection frameworks. But surely *this* meta-selection framework too is

²³The probability function would presumably be a logical probability for Carnap. Late in his career, though, he became increasingly sympathetic to subjective probabilities. This issue need not detain us here.

subject to the principle of tolerance, and so we go up the levels.

What should we say about this kind of regress? We could try to block it by adopting what we might call *framework foundationalism* (or just *foundationalism* for short). The most straightforward and perhaps initially most plausible way to block the regress is to insist that while tolerance applies to first-order frameworks, it does not extend to our selection framework. We thus in effect ascribe to our selection framework a regress-blocking foundational role. But what can be said in favor of attributing such a privileged status to it? It might be claimed that framework choice *just is* decision-theoretic in nature. And since standard decision theory *is*, by and large, the only game in town, our selection framework, which is based on standard decision theory, is equally unique. In other words, unlike the case of logico-mathematical formalisms in first-order frameworks, where there are any number of options to choose from (classical, intuitionistic or quantum logical systems, to name but a few), there simply are no serious alternatives to standard decision theory and hence to our selection framework.²⁴

To bolster her claim, the proponent of this foundationalist line might point to various forms of representation theorems.²⁵ Representation theorems of the kind in question ‘demonstrate the correctness’ of decision theory by showing that agents who fail to be expected utility maximizers must have preferences that are, in a sense to be specified, rationally defective.²⁶ So, in stark contrast with the case of first-order linguistic frameworks, not only can the question of correctness be raised, it can also be answered, since there is but one demonstrably correct second-order selection framework. If this approach is on the right track, it would deliver a principled delineation between first-order frameworks and the unique selection framework and so block the regress.

But how plausible is this foundationalist line of thought? Not very, I would argue. First of all, it seems dubious, to say the least, that there really are no

²⁴Notice that despite the uniqueness that our decision-theoretic selection framework enjoys on this view, it still retains a great deal of flexibility. Two deliberators might agree on what the world is like in some relevant respect—i.e. they might share the same probability function—and yet they might operate with different utility functions because they pursue different goals and so put a higher premium on one or another component of the utility function (fruitfulness as opposed to immunity to inconsistency, say) or because they consider different components altogether. Conversely, they might agree on a utility function but differ with respect to their probability assessments.

²⁵Buchak (forthcoming) and Fishburn (1981) provide helpful surveys.

²⁶Such theorems generally take the form: if an agent’s preferences obey certain constraints, then these preferences can be represented as resulting from a unique (up to linear transformations) utility function and a unique probability function by maximizing expected utility. The constraints (e.g. transitivity, connectedness, etc.) on preferences vary with different variants of representation theorems.

viable alternatives for the selection framework. Yes, the foundationalist declared that there simply is nothing more to framework choice than selecting whichever course of action sums to the greatest expected utility, but that may just be a failure of imagination on our part.²⁷ More importantly, though, there *are* alternatives to standard decision theory even if we stick to expected utility theory. Decision theories based on non-classical probability functions are a case in point. The idea can be put as follows. Standard probability theory is based on classical logic.²⁸ This generates the familiar tight connections between probability theory and classical propositional logic: e.g. a probability function assigns probability 1 to all classical tautologies and 0 to all contradictions; the probability of a proposition cannot be strictly less than the probability of a proposition that entails it, and so on. However, there are numerous well-known phenomena for which non-classical systems of logic arguably recommend themselves: constructive mathematics, quantum mechanics, apparently non-trivial inconsistent theories, semantic paradoxes, vague propositions, etc. To the extent that we should abandon classical logic in these cases, we should also abandon classical probability theory—and espouse one of its non-classical counterparts instead. In this way we arrive at a vast space of ‘non-classical’ probability theories.²⁹ But now simply plug these non-classical probability functions into standard decision theory. By doing so, we find ourselves with an equally vast array of *non-standard* decision theories. Hence if there is a reason not to treat selection frameworks in a spirit of Carnapian tolerance, it cannot be because there are no alternative frameworks to consider. And if that’s case, it seems that in suspending the principle of tolerance at the level of selection frameworks we run the risk of making ourselves guilty of precisely the kind of dogmatism that Carnap deemed ‘worse than futile’ because it ‘obstructs scientific progress’ (Carnap 1950, p. 40).

At this point, the foundationalist is likely to play the representation theorem card: ‘To be sure, there are various deviant decision theories one might concoct, and so it is true that our selection framework based on standard decision theory is not strictly speaking the only conceivable framework. Still: it provides a suitable foundation because it has a privileged standing among alternative frameworks. And the reason it enjoys such a standing is that it features standard decision theory, which is demonstrably correct in virtue of the representation theorems we can prove for it’. Now, as far as I know, it is an open question whether analogous representation

²⁷Certainly, it takes very little imagination to conceive of framework selection as being guided by choice principles other than expected utility maximization. Why should there not be cases in which, say, mini-max or some other principle would be better suited?

²⁸Or if you prefer: is defined on a Boolean algebra.

²⁹See (Williams forthcoming) for a survey.

theorems can be established for non-standard decision theories (though I am not sure anyone has investigated the matter in a serious way).³⁰ Even so: it is highly controversial that representation theorems really do afford anything like a demonstration of correctness for standard decision theory. In a recent paper, Meacham and Weisberg (2011) have mounted a persuasive case to the effect that representation theorems are unfit to play the foundational role that has often been attributed to them. Without pursuing the point further here, I take it that the worries raised by Meacham and Weisberg cast serious doubt on the idea that an appeal to representation theorems gives standard decision theory a leg up over its non-classical rivals, even if similar theorems turn out not to be provable for non-standard decision theories.

But even *if* representation theorems did have the significance the foundationalist needs them to have, this still would not vindicate our selection framework's claim to uniqueness. There *are* other ways—besides tinkering with the type of probability function built into our decision theory—of arriving at alternative selection frameworks and so undermining the uniqueness claim. Indeed, it is clear that this must be so, since a linguistic framework that is to play the role of a selection framework—indeed any linguistic framework—must have a logic built into it alongside any mathematical resources needed to do decision theory.³¹ But selection frameworks will differ if they comprise different relations of logical consequence (even if they all operate with classical probability theory). After all, the consequence relation induced by the framework's *L*-rules specifies which modes of deductive inference are acceptable within the selection framework. To take a trivial example, I need logic to get from 'The linguistic framework *L* has the highest expected utility given my utility function, my probability function and my theoretical goals' and 'If a linguistic framework has the highest expected utility, I ought to select it' to 'I ought to select *L*.' If this is right, we are faced with no less 'boundless [an] ocean of unlimited possibilities' (Carnap 1937, p. xv) when it comes to the logical consequence relations that may figure in our selection frameworks than we are with regard to first-order frameworks.³²

³⁰Though analogues of arguments for probabilistic coherence—Dutch book and accuracy dominance-style arguments—are available for many non-classical logics. See (Williams 2012).

³¹A logic is just one further component such a framework would have to include. Another example is given by the possible worlds built into our decision theoretic framework. They might be interpreted in a more or less realist or 'ersatzist' spirit. But of course the realism/anti-realism question about possible worlds is a paradigmatic example of a set of philosophical theses that Carnap would interpret as proposals for adopting the respective language forms as opposed to truth-apt claims. So this also shows that even if we were to grant that classical decision theory alone could lay claim to correctness, we would still be faced with several selection frameworks to choose from—some realist, some ersatzist.

³²One interesting (though rather un-Carnapian) approach consists in considering the intersection

There is a final attempt the foundationalist might try at this point. In order to motivate the thought that there is a privileged selection framework that blocks the regress, it might be thought that the selection framework is not arrived via an active deliberative process. Rather it is just ‘given to us’ in some way. There are different stories one might try to tell as to how it might have been given to us. For instance, we might have arrived at it through a communal process of tacit reflective equilibrium building (see e.g. Resnik (1985) and Thagard (1982)), or through the emergence of implicit conventions (Tennant 1987b). Perhaps certain norms are also the product of evolutionary adjustments of our cognitive systems to the world we live in (Maddy 2007). I will make no attempt to develop any of these possible explanations. Nor will I take a stand as to how plausible any of them are. All that matters for us here, is that any of the explanatory proposals mentioned is bound to undercut the centerpiece of Carnap’s philosophical thinking: the principle of tolerance. For if the Carnapian wished to follow the foundationalist in restricting the principle of tolerance to first-level frameworks (but taking it to be inapplicable at the meta-level of our selection framework), she would face the question why her favored explanation as to why the selection framework enjoys a special standing does not carry over straightforwardly to first-order frameworks. If a selection framework can be ‘given to us’ through impersonal, non-deliberative processes, why should the same not be true at ground-level theorizing?

6 Soft foundationalism to the rescue?

I believe these considerations give us very strong reasons for rejecting framework foundationalism. The notion that there should be a unique privileged linguistic framework that sets forth the norms governing the process of framework selection simply cannot be reconciled with some of Carnap’s most fundamental philosophical commitments. It seems safe to say that any attempt to locate a regress-stopping foundation at a higher level (not just at the second level) is bound to fall victim to analogous objections. The regress worry cannot be met by foundationalist manoeuvres. But what, then, can the Carnapian say about it? In this final section I will sketch what I take to be the best response for the Carnapian.

So far we have examined what we might call *hard* framework foundationalism. On this view, the regress bottoms out in a unique privileged framework that issues in

of all of the logical systems containing a given set of logical constants. Since the system so arrived at is arguably be the weakest (sensible) logic, it enjoys a special foundational status. Neil Tennant undertakes the important task of developing such a ‘core logic’, see (Tennant 1987a, 1997).

norms, which then set standards of correctness for all lower-level framework choices. But what if we drop the notion that the foundational framework has a privileged standing by virtue of having a claim to uniqueness (because there are no alternatives or because it has a special property that uniquely qualifies it)? What if we say instead that the framework in question may serve as a regress-blocking foundation just because we treat it as such (in a particular context)? After all, we must start somewhere! So why not, in Quine's words, start 'at home'—with the selection framework we deem appropriate for the foundational role? Call this approach *soft foundationalism*.

An analogy with the epistemic regress problem may help clarify the idea. Hard framework foundationalism stands to soft foundationalism just as foundationalism in epistemology stands to what Gilbert Harman (2003) has dubbed a 'general foundations theory'.³³ Unlike the foundationalist, who seeks to identify a privileged proper subset of our beliefs that could constitute the bedrock upon which to erect our entire belief system, the general foundations theoretician favors an innocent-until-proven-guilty policy of belief maintenance: we begin with whatever beliefs we find ourselves with at a given moment in time. We grant each of these beliefs a default justification. As we proceed, we adopt a conservative attitude to belief revision, retaining our beliefs unless and until we encounter suitably strong countervailing evidence. Unlike the special foundations theorist, who is conservative only with respect to the privileged foundational subset of her beliefs, which constitute the bottom-most elements of the justificatory chain that reaches up to the remainder of her (justified) beliefs, the general foundations theorist is conservative about *all* her beliefs. On this picture, then, our belief system as a whole does not rest on an ultimate foundation. Like Neurath's boat, it is under permanent reconstruction under the influence of incoming evidence and minimal compensatory coherence-preserving adjustments. Thus whatever its other virtues, the general foundations theory dodges the regress problem.

Soft foundationalism proposes an analogous move with respect to the regress problem generated by Carnap's voluntarism about norms of rationality. Our framework selections, much like our belief system, do not stand in need of any ultimate foundation. Rather, we can make do with a provisional (in principle revisable) foundation. And it is precisely in this light that we should view our selection framework. We should be open to future reforms of its norms. Perhaps we even occasionally help ourselves to alternative selection frameworks for particular local framework choices. In order to make such decisions between our preferred selection framework and an alternative, we may on occasion have to move up a level and appeal to a third-level

³³Harman refers to foundationalist accounts as 'special foundations theories'.

framework. All the same, we never have to make the infinitely many framework choices all at once because we can fall back on a provisional foundation.

Is a weak foundation a sturdy enough basis for Carnap's tolerance-based philosophy? What are we to make of such an approach? Here is a potential drawback. On the general foundations approach in epistemology I run the risk of finding myself with wildly confused initial beliefs (say, if I grow up in a community of voodoo-practicing, racist flat-earthers) that it might be extremely difficult for me to shed even through rational processes of belief revision. This is a well-known issue for subjective Bayesians. If my initial beliefs are sufficiently whacky, it may be that no amount of conditionalizing, will be able to reign them in, i.e. make it the case that they eventually converge on some intersubjectively agreed upon standard.³⁴ Similarly, on the soft foundations view, there is no guarantee that the norms we find ourselves with are not equally badly misguided. That is, we could imagine a hypothetical inquirer whose principles for selecting linguistic frameworks in view of her theoretical ends are so poor that a proper practice of scientific inquiry in which the inquirer's theoretical apparatus is continually improved in light of the results yielded just does not get off the ground. How worried should the Carnapian be about this hypothetical scenario?

Not overly, I think, but this requires some qualification. The question raised by this hypothetical scenario is an *explanatory* one: if we assume that there are no constraints whatsoever on the selection framework our inquirer tacitly appeals to (and hence if we assume that our agent may persistently choose frameworks that are wildly at odds with her scientific pursuits), how is it possible for a productive scientific practice to emerge? What justifiably troubles us, I take it, is that our soft foundations interpretation of Carnapian framework selection makes it look miraculous that agents proceeding in this way should ever arrive at a successful scientific practice. Surely, it cannot be merely by fluke that we hit upon a workable set of norms for selecting and devising frameworks at some point in the course of the development of scientific thought. But our interpretation provides no explanation for how we arrived at a successful scientific practice given that, as it were, we had to make up to the rules for scientific theorizing as we went along. This is a serious and difficult question, and Carnap provides no answer to it.

But Carnap is to be faulted for not providing an answer to the explanatory question? The answer, I think, is that he is not to be faulted. The explanatory question is simply not Carnap's concern (at least not his primary concern). What drives the explanatory question is the possibility that the norms for selecting and devising lin-

³⁴See Hájek and Hartmann (2010) for a compact but helpful discussion of the reach of convergence theorems for subjective Bayesianism alluded to here.

guistic frameworks a scientific community is equipped with at a hypothetical starting point of inquiry are massively off the mark, so far off the mark that even continuous adjustments will be insufficient to bring about a successful scientific practice. But Carnap does not (and need not, given his immediate aims) entertain this hypothetical for the purposes of his project of formulating a logic of science. The starting point for Carnap's project is one that takes a relatively successful scientific practice as a given. The scientific project is up and running. The task now is to devise a logic of science, whose task, as Friedman explains,

is not so much to describe the nature of science or scientific method as it has been practiced so far as to open up the possibility for a new kind of ongoing philosophical interaction with the sciences, which, in Carnap's eyes, promises to be particularly fruitful for both. Armed with the new logico-mathematical tools of modern logic (especially the new tools of metamathematics), the philosopher—that is, the logician of science—can participate, together with the scientists themselves, in the articulation, clarification, and development of formal inferential frameworks for articulating empirical theories and testing them by experimental methods (Friedman 2006, p. 50)

When viewed from this perspective of Carnap's bigger project, soft foundationalism can be seen to offer an interpretation that accords well with Carnap's broader commitments: we start with whatever we have. That is, our default selection framework is the one that articulates the rational norms and principles that are operative in guiding current scientists in choosing and devising the forms of language within which to conduct and articulate their inquiries. But Carnap's voluntarism about linguistic frameworks, ground-level and higher-level alike, guarantees that we are free to abandon and to revise our current set of rational norms for choosing frameworks whenever we see ways in which to improve them. Therefore, there is nothing special about the selection framework we work with at any given time; it has no more claim to correctness than any other linguistic framework.

Thus, once we distinguish the explanatory question from the methodological question that was Carnap's primary concern, we appreciate that soft foundationalism does provide an adequate response to the latter question. Notice, though, that is not to say that Carnap's commitment to tolerance makes it impossible for him to answer the aforementioned explanatory question. In fact, his tolerance-based project would appear to be compatible with a range of possible explanations for the emergence of a workable set of rational principles that put the scientific enterprise on its tracks.

Where does this leave us with respect to the question of the rationality of frame-

work choice? What we have seen in section 4, is that the Carnapian is committed to framework choice being subject to rational norms. Soft foundationalism now provides the final piece of the puzzle as to how the demand for the rationality can be reconciled with Carnap's other commitments—in particular, his commitments to tolerance and to non-cognitivism. The notion of rationality we arrive at is, admittedly, a thin one. Carnap refers to instrumental rationality as 'relative rationality'. What we find ourselves with now is, in fact, a notion of rationality that is *doubly* relative. It is relative in Carnap's intended sense of being relative to a value or goal, the rationality of which is not itself in question; but, as we have seen, it is also relative to a selection framework. It is in this sense that the rationality in question is indeed minimal. However, given that Carnap's concern is the methodological (not the explanatory) question, this minimal notion of rationality is all he needs.

7 Conclusion

Let us briefly take stock. My aim in this paper was to address the question as to how, on Carnap's view, it is possible to account for the rationality of the process by which one devises and adopts a linguistic framework. Since for Carnap rational norms are always tied to a linguistic framework, it would seem that there would have to exist a selection framework that articulates the norms that rationally constrain the process of framework choice. But this poses an interpretative conundrum. If such a selection framework exists, framework selection would seem to be an internal, theoretical question. But as Carnap has famously emphasized, he regards framework choice as an external, practical matter. I showed that this apparent puzzle can be resolved by distinguishing between the theoretical question of determining the means to one's theoretical ends from the practical question of determining the ends themselves. Having thus created the much-needed logical space for the concept of a selection framework, I raised the question of the nature and status of a selection framework. As for the nature of selection frameworks, I have argued that they are most plausibly conceived of in decision-theoretic terms. Moreover, I showed that the question of the status of selection frameworks is beset by a potential regress worry. I considered two possible strategies for dodging the regress: hard and soft foundationalism. I argued that hard foundationalism fails to respect undermines Carnap's commitment to the principle of tolerance. Instead, the solution to my initial question lies with the soft foundations approach according to which there is no privileged set of norms of framework selection. We simply make do with whatever set of rational principles we are currently operating with, improving our methods as we go along. Our 'central question' thus receives a thoroughly Carnapian answer.

References

- S. Awodey and A. W. Carus. From Wittgenstein’s prison to the boundless ocean: Carnap’s dream of logical syntax. In P. Wagner, editor, *Carnap’s Logical syntax of language*, pages 79–108. Palgrave Macmillan, London, 2009.
- D. Bonnay. Carnap’s criterion of logicity. In P. Wagner, editor, *Carnap’s Logical syntax of language*, pages 147–166. Palgrave Macmillan, London, 2009.
- L. Buchak. Decision theory. In A. Hájek and C. Hitchcock, editors, *Oxford handbook of probability and philosophy*. Oxford University Press, Oxford, forthcoming.
- R. Carnap. Theoretische Fragen und praktische Entscheidungen. *Natur und Geist*, 2:257–260, 1934.
- R. Carnap. Testability and meaning. *Philosophy of science*, 3 and 4:419–471 and 1–40, 1936 and 1937.
- R. Carnap. *The logical syntax of language*. Routledge, London, 1937.
- R. Carnap. *Foundations for logic and mathematics*, volume 1 of *International encyclopedia of unified science*. The University of Chicago Press, Chicago, 1939.
- R. Carnap. Empiricism, semantics and ontology. *Revue internationale de philosophie*, 4:20–40, 1950.
- R. Carnap. Value concepts. ASP/RC 089-14-01, 1958.
- R. Carnap. *Logical foundations of probability*. The University of Chicago Press, Chicago, 2 edition, 1962.
- R. Carnap. Carnap’s intellectual autobiography. In P.A. Schilpp, editor, *The philosophy of Rudolf Carnap*, pages 3–81. Open Court, Chicago, 1963a.
- R. Carnap. Abraham Kaplan on value judgments. In P.A. Schilpp, editor, *The philosophy of Rudolf Carnap*, pages 999–1013. Open Court, Chicago, 1963b.
- A. W. Carus. Carnapian rationality. Typescript of April 2014, 2014.
- R. Creath. Quine’s challenge to Carnap. In M. Friedman and R. Creath, editors, *The Cambridge companion to Carnap*, pages 316–335. Cambridge University Press, Cambridge, 2007.

- R. Creath. The gentle strength of tolerance: *The logical syntax of language* and Carnap's philosophical programme. In P. Wagner, editor, *Carnap's Logical syntax of language*, pages 203–214. Palgrave Macmillan, 2009.
- P. Fishburn. Subjective expected utility: A review of normative theories. *Theory and decision*, 13:139–199, 1981.
- M. Friedman. Analytic truth in Carnap's *Logical syntax of language*. In *Reconsidering logical positivism*, pages 165–176. Cambridge University Press, Cambridge, 1999a.
- M. Friedman. Carnap and Wittgenstein's *Tractatus*. In *Reconsidering logical positivism*, pages 177–197. Cambridge University Press, Cambridge, 1999b.
- M. Friedman. *The dynamics of reason*. CSLI Publications, Palo Alto, CA, 2001.
- M. Friedman. Carnap and Quine: Twentieth-century echoes of Kant and Hume. *Philosophical topics*, 34:35–58, 2006.
- A. Hájek and S. Hartmann. Bayesian epistemology. In J. Dancy, E. Sosa, and M. Steup, editors, *A Companion to Epistemology*, pages 93–106. Blackwell, Oxford, 2010.
- G. Harman. Reflections on *Knowledge and its limits*. *The philosophical review*, 111: 417–429, 2003.
- A. Kaplan. Logical empiricism and value judgments. In P.A. Schilpp, editor, *The philosophy of Rudolf Carnap*, pages 827–856. Open Court, Chicago, 1963.
- P. Koellner. Carnap on the foundations of mathematics. Unpublished.
- P. Maddy. *Second philosophy*. Oxford University Press, Oxford, 2007.
- C. J. G. Meacham and J. Weisberg. Representation theorems and the foundations of decision theory. *Australasian journal of philosophy*, 89:641–663, 2011.
- G. Priest. *Doubt truth to be a liar*. Oxford University Press, Oxford, 2006.
- W.V.O. Quine. Comment. In W. Salmon and G. Wolters, editors, *Logic, language and the structure of scientific theories: Proceedings of the Carnap-Reichenbach centennial*, pages 345–352. University of Pittsburgh Press, Pittsburgh, 1994.
- G. Resnik. Logic: Normative or descriptive? The ethics of belief or a branch of psychology. *Philosophy of science*, 52:221–238, 1985.

- A. Richardson. Carnapian pragmatism. In M. Friedman and R. Creath, editors, *The Cambridge companion to Carnap*, pages 295–315. Cambridge University Press, Cambridge, 2007.
- T. Ricketts. Rationality, translation and epistemology naturalized. *Journal of philosophy*, 79:117–136, 1982.
- T. Ricketts. Tolerance and logicism: Logical syntax and the philosophy of mathematics. In M. Friedman and R. Creath, editors, *The Cambridge companion to Carnap*, pages 200–225. Cambridge University Press, Cambridge, 2007.
- N. Tennant. *Anti-realism and logic*. Oxford University Press, Oxford, 1987a.
- N. Tennant. Conventional necessity and the contingency of convention. *Dialectica*, 41:79–95, 1987b.
- N. Tennant. *The taming of the true*. Oxford University Press, Oxford, 1997.
- P. Thagard. From the descriptive to the normative in psychology and logic. *Philosophy of science*, 49:24–42, 1982.
- J. R. G. Williams. Generalized probabilism: Dutch books and accuracy domination. *Journal of philosophical logic*, 41:811–840, 2012.
- J. R. G. Williams. Non-classical logic and probability. In A. Hájek and C. Hitchcock, editors, *Oxford handbook of probability and philosophy*. Oxford University Press, Oxford, forthcoming.