36

Experimental Philosophical Logic

DAVID RIPLEY

The purpose of this chapter is to explore the intersection of experimental philosophy and philosophical logic, an intersection I'll call *experimental philosophical logic*. In particular, I'll be looking for and sketching some ways in which experimental results, and empirical results more broadly, can inform and have informed debates within philosophical logic. Here's the plan: first, I'll lay out a way of looking at the situation that makes plain at least one way in which we should expect experimental and logical concerns to overlap. Then, I'll turn to the phenomenon of *vagueness*, where we can see this overlap explored and developed from multiple angles, showing just how intimately related experiment and logic can be. Finally, I'll canvass some other cases where we have similar reasons to expect productive interactions between experimental methods and formal logic, and point to some examples of productive work in those areas.

36.1 Logic, Pure and Applied

Let's open by briefly considering a distinction between *pure* and *applied* logic. This distinction is analogous to the one between pure and applied algebra, or between pure and applied topology, or between pure and applied versions of any branch of mathematics. (I don't pretend that any of these distinctions is precise, or that there are no problem cases; the gist is all that matters.)

Roughly, pure logic is an exploration of the properties and relations occupied by logical systems *in themselves*, without attending to any particular use they may or may not have. Typical questions within pure logic: is proof system *X* sound and complete for model theory *Y*?; is such-and-such a logical system decidable? compact? finitely axiomatizable?; is this rule admissible in that system?; and so on. Pure logic is most naturally thought of as a subfield of *mathematics*, although it is of course also pursued by researchers in philosophy, linguistics, computer science, electrical engineering, and so on, in pursuit of our own varied research interests.

A Companion to Experimental Philosophy, First Edition. Edited by Justin Sytsma and Wesley Buckwalter.

On the other hand, and even more roughly, applied logic is the use of logic in cases where the main questions at hand are not in the first instance logical; this is then an application of logic to whatever domain those questions *are* drawn from. (This is particularly rough since logic can of course be applied in pursuit of *logical* questions as well; but I set this complication aside here. The goal is not to have a precise characterization of applied logic, but just to get the rough idea on the table.) Examples of questions sometimes or often answered in part by applying logic: what output will this circuit produce given that input?; which conditional sentences are true?; does anything not exist?; and so on. Questions fruitfully pursued by applying logic are often not in the first instance questions about logic at all; they can occur in pursuit of just about any imaginable research agenda. As it happens, most of the questions that arise in what's called *philosophical logic* are of this sort: they are questions in applied logic.

36.1.1 Setting Some Things Aside

I want to use this rough and ready distinction to whittle down the topic of this chapter a bit, to say a bit about what I won't say anything more about. In the first place, there are a number of questions involving logic that don't fall at all neatly into the divide between pure and applied. For example, there are mathematical questions about the relations between logic and algebra, topology, recursion theory, category theory, and so on; I will ignore these.

There also remain philosophical questions about logic that don't fit into the divide, to do with logic's epistemological, metaphysical, or normative statuses; with whether logical claims are true by stipulation when true; with what makes something logic rather than not logic; and so on. These questions all have relatives that are about, say, arithmetic instead of logic. But those questions *about* arithmetic are not themselves *part of* arithmetic; and these questions *about* logic are not themselves *part of* logic – pure, applied, or otherwise. I will ignore all of these questions here, not because I think they are unimportant, or because I think that empirical data must have nothing to say to them, but because I am focusing on ways to pursue *logic* empirically. Ways to engage in these *other* pursuits (epistemology, metaphysics, etc.) empirically are not my topic. (I don't pretend that *this* line is bright or clear either.)

I will also set aside *pure* logic. Again, this is not to say that something like experiment cannot help with pure logic. For example, computers often allow us to explore questions in pure logic in a way that at least can *feel* experimental. In this connection (Meyer, 1983, 450) offers enthusiastic thanks "... to the Beast itself, which bids fair to at last make of Logic an *empirical* science – putting the *quietus* to the Kantian *a priorism* which has lingered too long in our subject, even after it had been vanquished almost everywhere else...", and he is onto something. (From context, one guesses that the Beast is a particular computer wrangled by John Slaney.) But this kind of thing is not what I will focus on here; these methods, while valuable, are sufficiently unlike other methods in experimental philosophy that I suspect they call for a separate treatment entirely.

36.1.2 Experimental Applied Logic

Where we are most likely to find productive overlap between experimental philosophy and logic is in *applied* logic. In particular, I reckon there are two features to keep an eye open for; when a particular question in applied logic exhibits both of these features, it seems likely enough that empirical methods will be helpful. (There may be other cases, but all the cases I will examine fit this mold.) Feature one: some part of what's at issue involves rigging up a logical system as a model of a particular target phenomenon. (Here, I use "model" in roughly the sense of Shapiro 2001 and Cook 2002, and not in the sense of "model theory" (for which see Hodges 1993).) Feature two: the target phenomenon being modeled is itself amenable to experimental exploration.

When a question exhibits both of these features, it should be clear how experimental methods can help. We need to know something about what a target phenomenon itself is like in order to figure out how to rig up a logical system that captures some of its interesting features. If we can use experimental methods to find out what the target phenomenon is actually like, then these methods can become intimately involved in constructing our logical models.

As it happens, many of the standard questions that arise in pursuit of philosophical logic exhibit exactly these two features. We should expect them to be ripe for empirical exploration. To draw out how these features work in practice, let's turn to an example: the case of vagueness.

36.2 Vagueness

From a certain perspective, vagueness can make it seem completely puzzling that we are able to successfully communicate at all. After all, almost everything we say is shot through with vague words, and there are notorious problems understanding how vague words can draw distinctions at all, given that they do not draw sharp distinctions. (For helpful overviews, see Williamson 1994; Keefe 2000; and Smith 2008.) From this same perspective, it can seem completely puzzling that we are able to successfully think and understand at all; after all, our thought too is shot through with vagueness.

Nowadays, these things do not seem puzzling for long. There is no shortage of fleshed-out theories of vagueness waiting to resolve our puzzlement, to reassure us either that distinctions simpliciter do not require sharp distinctions, or that sharp distinctions (in the senses needed) are in fact present in our talk and thought. Whole forests have met the ax under the last, say, 40 years' onslaught of treatments of these phenomena. Most of these treatments include a substantial logical component: part of the goal is often to construct a logical system that shows us something about how vague talk and thought behave.

If you want to know how it is *possible* that our talk and thought can succeed in drawing distinctions despite their vagueness, you can take your pick of these theories. Myself, I'm reasonably optimistic about most of them for this purpose; at this point there's been enough argy-bargy that we've got a nice bunch of reasonably coherent and well-developed views on the table. (Any attempt to catalog or categorize these is going to mangle some and omit others, and would be beside the point here anyhow.) Each gives us a picture of how it is *possible* to avoid being dragged along the sorites argument into disaster.

But proponents of these views at least sometimes seem to understand the views as *rivals*. As answers to the "how possible" question, they are not; such questions admit multiple answers easily. It would be possible like *this*, or like *that*, or like *the other*. (To be sure, we could disagree about whether, say, *the other* is really possible, or whether *that* is really a way to do what we're trying to do. But there is no tension in proposing multiple *possible* approaches.)

There are at least two questions in the area, though, that are more likely to generate conflict between distinct hypotheses: first, how *do* we manage to draw distinctions using vague talk and thought? and second, how *should* we do it? Following Scharp (2013) (on a different topic), call these the *descriptive* and *prescriptive* questions, respectively. To focus yet more finely, I'll set the prescriptive question aside; it is the descriptive question that is most clearly open to experimental exploration.

36.2.1 *The Descriptive Question*

It is the descriptive question that is most clearly amenable to empirical treatment. To find out how we in fact do get a job done, armchair reflection can only go so far. For example, consider *epistemicist* theories of vagueness as possible answers to the descriptive question.

According to epistemicist theories (e.g., Sorensen 2001; Williamson 1994), vagueness is a purely epistemic phenomenon. On these theories, the only sense in which vague predicates lack sharp boundaries is that they lack *known* sharp boundaries, and similarly for vague concepts. That is, the sharp boundaries are in fact present; we simply do not know where they are.

If this is right, we might expect it to affect the way we deal with vague predicates. For example, Bonini *et al.* (1999, 387) adopt the following hypothesis: "a typical speaker *S* of a natural language... mentally represents vague predicates in the same way as other predicates with sharp true/false boundaries of whose location *S* is uncertain." It should be clear that this hypothesis is both intimately related to epistemic theories of vagueness and open to experimental study.

(Of course, epistemicism per se does not *require* this hypothesis. You could think, for example, that we are not only ignorant of where the sharp boundary is, but also ignorant of the fact that there is one; this latter ignorance might lead us to mentally represent vague predicates in a different way from predicates with boundaries which are unknown but known to be precise. But this latter hypothesis, like the simpler one I consider more fully, is both amenable to experimental exploration and logically tractable. It too would provide a fruitful, if more complex, site for experimental philosophical logic.)

Indeed, Bonini *et al.* (1999) reports a series of studies intended to gather evidence for this hypothesis. For example, in one (their Study 4), participants were asked to fill in the blank in the following:

> When is a man tall? Of course, very big men are tall and very small men are not tall. We're interested in your view of the matter. Please indicate the smallest height that in your opinion makes a man tall.
> A man is tall if his height is greater than or equal to ________ centimetres.

Different participants were asked to fill in the blank in a corresponding question involving "not tall" instead of "tall". (Other predicates were used as well.) Judgments of the lower threshold it takes to be tall were significantly higher ($p < 0.001$) than judgments of the higher threshold it takes to be not tall. At first, this might suggest a "gap" between men that are tall and men that are not tall, at least in participants' minds (which, recall, is what's in question here!) – but the paper also reports similar discrepancies in responses to questions like "A man is at least of average height among 30-year-old Italians if his height is greater than or equal to ________ centimetres" and "A man is not as tall as average among 30-year-old Italians if his height is less than or equal to ________ centimetres". In these latter questions, no vagueness at all is present; instead, participants can be assumed to be ignorant of the relevant (fully precise) facts. So the hypothesis that speakers represent vague predicates just as they do precise predicates with unknown boundaries is compatible with the results reported in Bonini *et al.* (1999).

36.2.1.1 Borderline Contradictions

It does not, however, seem to sit well with other results. For one thing, participants in experiments reported in works by Alxatib and Pelletier (2011), Ripley (2011a), Serchuk, Hargreaves, and Zach (2011), and Égré, de Gardelle, and Ripley (2013) all agree, to varying degrees, with contradictory-looking sentences (like "The man is both tall and not tall") in borderline cases of vague predicates. (Call these "borderline contradictions", just to have a name for them; it is contentious whether they are actually contradictory.) Yet it seems unlikely (to say the least) that speakers will agree to such contradictory-looking sentences whenever they are ignorant. (Of course, this itself is open for experimental exploration; as far as I know it has received none.) These results cast doubt on Bonini *et al.*'s (1999) main hypothesis.

Participants' agreement to these borderline contradictions is itself an interesting topic, and a number of hypotheses as to the source of such agreement are plausible. (A number of natural hypotheses also fail to be plausible when fully thought through. See Ripley 2011a for discussion of a number of potential hypotheses, as well as works by Alxatib and Pelletier (2011), Cobreros *et al.* (2012), and Alxatib, Pagin, and Sauerland (2013) for hypotheses of particular interest, drawing on three- and four-valued logics in the first two cases, and on fuzzy logics in the last case.)

Borderline contradictions also give an example where logic can inform experiment, rather than the other way around: Ripley (2011b) argues that, when it comes to explaining participants' agreement with borderline contradictions, two seemingly incompatible theses – contextualism and dialetheism – in fact have identical experimental predictions, absent concrete operationalizations of "context." (On a contextualist explanation of the phenomenon, participants are not really agreeing with contradictory claims; there is some shift in context that dissolves the apparent contradiction, so the thing participants are agreeing with is not contradictory. On a dialetheist explanation, participants are agreeing to contradictory claims; no shift in context is required.)

The core of this argument is purely logical: for every contextualist model there is a prediction-equivalent dialetheist model, and vice versa; this is established in the usual (nonexperimental) ways. But it is clearly of import for setting experimental hypotheses about borderline contradictions: such hypotheses must either include some operationalization of "context," or fail to distinguish between contextualism and dialetheism. Similarly, it is of import for devising experiments to try to decide between these hypotheses, as it makes plain that the interpretation of such experiments must depend on particular choices about what a context is (and in particular, when contexts change.) Since contextualism and dialetheism are (or at least still appear to be) distinct hypotheses about participants' approach to borderline contradictions, this logical result tells us something about how we can (and can't) improve our evidential status via experiment.

Any hypothesis about these borderline contradictions must also accommodate the phenomenon first reported by Alxatib and Pelletier (2011), and later also by Égré, de Gardelle, and Ripley (2013), that a number of participants agree to borderline contradictions *while disagreeing* with their conjuncts. Such participants might, for example, agree with "The man is both tall and not tall," while disagreeing with both "The man is tall" and "The man is not tall." For potential explanations of this phenomenon (and related), see works by Alxatib and Pelletier (2011), Sauerland (2011), Cobreros *et al.* (2012), Alxatib, Pagin, and Sauerland (2013), and Cobreros *et al.* (2015).

36.2.1.2 Instability

There is different evidence that pushes against Bonini *et al.*'s (1999) hypothesis in work by Hampton *et al.* (2012). It has long been known (see, e.g., McCloskey and Glucksberg 1978) that participants' forced judgments about whether a vague predicate applies to a borderline case are *unstable*: the very same participant may well judge one week that it does apply and the next week that it does not. Unsurprisingly, as Hampton *et al.* (2012) offer evidence for, forced judgments of unknown matters of fact are also unstable.

But there is a way to reduce the instability of unknown factual judgments: allow participants to avoid the forced choice and say that they don't know. The availability of this option, in Hampton *et al.*'s (2012) experiments, significantly reduced the instability of these factual judgments. This makes sense: participants who had this third choice were no longer being forced to guess wildly. If borderline cases of vague predicates are mentally represented just like unknown cases of precise predicates, then the same idea should work for vague predicates: allowing participants to avoid the forced choice by saying that they don't know whether the vague predicate applies should increase the stability of their judgments. But in Hampton *et al.*'s, (2012) experiments, this does not happen; judgments of borderline cases of vague predicates are no more stable when "don't know" is an option than they are without it.

There is evidence, then, that speakers' interpretations of vague predicates differ in important ways from their interpretations of precise ones, even precise ones with unknown boundaries. If we are interested in using logical tools to frame hypotheses about such interpretations, we should pay attention to the experimental evidence available, and gather more where it would help. (For example, it would be nice to have more data than we currently have about speakers' responses to so-called "penumbral connections" (for which see, e.g., Fine 1975).) We still do not know, of course, exactly how speakers manage to use vague predicates to draw distinctions, but it seems clear at this point that interaction between logical and empirical methods can yield more than either would on their own.

36.3 Other Cases

In this section, I'll canvass some other possible and actual uses of experimental methods to approach questions in philosophical logic. In each case, the pattern is the same: the question is about which logical system or system best captures certain aspects of a particular phenomenon, and the phenomenon in question is itself open to experimental investigation.

36.3.1 Conditionals

Conditionals provide another example of a paradigm area of research in philosophical logic that is wide open for empirical exploration. As in the case of vagueness, we have a wide array of logical systems on offer exhibiting very different behavior from each other. My discussion here is necessarily very brief. For a good next stop, see work by Douven (this volume).

Consider some of the different behavior exhibited by these different systems. (I'll write $\rightarrow$ for a conditional, $\wedge$ for conjunction, $\vee$ for disjunction, $\neg$ for negation, and A, B, C for arbitrary clauses.) In some systems, $(A \vee B) \rightarrow C$ entails $(A \rightarrow C) \wedge (B \rightarrow C)$ (Lycan 2001; Anderson and Belnap 1975; Grice 1975); in some it does not (Lewis 1973; Stalnaker 1968). In some systems, $(A \rightarrow B) \vee (A \rightarrow \neg B)$ is always true (Stalnaker 1968; Grice 1975); in others, it is not (Lewis 1973; Lycan 2001; Anderson and Belnap 1975)). In some systems, $(A \wedge B) \rightarrow C$ entails $(A \rightarrow C) \vee (B \rightarrow C)$ (Grice 1975); in others, it does not (just about all others). And so on.

These systems themselves are all citizens in good standing of logic-land; there need be nothing wrong with any of them in the abstract. But they are often intended as hypotheses about at least some natural-language conditional constructions; even when they are not, we can interpret them as though they were, to see if any light is shed. Here the differences in logical behavior give us a way to evaluate these theories. Do or don't speakers take "If A or B, then C" to entail "If A then C, and if B then C"? Do or don't speakers take "Either if A then B, or if A then not B" to always be true? And so on. By exploring these questions, we can make progress on developing the formal semantics of conditionals, a topic that remains up to its ears in both logic and empirical import.

Interestingly, there is research that seems to indicate that *all* of the aforementioned theories of conditionals come strikingly apart from reported intuitions, at least about counterfactual conditionals. All of these theories involve systems in which $A \rightarrow (B \wedge C)$ entails $A \rightarrow B$, in which $A \rightarrow B$ entails $A \rightarrow (B \vee C)$, and in which $A \rightarrow (B \wedge C)$ entails $(A \wedge B) \rightarrow C$. But Miyamoto, Lundell, and Tu (1989) report an experiment involving counterfactual conditionals in which participants judge otherwise, being more likely to agree with the premise than with the conclusion in each case.

Of course, such an experiment does not automatically show that all these theories are in error, even if its results turn out to be robust. There is much more feeding into participants' reported judgments than just their competence with conditionals, and it could be that other factors are in part responsible for the observed results. A hypothesis along these lines would have to stand or fall on its own merits, of course, but that is the usual situation. (One factor likely to be involved is the way in which participants reason under supposition; see, e.g., Byrne and Tasso 1999; Thompson and Byrne 2002.)

36.3.2 Wason Selection Task

Logic is not only useful for exploring hypotheses about truth conditions and entailments but also for exploring hypotheses about *reasoning*. For example, consider the famous Wason selection task (Wason 1968, 1969). In this task, participants are shown one side of each of four cards and told that each has a letter on one side and a number on the other; they can see one card with a consonant showing, one with a vowel, one with an odd number, and one with an even number. They are then asked which cards they would need to turn over to decide whether the rule "If there is a vowel on one side, then there is an even number on the other side" is true of the cards.

Although theories of conditionals disagree about many things, they tend to have at least this in common: taking a conditional with a true antecedent and false consequent to be false. On any of these theories, then, unless participants have some other way of concluding that the conditional rule they are asked to evaluate is false, we might expect them to turn at least the card with the vowel showing *and* the card with the odd number showing: the one known to make the antecedent true (in case it makes the consequent false, and so settles the question) and the one known to make the consequent false (in case it makes the antecedent true, and so settles the question). But while participants do tend to turn the card with the vowel showing (more turn this card than any other), they tend *not* to turn the card with the odd number. It is not even the second-most-chosen card; that honor goes to the card with the *even* number showing.

Why do participants make these choices in this task? There are a number of potential explanations, and one of the key dimensions they differ along is the extent to which they see logic as potentially capturing the key reasoning resources participants draw on in reaching their answers.

For example, it is often pointed out that participants respond very differently if instead of being asked about vowels and even numbers, they are asked to check the rule "If someone is drinking alcohol, they are over 18," when faced with someone of unknown age drinking alcohol, someone of unknown age not drinking alcohol, someone over 18 drinking an unknown liquid, and someone under 18 drinking an unknown liquid. (Here participants tend to check the age of the person known to be drinking alcohol, and the liquid of the person known to be under 18; unlike predominant responses to the original selection task, this choice reflects what most theories of conditionals would recommend.)

There are clear structural similarities between the age-and-alcohol version of the task and the letter-and-number version of the task, and yet they produce strikingly different responses. Because of this, some authors (e.g., Cosmides and Tooby 1992) have concluded that logical tools will not shed much light on participants' responses in these tasks, since logical tools are sensitive only to structure and not content, while participants seem to be responding more to content than to structure. On the other hand, some authors (e.g. Stenning and van Lambalgen 2008) argue that there are different logical forms in play in these two versions of the task leading to the difference in responses, and that the difference in content is mainly important for the cues it gives participants about which logical form is most appropriate.

This is a slightly different case from the cases of vagueness and conditionals. Here it is not even agreed that setting up a logical treatment of the phenomenon in question is an appropriate way to proceed. Rather, we have some authors arguing that it is not and others arguing that it is. Even deciding *this* question, though, will of course require interplay between logical and empirical tools: if it turns out that a logical approach does not in the end give us a good way to understand participants' responses to the task, this will only emerge by considering what a logical approach could in fact offer.

36.3.3 Others

These examples hardly exhaust the situation; there are many other areas in which logic and experiment can interact to the benefit of both. I take this opportunity to list a few further examples.

Another task in which participants seem to reason in surprising (at least to some) ways is the so-called "suppression task," reported and discussed by, for example, Byrne (1989) and Dieussaert *et al.* (2000). Logical approaches to this task are explored in works by, for example, Stenning and van Lambalgen (2008) and Dietz *et al.* (2014).

Experimental methods that can help us explore hypotheses involving "truth-value gaps" are investigated by Križ and Chemla (2014). (One method they investigate involves separately asking for judgments of truth and judgments of falsity, with each judgment binary; while another involves a single trinary judgment.)

The phenomenon of "existential import" – whether "All *A*s are *B*s" does or does not entail "There is an *A*" (and so "At least one *A* is a *B*") – is explored experimentally by Begg and Harris (1982), Newstead and Griggs (1983), and Rips (1994). As this is one of the most striking differences between ancient and modern logical approaches to quantification, it's important to have a clear picture of the empirical situation: have logicians, in turning away from existential import, come closer to or gone farther from the way speakers use quantifiers? For discussion, see, for example, Geurts 2007.

There are also "knight-knave" tasks (Rips 1989), involving reasoning about truth and falsity directly. Particularly interesting in this connection is the finding by Elqayam (2006) that participants treat the liar paradox differently from the truth-teller paradox. This contrasts strikingly with much of the philosophical literature on these paradoxes.

36.4 Worries

In this closing section, I want to consider some possible worries about whether experimental methods really can help in the practice of philosophical logic. While I am hesitant to attribute exactly these worries to anyone in particular, I think they are representative of the kinds of concerns it is natural to have about the kinds of projects I have been discussing, and I have heard at least a relative of each of these raised in conversations about these topics.

My responses to these worries will for the most part be sympathetic; I think most of them really are on to something. We should expect some difficulties to arise in the pursuit of these projects, as with any other, and it is good to be aware from the outset of some predictable snags. But if these all provide reasons for caution, none provides a reason to despair; the appropriate response, as ever, is to do what we can to take account of these risks and move forwards.

Worry: People, mostly, are bad at logic. If our concern is to decide on a logic for a particular domain (vagueness, conditionals, etc.), we should not expect to get good advice from logically untrained intuitions.

Reassurance: This is the local flavour of the familiar *expertise* objection that arises in all areas of experimental philosophy. For discussions of the general situation, see Knobe and Nichols 2007; Alexander this volume.

By looking at the examples I have given, we can see how badly misplaced this objection is *here*, whatever its status elsewhere. To my knowledge, nobody has proposed to explore logical hypotheses about vagueness, or about conditionals, or about reasoning, by asking speakers about the hypotheses themselves. Imagine the survey that would be! "Please report your intuition: Is it the case that a conditional is true at a possible world w iff its consequent is true in the most similar possible world w' to w such that its antecedent is true at w'?"

Rather, participants are typically asked to judge the truth values of object-level sentences, or to engage in reasoning, without any funny business intervening, without time pressure, and so on. Untutored speakers cannot be assumed to be experts on the proper logical theory of vagueness; they *can* be assumed to be experts on whether they take anybody to be neither tall nor not tall (and if anybody, who?). It is precisely questions like the latter that are in play here. The job of a logical hypothesis is to play a role in predicting, explaining, and systematizing these judgments, not to itself serve as the topic participants are asked to reflect upon.

Worry: Logical approaches produce tidy, clean, and elegant results. These will always come off badly when held up against experimental data, which is messy and noisy.

Reassurance: This worry is on to something important. A direct comparison between a logic-based hypothesis and the results of a particular experiment will indeed almost always look like "no match." Theorems and statistical tendencies are different kinds of patterns altogether. Even if a particular logical system correctly describes (certain aspects of) a speaker's competence with vague predicates, or with conditionals, that speaker still might not give the responses in an experimental environment that a naive reading of the logical system might seem to predict. There are always other factors involved in generating participants' responses.

This creates some difficulty in interpreting experimental data as evidence for or against a particular hypothesis about speaker competence. But while we shouldn't neglect this difficulty, we shouldn't overstate it either.

A natural way to proceed, given a particular range of logical hypotheses and particular experimental results, is to see what auxiliary hypotheses would need to be conjoined to each logical hypothesis to give the most plausible account of the observed results. These auxiliary hypotheses can then be evaluated in their own right, and their relative plausibility will tell us something about the relative plausibility of the logical hypotheses that require them. This is the method adopted in (Ripley 2011a), for example. There is nothing distinctively logical about this method; it is what we in general must do to relate theory to experiment.

Worry: We already know what the right logic is, by a priori means. (Often classical logic is intended here, although this is inessential.) No amount of information about what sentences people agree to, or what entailments they think hold, can change that; either they agree with the correct logic, or they are simply wrong. And while their rightness or wrongness might be an interesting topic of discussion in its own right, it is not what we care about when we are doing logic.

Reassurance: There might be some sense in which we can determine what the right logic is by a priori means. But surely we can't determine what the right logic is *for describing a particular natural phenomenon* by a priori means; that depends on what the phenomenon itself is like. Similarly, we cannot even determine *whether* logical systems provide the right kind of tools for describing a particular natural phenomenon by a priori means; we must first know something about the phenomenon in question before we decide whether logic can help with it.

Someone sympathetic to this worry perhaps sees the use of logic in formal semantics, or formal pragmatics, or the study of actual reasoning as "off-label" uses, as possibly helpful for understanding the phenomena in question, but as straying from the real point of logic (or, perhaps, of Logic). Indeed, in so straying, one might worry, we are in danger of resurrecting an old discredited psychologism about logic: thinking that logic is, in the end, a description of certain psychological features.

But none of the research I've pointed to requires taking any view at all about what logic *is*. That kind of question is among the questions that can simply be set aside while we worry about how best to apply logical tools. Logical and experiment can fruitfully interact in a variety of applications, and that is all we need for experimental philosophical logic to be worthwhile. There is no sense in which these uses need to be privileged above other (perhaps a priori, perhaps prescriptive) uses of logic for them to be fruitful. Off-label uses are still uses. Moreover, it should be clear that these uses of logic, whether primary or not, are very much at the heart of philosophical logic: the questions of vagueness, conditionals, and the like that occupy philosophical logicians have sizeable empirical components.

36.5 Conclusion

We should expect experiment and logic to fruitfully interact whenever a field of inquiry involves rigging up a logical system to capture some experimentally-explorable phenomenon; in these cases, logical approaches will help us decide which aspects of the phenomenon to experimentally explore, and experimental approaches will help us choose which logics best capture the phenomenon.

These conditions are present in a number of areas of philosophical logic; although I have focused here on a few examples, these hardly exhaust the potential applications.

References

Alexander, J. This volume. Philosophical Expertise.

Alxatib, S., P. Pagin, and U. Sauerland. 2013. Acceptable Contradictions: Pragmatics or Semantics? A Reply to Cobreros et al. *Journal of Philosophical Logic* 42(4): 619–634.

Alxatib, S., and J. Pelletier. 2011. The Psychology of Vagueness: Borderline Cases and Contradictions. *Mind and Language* 26(3): 287–326.

Anderson, A.R., and N.D. Belnap. 1975. *Entailment: The Logic of Relevance and Necessity*, Vol. 1. Princeton, NJ: Princeton University Press.

Begg, I., and G. Harris. 1982. On the Interpretation of Syllogisms. *Journal of Verbal Learning and Verbal Behavior* 21: 595–620.

Bonini, N., D. Osherson, R. Viale, and T. Williamson. 1999. On the Psychology of Vague Predicates. *Mind and Language* 14: 377–393.

Byrne, R.M. 1989. Suppressing Valid Inferences with Conditionals. *Cognition* 31(1):61–83.

Byrne, R.M., and A. Tasso. 1999. Deductive Reasoning with Factual, Possible, and Counterfactual Conditionals. *Memory and Cognition* 27(4): 726–740.

Cobreros, P., P. Égré, D. Ripley, and R. van Rooij. 2012. Tolerant, Classical, Strict. *Journal of Philosophical Logic*, 41(2): 347–385.

Cobreros, P., P. Égré, D. Ripley, and R. van Rooij. 2015. Nonmonotonicity and Partiality: Pragmatic Interpretations of Vague Expressions. *Journal of Philosophical Logic* 44(4): 375–393.

Cook, R.T. 2002. Vagueness and Mathematical Precision. *Mind* 111(442): 225–248.

Cosmides, L., and J. Tooby. 1992. Cognitive Adaptations for Social Exchange. In *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*, edited by J.H. Barkow, L. Cosmides, and J. Tooby, 163–228. Oxford: Oxford University Press.

Dietz, E.-A., S. Hölldobler, and C. Wernhard. 2014. Modeling the Suppression Task under Weak Completion and Well-Founded Semantics. *Journal of Applied Non-Classical Logics*, 24(1–2): 61–85.

Dieussaert, K., W. Schaeken, W. Schroyen, and G. d'Ydewalle. 2000. Strategies During Complex Conditional Inferences. *Thinking and Reasoning* 6(2):125–161.

Douven, I. This volume. Experimental Approaches to the Study of Conditionals.

Égré, P., V. de Gardelle, and D. Ripley, 2013. agueness and Order Effects: Evidence for Enhanced Contrast in a Task of Color Categorization. *Journal of Logic, Language, and Information* 22(4): 391–420.

Elqayam, S. 2006. The Collapse Illusion Effect: A Semantic-Pragmatic Illusion of Truth and Paradox. *Thinking and Reasoning* 12(2): 144–180.

Fine, K. 1975. Vagueness, Truth, and Logic. *Synthese* 30: 265–300.

Geurts, B. 2007. Existential Import. In *Existence: Syntax and Semantics*, edited by I. Comorovski, and K. von Heusinger, 253–271. Dordrecht: Springer.

Grice, P. 1975. Logic and Conversation. In *The Logic of Grammar*, edited by D. Davidson and G. Harman, 64–75. Encino, CA: Dickenson.

Hampton, J.A., B. Aina, J.M. Andersson, H.Z. Mirza, and S. Parmar. 2012. The Rumsfeld Effect: The Unknown Unknown. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 38(2): 340–355.

Hodges, W. 1993. *Model Theory*. Cambridge: Cambridge University Press.

Keefe, R. 2000. *Theories of Vagueness*. Cambridge: Cambridge University Press.

Knobe, J., and S. Nichols. 2007. An Experimental Philosophy Manifesto. In *Experimental Philosophy*, edited by J. Knobe and S. Nichols, 3–14. Oxford: Oxford University Press.

Križ, M., and E. Chemla. 2014. Two Methods to Find Truth Value Gaps and Their Application to the Projection Problem of Homogeneity. MS University of Vienna and LSCP.

Lewis, D. 1973. *Counterfactuals*. Oxford: Blackwell.

Lycan, W.G. 2001. *Real Conditionals*. Oxford: Oxford University Press.

McCloskey, M.E., and S. Glucksberg. 1978. Natural Categories: Well-Defined or Fuzzy Sets? *Memory and Cognition* 6(4): 462–472.

Meyer, R.K. 1983. A Note on R-matrices. *Notre Dame Journal of Formal Logic* 24(4): 450–472.

Miyamoto, J.M., J.W. Lundell, and S. Tu. 1989. Anomalous Conditional Judgments and Ramsey's Thought Experiment. In *Proceedings of the Eleventh Annual Conference of the Cognitive Science Society*, 212–220. Hillsdale, NJ: Erlbaum.

Newstead, S.E., and R. Griggs. 1983. Drawing Inferences from Quantified Statements: A Study of the Square of Opposition. *Journal of Verbal Learning and Verbal Behavior* 22: 535–546.

Ripley, D. 2011a. Contradictions at the Borders. In *Vagueness and Communication*, edited by R. Nouwen, R. van Rooij, H.-C. Schmitz, and U. Sauerland, 169–188. Berlin: Springer.

Ripley, D. 2011b. Inconstancy and Inconsistency. In *Understanding Vagueness: Logical, Philosophical, and Linguistic Perspectives*, edited by P. Cintula, C. Fermüller, L. Godo, and P. Hájek, 41–58. London: College Publications.

Rips, L.J. 1989. The Psychology of Knights and Knaves. *Cognition* 31(2):85–116.

Rips, L.J. 1994. *The Psychology of Proof: Deductive Reasoning in Human Thinking*. Cambridge, MA: MIT Press.

Sauerland, U. 2011. Vagueness in Language: The Case Against Fuzzy Logic Revisited. In *Understanding Vagueness: Logical, Philosophical, and Linguistic Perspectives*, edited by P. Cintula, C. Fermüller, L. Godo, and P. Hájek, 185–198. London: College Publications.

Scharp, K. 2013. *Replacing Truth*. Oxford University Press, Oxford.

Serchuk, P., I. Hargreaves, and R. Zach. 2011. Vagueness, Logic, and Use: Four Experimental Studies on Vagueness. *Mind and Language* 26(5): 540–573.

Shapiro, S. 2001. Modeling and Normativity: How Much Revisionism Can We Tolerate? *Agora* 20(1): 159–173.

Smith, N.J.J. 2008. *Vagueness and Degrees of Truth*. Oxford: Oxford University Press.

Sorensen, R. 2001. *Vagueness and Contradiction*. Oxford: Oxford University Press.

Stalnaker, R.C. 1968. A Theory of Conditionals. In *Studies in Logical Theory*, edited by N. Rescher, 98–112. Blackwell: Oxford.

Stenning, K., and M. van Lambalgen. 2008. *Human Reasoning and Cognitive Science*. Cambridge, MA: MIT Press.

Thompson, V.A., and R.M. Byrne. 2002. Reasoning Counterfactually: Making Inferences about Things That Didn't Happen. *Journal of Experimental Psychology: Learning, Memory and Cognition* 28(6): 1154–1170.

Wason, P.C. 1968. Reasoning about a Rule. *Quarterly Journal of Experimental Psychology* 20: 273–281.

Wason, P.C. 1969. Regression in Reasoning? *British Journal of Psychology* 60: 471–480.

Williamson, T. 1994. *Vagueness*. London: Routledge.