



HAL
open science

The experimental philosophy of logic and formal epistemology

Igor Douven, Shira Elqayam, Karolina Krzyzanowska

► **To cite this version:**

Igor Douven, Shira Elqayam, Karolina Krzyzanowska. The experimental philosophy of logic and formal epistemology. Alexander Max Bauer & Stephan Kornmesser. The Compact Compendium of Experimental Philosophy, Walter de Gruyter, In press. hal-03921661

HAL Id: hal-03921661

<https://hal-cnrs.archives-ouvertes.fr/hal-03921661>

Submitted on 4 Jan 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

The Experimental Philosophy of Logic and Formal Epistemology

Igor Douven

IHPST / CNRS / Panthéon–Sorbonne University

Shira Elqayam

School of Applied Social Sciences, De Montfort University

Karolina Krzyżanowska

ILLC, University of Amsterdam

Abstract

Classical logic was long believed to provide the norms of reasoning. But more recently researchers interested in the norms of reasoning have shifted their attention toward probability theory and various concepts and rules that can be defined in probabilistic terms. In philosophy, this shift gave rise to formal epistemology, while in psychology, it led to the New Paradigm psychology of reasoning. Whereas there has traditionally been a clear division of labor between philosophers and psychologists working on reasoning, the past decade has seen an increasing collaboration between philosophers and psychologists, from which an experimental philosophy of logic and formal epistemology emerged. An area in which the fruits of this collaboration have been particularly in evidence is the research concerned with conditionals and conditional reasoning. This chapter showcases contributions to this area to underline the value of the said branch of experimental philosophy more generally.

Keywords: conditionals; experimental philosophy; formal epistemology; logic; New Paradigm psychology of reasoning; pragmatics; probability; semantics.

1 Introduction

When we argue, we typically present some premises as warranting a given conclusion: so-and-so is true, thus/hence/therefore/ . . . such-and-such is true as well. Although people start to use words like “thus,” “therefore,” and their ilk early in life, and they use them frequently, the use of these words can easily spark controversies, disputes of the following form being nothing out of the ordinary:

Speaker 1: A, B, and C, hence D.

Speaker 2: No! I grant you A, B, and C. But D doesn't follow!

How are we to arbitrate such disputes? Can anything systematic be said about when “thus,” “hence,” and so on, are used correctly and when incorrectly?

Classical logic was long believed to provide the norms of reasoning, so that “thus,” “hence,” and kindred terms were thought to be used correctly as long as their use conformed to the laws of logic. For instance, in the preceding dispute, the first speaker used “hence” correctly if, and only if, D can be derived, using the rules of logic, from premises A, B, and C, or equivalently (by the completeness theorem for logic), if, and only if, the joint truth of A, B, and C guarantees the truth of D.

But in everyday English, it can be perfectly fine to present an argument using “thus” to indicate that a conclusion follows from certain premises even if the truth of the premises does *not* guarantee the truth of the conclusion. For example, it can be perfectly fine to argue

- (1) So far, Alice passed all her math exams with flying colors, thus she'll pass her upcoming math exam as well.

We may have reason to suspect that the upcoming exam is particularly challenging, or that this time Alice has not been able to prepare properly. But barring such reasons, the use of “thus” in (1) is pre-theoretically entirely appropriate.

Here, we are reminded that modern classical logic was primarily devised (by Frege, Russell, Whitehead, and others) to facilitate mathematical reasoning, and that our everyday reasoning is in many respects different from that. (Throughout the paper, by “logic,” we mean “classical logic,” given that virtually all experimental work on logical reasoning has been concerned with classical logic; in fact, it has been predominantly concerned with classical propositional logic.¹) In particular, our everyday reasoning does not normally consist of deriving, one step at a time, theorems from a system of axioms, where each step can be seen to be fully secure. Rather, we are typically trying to reason our way to a conclusion taking into account all sorts of uncertainties. For instance, in reasoning why we expect Alice to pass the upcoming exam, we may not be able to rule out entirely that Alice was unable to prepare for the exam in her usual thorough manner. So, we may not be 100 percent certain that Alice will pass the exam even though she passed all previous ones with flying colors. And yet, the argument can be valid in a pre-theoretic sense.

Backed by such observations, philosophers have proposed to regard probability theory, and possibly principles definable in probabilistic terms (such as, e.g., rules for responding to new information), as embodying the norms of reasoning. For instance, it could be argued that the first speaker in the dispute stated above used “hence” correctly if, and only if, the probability of D conditional on the conjunction of A, B, and C is close to 1. From the 1980s onward, philosophy saw a broader shift in attention from logic to probability theory, leading to the field now commonly called “formal epistemology,” which among other things studies the norms of correct non-deductive reasoning. Note that this is not to dismiss logic as being irrelevant to reasoning: probability theory *builds* on logic. It is just to say that logic can be taken to give the norms of reasoning only under very special circumstances, in which we are not dealing with uncertainties.

There has been a parallel development in the psychology of reasoning. Peter Wason, who many regard as the founding father of that field in its modern form, took for granted that

¹This is not to say that there are no exceptions; see, e.g., Stenning and van Lambalgen (2012) on non-monotonic logic.

classical logic embodies the standards of correct reasoning and was interested in whether, and the extent to which, people are able to live up to those standards. He is most famous for reporting experimental work on reasoning, apparently showing that people do quite badly in this respect (Wason 1968). In what is known as “Wason’s selection task,” replicated many times since, he showed participants four cards and told them those cards had a letter on one side and a number on the reverse. In the abstract, indicative version of the task, participants were given an indicative conditional rule of the form “If [antecedent condition] on one side, than [consequent condition] on the other side,” such as “If there is an A on one side of the card, then there is a 2 on the other side of the card,” then shown four cards, for example, cards showing an “A,” a “K,” a “2” and a “7.” Their task was to turn over all the cards, and only the cards, that would allow them to find out if the rule is true or false. Given that the cards that show the “A” and the “7” are the only ones that could provide falsifying evidence, logic suggests that those ought to be turned. That is, the normatively sanctioned selection is the [antecedent condition] card and the not-[consequent condition] card. But most participants select either the cards showing the “A” and the “2” or the “A” card alone. The standard explanation is that participants *match*: they select the cards named in the rule.

However, in psychology, too, researchers came to question the assumption that the laws of logic constitute the standards of correct reasoning. The first cracks in the wall came from examining deontic conditionals, such as “If a person drinks beer, then this person must be over 18 years of age”(Griggs and Cox 1982). This version substantially facilitated performance, with 75 percent of participants selecting the normatively sanctioned [antecedent condition] \wedge \neg [consequent condition] cards, instead of the more usual 10 percent in the indicative abstract task. It turned out that, contrary to the previous hypothesis, it is not familiarity alone that facilitates performance; it is the use of a deontic operator alongside clearly identified utility (see also Bonnefon 2009).

More substantively still, Nick Chater, Mike Oaksford, David Over, and several other researchers started to draw attention to the fact that reasoning usually takes place in a sea of uncertainty and that hence we should rather be looking for norms governing uncertain reasoning, which logic does not cover. Their proposal to look at probability theory instead led to the emergence of the so-called New Paradigm psychology of reasoning (Over 2009, Elqayam and Over 2013). An early success of this approach was Oaksford and Chater’s (1994, 1996) work on the selection task, which argued that the predominant response in Wason’s experiment was the one we should expect to find if people followed something close to an optimal probabilistic strategy for seeking information. The core idea is that people will tend to interpret the task before them as one of discovering a statistical dependence between “A” cards and “2” cards, and that by turning the “A” card and the “2” card in front of them, they are gaining the most informative evidence regarding such a dependence; at least, that is so given assumptions about people’s priors and about how best to measure information gain, which Oaksford and Chater argue to be plausible in the context of Wason’s task. One advantage of the New Paradigm is that the focus on probability and utility enables better cross-disciplinary communication with the vast and important literature on judgment and decision making. In other words, reasoning and decision-making are two facets of human thinking.

The clash between what logic would seem to prescribe and apparently sound everyday thinking and reasoning was particularly evident in studies on how people reason with indicative conditionals—sentences of the form “If A, [then] B” or “B if A,” with the antecedent being

in the indicative mood—and how they evaluate the truth values and probabilities of conditionals.² Logic gives us the so-called material conditional,³ which is true if its antecedent is false or its consequent is true; otherwise it is false. However, not many people are willing to infer

(2) If Jeff Bezos went broke, he is a billionaire,

from

(3) Jeff Bezos is a billionaire.

Yet, if conditionals are to be interpreted materially, the inference is valid. Furthermore, on the same supposition, and given Bezos' net worth, (2) is *true*, which few people may be inclined to agree with. Relatedly, whereas most people will be fully convinced that Bezos is a billionaire, few would want to assign (2) a probability of 1, even though anyone who is fully convinced of the consequent *should* do that, again assuming that natural-language conditionals are material conditionals.

Of course, once we have abandoned the idea that the laws of classical logic are the norms of reasoning, it becomes natural to explore other interpretations of the ordinary English conditional beyond the material one. That is what both philosophers and psychologists have done. In spite of all the work that went into this, however, there is little to no consensus on even the most basic questions concerning conditionals. What is the majority view on what the truth conditions of conditionals are? What is the majority view on whether conditionals have truth conditions to begin with, on the conditions under which we can assert or accept a conditional, on how we ought to evaluate the probabilities of conditionals, on how we ought to respond to the receipt of conditional information, and on and on? The answer is always the same, to wit, that there is no majority view. Importantly, it is not that philosophers tend to hold views very different from those held by psychologists. Rather, there is widespread disagreement on these matters in both camps.

Given that this has been the situation for years, one starts wondering whether there is any hope of making progress on our theoretical understanding of conditionals. Is there any one method that recommends itself here? We want to explain why we are betting on an experimental approach.

Philosophers and psychologists working on conditionals have largely focused on the same questions: questions regarding the truth conditions of conditionals, questions regarding their acceptability or assertability conditions, and questions regarding their probabilities. However, for many years, the two research communities used different methods to address these questions. Where philosophers tended to rely on conceptual analysis and formal modelling, psychologists mostly used empirical methods. But along with a growing general concern among philosophers about the reliability of conceptual analysis (which gave rise to experimental philosophy), the method came to appear especially unsuited for addressing the key

²We will only be concerned with indicative conditionals and therefore refer to them simply as “conditionals” throughout this chapter. Note that we include in this deontic conditionals, as long as they are not in the subjunctive mood.

³Strictly speaking, classical logic defines the relation of the material implication, which was not intended to capture the meaning of the natural language conditional. For more on this, see Kyburg, Teng, and Wheeler (2007). There are, however, logics of conditionals that are devised specifically to account for empirical findings on how people use conditional sentences. See, e.g., Crupi and Iacona (2020), or Berto and Özgün (2021), for recent developments.

questions about conditionals as experimental work showed some of the main philosophical accounts of conditionals, seemingly backed by sound intuitions, to be inconsistent with real-world data about how people use conditionals.

There is always the option for philosophers to dismiss such findings by saying that all these do is showing that ordinary people get confused by conditionals and have a tendency to be mistaken about their truth value, or to assign a wrong probability to them, or to deem them acceptable or assertable when in fact they are not and vice versa. But this response is not only uncharitable to ordinary people, it also makes one wonder what the point could be of having norms putatively governing conditionals if ordinary people are not able to generally stick to them, not even approximately. More importantly, it is unclear where the norms are to derive from. For example, advocates of the so-called material conditional account, according to which the semantics of the conditional is that of the material conditional, may say that one ought to respect Modus Ponens and Modus Tollens, and that one's probability for "If A, B" ought to equal one's probability for "not-A or B," that one should be willing to accept "If A, B" as soon as one is willing to accept at least one of "¬A" and "B," and so on. But it is not as though philosophers had produced an argument to the effect that, unless we bring our use of conditionals in line with the material conditional account, then we are liable to something like a Dutch book argument, say, or we are likely to engage in other behavior supposedly betokening irrationality on our part. The same holds true of any other semantics of conditionals that philosophers have developed.

In other words, there is no practical rationality justification to support the normative rationality of the material conditional—or of classical logic, for that matter. But one might ask, then, whether the use of empirical methods does not wipe out the boundary between philosophy and psychology. Why do we need two disciplines, in that case? Are we advocating for philosophers to forgo conceptual analysis, or for psychologists to become empirical philosophers? The answer, as far as we can tell, is that psychologists and philosophers differ in the research questions they tend to ask. Marr (1982) famously distinguished between three levels of analysis in conceptualizing cognitive systems: first, the computational level of analysis, the level of what the system does and why; second, the algorithmic level of analysis, the level of how this is done; and third, the implementational level of analysis, the level of hardware or wetware supporting the system. Leaving aside the latter for now, we can identify both the overlaps and the differences between psychological and philosophical research questions in this domain. Both psychologists and philosophers are interested in the computational level of analysis, especially in characterizing what the system does. In this case, for example, the semantics of conditionals as a description of how humans use conditionals in language. The difference between psychologists and philosophers is twofold: first, psychologists tend to ask more research questions about processing and representations, focusing on the algorithmic level of analysis. Second, within the computational level of analysis, philosophers ask more normative research questions (what ought we to do or to think?), whereas psychologists ask more function-related questions (what is this for?). That said, we note that psychology of reasoning is unusual within psychological science in that it attends to normative issues as well, although this has been criticized (Elqayam and Evans 2011).

If a priori theorizing about conditionals has not brought the returns that philosophers were hoping for, then perhaps philosophers should also start investing more heavily in empirical approaches. To be sure, we already said that the empirical work done by psychologists did

not bring us any closer to a consensus view. But much is still to be explored. So far, psychologists have mostly focused, on one hand, on the material conditional account, and on the other hand, on probabilistic approaches to conditionals. In this chapter, we want to make a case for an experimental philosophy of conditionals and conditional reasoning by drawing attention to a recent development, which takes important cues from philosophical work on conditionals that, we believe, has been unduly ignored. In this development—which has gone under the banner of “inferentialism”—philosophers and psychologists have joined forces from the start, and philosophical theorizing has, from the start, gone hand in hand with experimentation. Although there are still some important open questions to be answered, the empirical results obtained so far show inferentialism to be a promising new theory of conditionals.

2 The experimental philosophy of conditionals: The case of inferentialism

Humans are hardwired, from an early age on, to be attuned to all sorts of connections in the world. We often use conditionals to store and transmit the fruits of those endeavors. A parent teaching his or her child that if A, B is, in a way, handing the child a license to infer B from A. At least in principle, should the child ever receive the information that A, it can immediately draw the conclusion that B is the case if it does not believe B already at that time. We say “in principle” because there can be countervailing considerations. Perhaps the child already knows B *not* to be the case when it comes to know that A is the case. Then it may want to abandon the conditional that if A, B. All of this is compatible with the idea that conditionals embody some sort of inferential connection between their antecedent and consequent.

Pre-theoretically plausible though it may be, the idea that conditionals embody inferential connections has never become mainstream, neither in philosophy nor in psychology. Indeed, it is glaringly absent from any of the better-known semantics of conditionals. For instance, according to the material conditional account, (2) is true, as we said, even though no parent would want to hand this conditional as an inference ticket to his or her child. Similarly for Stalnaker’s (1968) possible worlds account, according to which a conditional is true if, and only if, its consequent is true in the world in which its antecedent is true that is closest to the actual world. On this account, any conditional with a true antecedent and a true consequent, however unrelated they are (e.g., “If Paris is the capital of France, Bezos is a billionaire”), comes out true. To mention a third popular proposal, according to Adams (1975), conditionals are neither true nor false, but they can be acceptable, provided their probability is sufficiently close to 1. For Adams and his followers, the probability of a conditional is the conditional probability of its consequent given its antecedent. If you are subjectively certain that a proposition A is false, then—Adams stipulates—the conditional probability of any self-consistent proposition given A is 1. So assume that you are subjectively certain that Bezos did not go broke—if he did, it would have been all over the news and you would have heard about it. Then (2) is acceptable on Adams’ account, which conflicts with our pre-theoretical judgment. And of course the account also faces the problem that Stalnaker’s faces, to wit, that any conditional about whose consequent one is subjectively certain has a conditional probability of 1 and thus is acceptable, irrespective of whether there is any pre-theoretically sensible connection between its antecedent and consequent.

This is not to say that no one ever ventured a semantics for conditionals starting from

the thought that, for a conditional to be true, its consequent must be inferable from its antecedent. The Stoic philosopher Chrysippus did (Kneale and Kneale 1962), as did, much later, Mill (1843/1872), who writes the following:

When we say, If the Koran comes from God, Mahomet is the prophet of God, we do not intend to affirm either that the Koran does come from God, or that Mahomet is really his prophet. Neither of these simple propositions may be true, and yet the truth of the [conditional] may be indisputable. What is asserted is not the truth of either of the propositions, but the inferability of the one from the other. (p. 91)

Still later, we find Ramsey (1929/1990) explicitly endorsing Mill's idea:

In general we can say with Mill that "If p , then q " means that q is inferable from p , that is, of course, from p together with certain facts and laws not stated but in some way indicated by the context. (p. 156)

We find related ideas in Ryle (1950) and Mackie (1973), and in psychology in Braine and O'Brien (1991).

Mill possibly being an exception (Skorupski 1989, p. 73 f.), what the aforementioned authors meant by a consequent being inferable from an antecedent is that the consequent follows *deductively* from the antecedent. But as Krzyżanowska, Wenmackers, and Douven (2014) point out, thus interpreted, the idea that a conditional means that its consequent is inferable from its antecedent is difficult to maintain. Douven and colleagues (2018) give the example

(4) If Betty misses her bus, she will be late for the movies,

which, as these authors argue, could well be true in a situation in which there is still a remote possibility that, after she missed the bus, Betty is transported from where she is now to the cinema to still make it in time for the movies.

It is easy to come up with further examples of plausibly true conditionals whose consequent does not follow deductively from its antecedent. That may be why the idea that conditionals embody inferential connections never gained much traction. However, as argued in Krzyżanowska, Wenmackers, and Douven (2014), there is no reason why someone attracted to the idea should want to commit to a reading of "inference" as meaning *deductive* inference.⁴ In its place, these authors propose a broader understanding on which a consequent is inferable from an antecedent if a compelling argument can be made for the consequent starting from the antecedent and whatever background assumptions are available in the context of evaluation. As they emphasize, an argument can be compelling without being conclusive. And for an argument to be compelling, it is not necessary that it consists only of deductive steps. It may include, or consists only of, inductive steps (roughly, steps based on statistical considerations), abductive steps (roughly, steps based on explanatory considerations), and perhaps other inferential steps as well (e.g., steps based on analogical considerations; see Carnap 1980, or Paris and Vencovská 2018).

To be more exact, the new proposal to cash out the idea that conditionals are intimately connected to inference—which was dubbed "inferentialism"—is both contextualist and three-valued and goes as follows: A conditional "If A, B" is true if there is a compelling argument

⁴See also Krzyżanowska (2015), Douven (2016, 2017), Douven et al. (2018, 2020), Douven, Elqayam, and Krzyżanowska (2022).



Figure 1: The soritical color series from the materials of Douven et al. (2018).

from A plus contextually determined background premises to B, with A being pivotal to that argument (i.e., with A removed, the argument would cease to be compelling), false if there is a compelling argument from A plus contextually determined background premises to the *negation* of B, and indeterminate otherwise. As Douven, Elqayam, and Krzyżanowska (2022) remark, the intuitive understanding is that any person who is justified in believing A becomes justified to believe B as soon as she becomes justified to believe “If A, B” (e.g., on the basis of testimony), supposing her being informed that if A, B, does not undermine whatever justifies her belief in A. The broad idea here is that compelling arguments allow one to carry over any justification one has for their premises to their conclusion.⁵

There is already considerable empirical support for inferentialism, some experiments also contrasting inferentialism with the earlier-mentioned mainstream semantics of conditionals. The following subsections discuss this support, which concerns the truth values assigned to conditionals (Section 2.1); probabilities assigned to conditionals (Section 2.2); and reasoning with conditionals (Section 2.3).

2.1 Truth assignments

Douven and colleagues (2018) report the outcomes of an experiment that was designed around the color patches seen in Figure 1. These patches form a so-called soritical series in that the colors of adjacent patches are very similar, while the patches get slightly greener as we move to the right, ending in a clearly green patch, and having started from a clearly blue one. The materials of the experiment consisted of conditionals pertaining to this series of patches. Each conditional had the form

If patch number i is X , then patch number j is X ,

and the participants were asked to evaluate several instances of this schema. Each instance referred to one of the patches 2, 7, 8, 9, 10, or 13, and X was either “blue” or “green,” depending on whether the participant had been assigned to the blue condition or to the green condition, a split between participants that was made for control purposes. An orthogonal split was that between the small and the large condition, which determined the values j could take. For participants in the small condition, the patch referred to in the consequent was either one or *two* steps away from the patch referred to in the antecedent, whereas for participants in the large condition, the distance between the patches was either one or *three* steps.

Douven and colleagues point out that, with each of the resulting conditionals, one can readily associate an argument. Consider, for instance,

⁵Related proposals are to be found in Oaksford and Chater (2010, 2013, 2014, 2017, 2020) as well as in van Rooij and Schulz (2019). These authors analyze the connection between a conditional’s component parts in terms of causality. It may be difficult to experimentally distinguish between these authors’ proposals and inferentialism for everyday conditionals, given that both inductive and abductive considerations tend to rest on causal relations (e.g., most explanations are causal explanations, and regularities that warrant inductive inferences are often grounded in some causal mechanism). However, Douven and colleagues (2018) found evidence for abstract conditionals, where causal relations cannot underlie the inference.

(5) If patch number 6 is green, then so is patch number 7.

This conditional can be backed by the following argument: Patches become greener as we move to the right in the color series, so from the premise that patch number 6 is green, infer that patch number 7 is green (given that patch number 7 is to the right of patch number 6). Or consider

(6) If patch number 6 is green, then so is patch number 5,

with which we can associate this argument: Adjacent patches are very similar in color, so from the premise that patch number 6 is green, infer that patch number 5 is green as well (given that patch number 5 and patch number 6 are adjacent).

As Douven and coauthors further point out, the arguments that can be associated with the conditionals in their materials can vary in strength. This is already clear from (5) and (6). Although both these conditionals refer to adjacent pairs of patches, in (5) the consequent patch is to the “greener” side of the antecedent patch, whereas in (6) the consequent patch is to the “bluer” side of the antecedent patch. While the argument we can associate with (6) is not weak, it is not quite as strong as the argument we can associate with (5). After all, for the former argument, there is a consideration that at least somewhat weakens the conclusion, which is not the case for the latter argument.

In the analysis of Douven and colleagues, there were two key determinants for argument strength in the context of their materials, to wit, *direction* (is the consequent patch to the left or to the right of the antecedent patch?) and *distance* (how close is the consequent patch to the antecedent patch?). As to why direction matters to argument strength, compare again (5) and (6) above: the argument associated with the former is stronger than that associated with the latter, given that in the former we move in the “greener” direction whereas in the latter we move in the “bluer” direction. To see why distance matters to argument strength, it is enough to compare (6) with

(7) If patch number 6 is green, then so is patch number 4.

We would associate essentially the same argument with (6) and (7), except that the argument associated with the latter is slightly weaker than that associated with the former because immediately adjacent patches are more similar to each other than patches that are separated by an intermediate patch.

The statistical analysis reported by Douven and colleagues shows that direction and distance were highly accurate predictors of the rates at which participants judged the conditionals in the materials to be true. That that would be so was predicted by inferentialism. Thus, the results of the experiment provide clear support for that semantics.⁶

Douven and colleagues (2020) further strengthen their case for inferentialism. In this paper, they compare inferentialism with the semantics of conditionals mentioned above—the material conditional account, Stalnaker’s possible worlds semantics, and Adams’ proposal—as well as some further semantics (notably, de Finetti’s three-valued semantics and some variants of that; see, e.g., de Finetti 1995 and McDermott 1996). The comparison shows inferen-

⁶To be more exact, the experiment had 532 participants, and the analysis, using mixed-effects models, revealed a main effect of direction, $\chi^2(1) = 201.66$, $p < .0001$, as well as a main effect of distance, $\chi^2(2) = 80.00$, $p < .0001$.

tialism to be clearly superior to any of the other semantics in predicting the data from Douven and colleagues (2018).⁷

We also briefly mention very recent work looking at truth evaluations of conditionals embodying *analogical* arguments. An analogical argument is one where the premise or premises support the conclusion in virtue of some similarity relation holding between items referred to in the premises and conclusion. Paris and Vencovská (2017) give the following example:

My son likes the movie *Toy Story*.
My son likes the movie *The Sound of Music*.

Plausibly, and supposing you have a son, how compelling you deem this argument will depend on how similar you deem *The Sound of Music* to be to *Toy Story*. Inferentialism predicts that, the more compelling you find the argument, the more likely you are to judge (8) to be true:

(8) If my son likes *Toy Story*, then he'll like *The Sound of Music*.

This idea inspired Douven et al. (2021) to run a two-part experiment, one part of which presented participants with various analogical arguments, and the other part of which presented them with conditionals each of which corresponded to one of the arguments presented in the first part. The analysis of the participants' responses strongly supported inferentialism, showing that how compelling an argument was to the eyes of a participant reliably predicted how likely that same participant was to judge the corresponding conditional to be true.⁸

2.2 Probabilities of conditionals

We sometimes express conditional probabilities using “if” instead of “given that” or “conditional on.” For example, it seems perfectly fine to say that the probability of throwing a 6 with a die if you throw an even number is $\frac{1}{3}$. From this observation, it is a small step to the idea that the probability of a conditional is the probability of its consequent given its antecedent, an idea now mostly referred to as either “Stalnaker’s Thesis” or the “Equation.” However, Lewis (1976) showed that, given seemingly incontrovertible assumptions, the Equation has all sorts of absurd consequences.⁹ This result stunned the community at the time, given that the Equation does sound reasonable. What made the result look even more stunning is that when psychologists turned their attention to the Equation, they did find strong empirical support for

⁷This conclusion was based on a re-analysis of the experiment from Douven et al. (2018). This analysis looked at combinations of and patterns in truth evaluations of the conditionals that were part of the materials of the said experiment as well as of their antecedents and consequents. The analysis, which was akin to (though more complicated than) a χ^2 test, compared frequencies of combinations and patterns that counted as hits or as misses, according to the various semantics. See Douven et al. (2022, Section 4) for details.

⁸In this experiment, there were 93 participants. Douven and colleagues conducted a Bayesian mixed-effects logistic regression analysis, the main finding of which was that for every one point increase on the 7-point Likert scale that these authors had used to elicit judgments of argument strength, one could expect a close to 75 percent increase in the odds that the conditional with the given argument’s premise as an antecedent and its conclusion as a consequent would be judged true. See Douven et al. (2021, Section 4.2) for details.

⁹Douven and Verbrugge (2013) show that the assumptions underlying Lewis’ argument may not all be innocuous, but we let that pass.

it: people’s conditional probabilities appeared to predict quite accurately their probability assignments to the corresponding conditionals.¹⁰

Both in light of its intuitive appeal and in light of the apparently massive empirical support, many believe the right response to Lewis’ result is to abandon the thought that conditionals express propositions. More exactly, according to this proposal conditionals are neither true nor false, and the “probabilities” we assign to conditionals are not standard probabilities of truth but rather degrees of acceptability. This blocks Lewis’ argument against the Equation, given that that requires the usual logical operators—which are *propositional* operators—to apply to conditionals. But it also comes at a steep cost, precisely because now it becomes puzzling how conditionals can truth-functionally combine with other parts of the language, as it would seem they can. For example, we have no difficulty understanding conjunctions or disjunctions of conditionals.

Recent experimental results suggest that there may be no need to address Lewis’ result head on, simply because there may be independent reason to abandon the Equation. Spohn (2015) may have been the first to question the empirical results supposedly supporting the Equation. As he rightly notes, all materials used in the relevant experiments consisted of conditionals whose antecedent was positively probabilistically relevant to their consequent, meaning that the probability of the consequent given the antecedent was higher than the unconditional probability of the consequent. This is not because the conditionals were explicitly selected to have that feature. It is rather that most or even all conditionals we encounter in quotidian speech *have* that feature, and psychologists working on conditionals were interested in people’s responses to *normal* conditionals. Would conditional probabilities still be found to predict the probabilities of the corresponding conditionals if experimenters presented participants with conditionals whose antecedent is negatively probabilistically relevant to their consequent, or is probabilistically irrelevant to the consequent?

Skovgaard-Olsen, Singman, and Klauer (2016) set out to answer this question. They found that conditional probabilities reliably predict people’s probability assignments to conditionals as long as the positive relevance condition holds, but that they fail to do so quite badly in those cases in which the said condition does not hold. While Skovgaard-Olsen and colleagues explain their findings by postulating that the relevance belongs to the “core meaning” of conditional, that is its semantic content, others, for instance, Over and colleagues (2007), suggested that “the use of a conditional *pragmatically suggests*, in certain ordinary contexts, that *p* raises the probability of *q* or that *p* causes *q*” (p. 92).

Elsewhere, we have explained why we prefer semantic over pragmatic explanations of the relevance effect in the data from Skovgaard-Olsen, Singmann, and Klauer (2016).¹¹ And as shown in Douven, Elqayam, and Mirabile (2021), inferentialism provides a perfectly good semantic explanation. These authors start by working out the implications of inferentialism for the probabilities of conditionals. Unpacking the truth conditions that inferentialism assigns to conditionals, Douven and coauthors note that, as standardly understood, probabilities are probabilities of *truth*. Thus, the probability of “If A, B” is the probability that “If A, B” is *true*, which is the probability that the truth conditions of “If A, B” are realized, which finally, as-

¹⁰See, e.g., Oaksford and Chater (2003, 2007), Over and Evans (2003), Evans and Over (2004), Oberauer, Weidenfeld, and Fischer (2007), Over et al. (2007), Douven and Verbrugge (2010, 2013), Fugard et al. (2011), Over, Douven, and Verbrugge (2013), as well as van Wijnbergen-Huitink, Elqayam, and Over (2015).

¹¹See Douven, Elqayam, and Krzyżanowska (2022); see also Krzyżanowska, Collins, and Hahn (2021) as well as Rostworowski, Pietrulewicz, and Będkowski (2021).

suming inferentialism, is the probability that there is a compelling argument from the conditional's antecedent (plus background knowledge) to the conditional's consequent.

Douven, Elqayam, and Mirabile (2021) further note that it is often not immediately obvious to us whether there is a compelling argument for a given proposition starting from a second proposition in conjunction with whatever our background knowledge happens to be. Right now, for instance, we are inclined to believe that a compelling case can be made for the claim that the United States will get the COVID-19 outbreak under control on the supposition that its government can convince at least 80 percent of the population to get a vaccine. At the moment, however, this is really only an inclination: we are by no means sure and would have to think the matter through more carefully and see in particular whether we are not overlooking factors that might contribute to a continuation of the pandemic even in a fully or near-to-fully vaccinated population (such as, most notably, the emergence of new variants of the SARS-CoV-2 virus that the current vaccines offer insufficient protection to). But if we now had to answer the question what probability we assign to

- (9) If the United States government can convince at least 80 percent of its population to get vaccinated, they will get the COVID-19 outbreak under control,

we would estimate the likelihood that we can make a compelling case for the consequent, starting from the antecedent plus background knowledge, and would give that as our answer. Douven, Elqayam, and Mirabile (2021) argue that we do this by relying on a heuristic of gauging the inferential strength between antecedent and consequent.

Douven, Elqayam, and Mirabile (2021) report two experiments designed to test this “inference heuristic.” Each experiment presented participants with three tasks, which used the same set of 50 conditionals. One task asked participants for their probability for each of those conditionals; a second task, which participants received a week after the first, asked them to indicate, for each conditional, how strongly in their opinion the consequent followed from the antecedent; and the third task, which the participants received a week after the second, determined their conditional probabilities corresponding to the conditionals, these conditional probabilities being measured via a probabilistic truth-table task in one experiment and by asking participants to engage in suppositional thinking in the other experiment. For instance, one of the conditionals they were presented in the first part was

- (10) If a cure for AIDS will be discovered, condom sales will drop.

They were asked to indicate how probable this conditional was, in their opinion. Then in the second part, participants were asked, among other things, “Suppose that a cure for AIDS is discovered. How strongly do you agree that it then follows that condom sales will drop?” They were supposed to answer on a 7-point Likert scale that ranged from “Strongly disagree” to “Strongly agree,” with “Neither agree nor disagree” as the midpoint. And then the corresponding question in the final task was, in the first experiment, to rate each of the following situations on a probability scale ranging from 0 to 100 percent:

It is TRUE that a cure for AIDS will be discovered and it is also TRUE that condom sales will drop	...
It is TRUE that a cure for AIDS will be discovered but it is FALSE that condom sales will drop	...
It is FALSE that a cure for AIDS will be discovered but it is TRUE that condom sales will drop	...
It is FALSE that a cure for AIDS will be discovered and it is also FALSE that condom sales will drop	...

100

Here, participants were instructed to make sure that the probabilities summed to 100. In the second experiment, they were not shown such probabilistic truth tables but asked to suppose that a cure for AIDS will be discovered and then to assess, under that supposition, the probability that condom sales will drop.

Across both experiments, inference strength judgments were found to accurately predict probability ratings, in line with what one would expect on the basis of inferentialism. Moreover, inference strength judgments were also found to predict probability ratings much more accurately than conditional probability ratings, an outcome that strongly favors inferentialism over Adams' (1975) account.¹²

2.3 Reasoning with conditionals

A third major research area concerning conditionals, next to their truth conditions and probabilities, is the inferences they license. While this is again an area of vast disagreement, virtually all researchers agree that the conditional operator should validate Modus Ponens (MP): from A and "If A, B" we should be allowed to infer B. After all, this is a rule we rely on quite routinely in our reasoning. However, as Krzyżanowska, Wenmackers, and Douven (2014) acknowledged right away, their position does *not* validate MP, simply because we may deem an argument from a true premise A to a conclusion B compelling (making "If A, B" true, given our background knowledge), but, unbeknownst to us, B may be false.¹³

As Krzyżanowska and co-authors also pointed out, however, the fact that inferentialism invalidates MP does not mean that, from an inferentialist perspective, there is anything wrong with our practice of relying on that rule of inference. To the contrary, from that perspective, the designated practice is perfectly fine. For consider that, typically, when we have a compelling argument from A to B, and A is true, then B will be true as well. As Schurz and Hertwig (2019) argue, we rely on compelling-but-inconclusive arguments much more frequently in our daily lives than we rely on deductively valid arguments. Thus, it would be a serious problem if the arguments we judge to be compelling were not highly truth-conducive. But if that is so, then, from an inferentialist perspective, MP is a highly reliable inference rule, which typically yields a true conclusion when applied to true premises. Also, as McGee (1985) argued, it suffices to account for the intuition that it is perfectly alright to rely on MP that MP be highly reliable, given that we cannot expect our intuitions about the validity of an inference rule to be sensitive to the difference between that rule being guaranteed to preserve truth and it preserving truth with near-certainty.

In fact, psychologists have been long looking at MP. In multiple experiments, they found that whereas MP was typically *highly* endorsed, it seldom was *universally* endorsed.¹⁴ That is due to processing factors and perhaps some noise—but not entirely. As New Paradigmers argued, experimenters can request their participants to suppose the premises of an argument,

¹²In the first experiment there were 118 participants, in the second there were 204. The data from both experiments were analyzed using Bayesian mixed-effects linear models. In both analyses, the best models were those with both inference strength responses and conditional probabilities as predictors of the probabilities of conditionals. However, also in both analyses, the former had a much bigger impact than the latter. For details, see Sections 2.2 and 3.2.

¹³We say "unbeknownst to us" because we are unlikely to regard any argument as compelling if we know its conclusion to be false.

¹⁴Specifically, endorsement rates ranged between 89 and 100 percent.

but they should still reckon with the possibility that a participant's own beliefs about those premises will have some impact on his or her judgments about whether a certain conclusion follows. Most notably, Stevenson and Over (2001) found that uncertainty about the major premise in an MP argument tends to diminish a participant's willingness to endorse the conclusion.

Mirabile and Douven (2021) examined the endorsement rates of the conclusions of MP arguments with an eye toward testing inferentialism. More specifically, they were interested in whether a participant's judgment of the strength of the argument embodied by the major premise of an MP argument would predict the likelihood that that participant would endorse the conclusion of the argument. They were further interested in contrasting the predictive power of such inferential strength judgments with that of the judged probability of the consequent of the major premise given its antecedent.

They conducted a three-part experiment whose main materials consisted of a number of MP arguments. One part of the experiment was meant to measure the conditional probabilities corresponding to the conditionals that served as major premises in those argument, another part was meant to measure the strength of the inferential connection between those conditionals' antecedents and consequents, and the third was meant to measure endorsement rates of the conclusions of the arguments. In their analysis, Mirabile and Douven found, in support of inferentialism, that whereas conditional probability was a good predictor of conclusion endorsement, argument strength was a significantly better predictor.¹⁵

It may be helpful to explain in a little more detail how Mirabile and Douven see their results as supporting inferentialism. As they point out, from an inferentialist perspective, conditionals look a bit like *pipes* or *conduits* in that, if accepted, a conditional allows one to transfer whatever grounds one has for believing its antecedent to its consequent. That, after all, is what a compelling argument does: transferring whatever grounds one has for believing the premises to grounds for believing the conclusion. As mentioned, however, that an argument is compelling does not imply that it is conclusive. Because of that—Mirabile and Douven argue—conditionals are to be thought of as *leaky* pipes, where the leakiness can vary in degree: the argument they embody may not be strong enough to carry over *all* the support we have or may have for the antecedent to the consequent. The part aimed at measuring inference strength can be thought of as having measured the degree of leakiness, and the degree of leakiness of the major premise of an MP argument turned out to predict with high accuracy whether a participant was to endorse the conclusion of that argument.

2.4 Open questions

Inferentialism, in the version described in the foregoing, is still a young position, and there remain a number of questions to be answered. We mention two in particular. One concerns the learning of conditionals. It consists of two sub-questions, to wit, the question of how people actually adapt their beliefs to the receipt of conditional information, and the question of how they ought to adapt their beliefs to the receipt of such information. The learning of con-

¹⁵In this experiment, there were 120 participants. In the analysis, Mirabile and Douven fit a number of Bayesian cumulative ordinal regression models to the data from these participants, all models having endorsement rates as dependent variable. The model that did best had both inference strength and conditional probability as predictors, but that model showed the former to have a much bigger impact on the data than the latter. For details, see Mirabile and Douven (2021, Section 6.2).

ditionals is badly understudied, both in philosophy and in psychology. As shown in Douven (2012), standard mechanisms for updating our beliefs can give counterintuitive results when applied to conditionals. A worked-out proposal for an update rule that does give satisfactory results for such applications is still lacking. A good approach may be to first gather more data on how people actually react when they learn a conditional. Work on this has just begun.

We know, for instance, that upon receiving a testimony of the form “If A, C”, which we will refer to as a “conditional testimony,” the participants increase their conditional probability ratings. Moreover, the participants’ posterior conditional probability ratings are higher when they receive the conditional testimony from a highly reliable speaker, for instance, a professor of medicine making assertions about a patient’s prognosis, than when the same conditional is asserted by a less reliable speaker, such as a medical student (Collins et al., 2020). But it is not only the conditional probability that people adjust upon receiving a conditional testimony. In a follow-up to the paper from Collins and colleagues, Krzyżanowska, Collins, and Hahn (2020) report that the perceived strength of the (probabilistic) relevance relation, estimated as the difference between $\Pr(C | A)$ and $\Pr(C | \neg A)$, also increases in response to someone’s assertion of a conditional, and the extent of this increase depends on the reliability of the speaker, too. Furthermore, Collins (2017) found that people’s conditional probability judgments increase to a greater extent when the same testimony comes from multiple sources than when the conditional is asserted by a single speaker. Collins and colleagues (2020) collected not only the conditional probability ratings but also the probability estimates of the relevant antecedents and consequents on their own. Perhaps unsurprisingly, the participants did not adjust their probability ratings for the antecedents and consequents if, prior to receiving the conditional testimony, they found them as likely as not, that is, both the prior and posterior probability estimates for these antecedents and consequents were close to 0.5. However, participants did increase their probability ratings for the antecedents and consequents whose prior probability was judged to be low.¹⁶ Interestingly, while these results present a rather unsurprising and intuitive data pattern, they turned out to pose a significant modelling challenge. In their extensive theoretical discussion, Collins and colleagues conclude that none of the mainstream theories of conditionals can account for all of the reported findings in a straightforward way.¹⁷

The other question we want to mention concerns nested conditionals. Philosophers have been struggling to come to grips with such conditionals, and empirical studies devoted to nested conditionals are far and few between. There is some work relevant to the so-called Import–Export principle, according to which a nested conditional of the form “If A then if B, C” is equivalent to the simple conditional “If A and B, then C,” but the results are mixed (Douven and Verbrugge 2013, van Wijnbergen-Huitink, Elqayam, and Over 2015), calling for further studies. Van Wijnbergen-Huitink, Elqayam, and Over (2015) used abstract conditionals such as “If the chip is square, then if it is large, it is white” in a betting task, with two tasks, a probability task and a categorical truth task, both presented in a betting format. They found no differences between the probabilities of the iterated versus the imported form. However, in the categorical task, more imported forms than iterated forms were consistent with the de-

¹⁶Collins and colleagues (2020) also looked at high probability antecedents and consequents, but found only a slight, non-significant decrease in the probability ratings. However, the prior probability estimates for what was supposed to be high probability antecedents and consequents were less extreme—that is, farther away from the end point of the scale—than the estimates for their low probability counterparts.

¹⁷Though see Hartmann and Hahn (2020) for a new formal proposal devised specifically to account for the empirical data on updating with conditionals.

fective truth table (according to which a conditional with an antecedent that is either false or has an indeterminate “truth value” is itself indeterminate), an effect attributed to processing difficulty.

3 Further topics

We have focused on the experimental philosophy of conditionals and conditional reasoning because, first, it is an area that has seen a lot of recent activity, and second, much of that activity consisted of collaborative projects involving both philosophers and psychologists. We firmly believe that such collaborations have the best chance of leading to high-quality work in experimental philosophy and so we recommend that philosophers interested in doing experiments try to team up with colleagues from the psychology department.

In this section, we briefly mention some other research examining empirically topics that are of direct concern to logicians and formal epistemologists. The first topic we want to mention is the study of the Liar paradox (“This sentence is false”—which is false if it is true and true if it is false). While the roots of the Liar go way back, it is in the twentieth century that philosophers such as Tarski (1944) and Kripke (1975) famously used it as a focal consideration in theories of truth. Notwithstanding the massive impact this work had on philosophy, little psychological work has been done with the Liar. Elqayam (2006; also Elqayam et al., 2008) embedded Liar-type propositions in truth-table tasks, with the Truthteller (“This sentence is true”) as a control. Conditionals with a Liar component tended to be evaluated as indeterminate, whereas conditionals with the Truthteller (equally indeterminate but not paradoxical) were “collapsed,” that is, they were treated as if the Truthteller were simply true. In terms of processing, this is evidence to the difficulty people have in effort-laden computations such as those required in tracking multiple iterations.

Another interesting example is what has been dubbed in the psychological literature “logical intuitions.” The idea is that people are able to provide fast, automatic logical responses (e.g., Handley et al., 2011), or at least identify when their responses went wrong (De Neys, 2012), even if they were powerless to revise them. The term “logical intuitions” can be misleading, insofar as it might create the misapprehension that humans might have intuitions for classical logic implanted in their brains. As more recent work identified (Ghasemi et al., preprint), such intuitions are anything but logical; rather, they are generated by fast processing of superficial cues, whose outputs happen to correlate with logical responding.

Another topic we want to mention is the experimental study of various forms of non-deductive reasoning, such as work on analogical reasoning (e.g., Spellman and Holyoak 1992), or informal reasoning and argumentation (e.g., Hahn and Oaksford 2007; Mercier and Sperber 2011). A substantial body of work has gone into probabilistic reasoning. As intimated, it is a core tenet of both formal epistemology and the New Paradigm psychology of reasoning that reasoning is, most fundamentally, probabilistic in nature. Psychologists have been mainly interested in the extent to which “probabilistic” can be taken literally, that is, the extent to which people obey the postulates of probability theory in their reasoning. Famously, Tversky and Kahneman (1983) published results showing that people sometimes assign a probability to a conjunction that exceeds the probability they assign to the (in their eyes) least probable conjunct, in violation of probability theory. Collaboration between philosophers and psychologists led to a series of papers defending a plausible explanation of Tversky and Kahneman’s

results (see Crupi, Tentori, and Gonzalez 2007, Crupi, Fitelson, and Tentori 2008, as well as Tentori, Crupi, and Russo 2013). According to these authors, people often attend more closely to probabilistic confirmation than to probability per se, but because the notions are so closely related, they can get easily conflated in people's minds.

Work on abductive reasoning—a form of non-deductive reasoning guided by explanatory considerations—is reported in Douven and Schupbach (2015a, 2015b).¹⁸ Philosophers have long argued that we may be justified in believing something because it best explains the evidence in our possession (see, e.g., Boyd 1985, McMullin 1992, Lipton 2004). More recently, Bayesian philosophers of science have rejected this idea as—according to them—it can lead to probabilistic incoherence, which they see as a token of irrationality (van Fraassen, 1989). Douven and Schupbach showed experimentally that judgments of explanation quality had a significant impact on their participants' belief updates and explained why those updates tended to deviate from Bayesian prescriptions. Whereas that might just go to show that Bayesianism is not descriptively adequate, Douven (2022) argues that people may be right to update their beliefs on the basis of explanatory considerations and that there is no reason to view this as a sign of irrationality, even if it leads to discrepancies with Bayesian norms.

We already briefly touched upon analogical reasoning, which relies on similarity relations. So far, there has been little contact between philosophers and psychologists studying this form of reasoning. Carnap (1980), Kuipers (1988), Niiniluoto (1988), and other philosophers have been mainly concerned with trying to formalize analogical reasoning, whereas psychologists have—much of it under the heading of “category-based induction”—experimentally investigated the role of similarity in inference (e.g., to what extent people's willingness to infer that cows have a certain property from the premise that horses have that property depends on how similar they judge cows to be to horses). Osta-Vélez and Gärdenfors (2020) make an explicit attempt to connect philosophical work on analogical reasoning with the aforementioned psychological research. The main result of their paper is an account of analogical reasoning based on the so-called conceptual spaces framework (Gärdenfors, 2000). Douven and colleagues (2022) report empirical support for this new proposal.

Finally, formal epistemologists have expended much time and effort on analyzing the notion of coherence. In mainstream epistemology, Bonjour (1985) and others had proposed that coherence was key to a theory of justification. Specifically, the idea was that the more coherent a set of beliefs is—the more those beliefs hang together—the more justified we are in holding those beliefs. This idea was challenged on the grounds that there was nothing to suggest that coherence is truth-conducive and that it hence was unclear that coherence can play any role in a theory of justification (see, e.g., Klein and Warfield, 1994). That, at the time, we only had an informal understanding of the notion of coherence made the challenge hard to address. Realizing this, various researchers set out to make coherence formally precise, which led to a great number of probabilistic measures of coherence; see, among many other publications all proposing different measures, Shogenji (1999), Olsson (2002), Bovens and Hartmann (2003), Fitelson (2003), Douven and Meijs (2007), Schippers (2014), as well as Schippers and Schurz (2017).

Most authors working on probabilistic measures of coherence supported their proposals by arguing that they gave verdicts about cases that aligned with our intuitions about those

¹⁸See also Lombrozo (2016), Walker et al. (2017), Wojtowicz and DeDeo (2020), Douven (2021), as well as Jern, Derrow-Pinion, and Piergiovanni (2021).

cases. It is unfortunate that, to this day, few attempts have been made to subject the various proposals to empirical testing. Two notable exceptions are Harris and Hahn (2009) and Koscholke and Jekel (2017). An exception of sorts is Angere (2008), which uses computer simulations to determine which of the measures is most conducive to the truth. Whereas computer simulations have not gained any prominence so far in experimental philosophy, it is arguable that they can give empirical support of sorts, much in the way in which they can in physics and the natural sciences generally (Galison, 1997).¹⁹

References

- Adams, Ernest W. (1975). *The Logic of Conditionals*. Dordrecht: Reidel.
- Angere, Stefan (2008). Coherence as heuristic. *Mind*, 117, 1–26.
- Berto, Francesco & Özgün, Aybüke (2021). Indicative conditionals: probabilities and relevance. *Philosophical Studies*, 178, 3697–3730.
- Bonjour, Laurence (1985). *The Structure of Empirical Knowledge*. Cambridge MA: Harvard University Press.
- Bonnefon, Jean-Francois (2009). A theory of utility conditionals: Paralogical reasoning from decision-theoretic leakage. *Psychological Review*, 116, 888–907.
- Bovens, Luc & Hartmann, Stephan (2003). *Bayesian Epistemology*. Oxford: Oxford University Press.
- Boyd, Richard N. (1985). Lex orandi est lex credendi. In P. Churchland & C. Hooker (eds.), *Images of Science* (pp. 3–34). Chicago: University of Chicago Press.
- Braine, Martin D. S. & O'Brien, David P. (1991). A theory of *if*: Lexical entry, reasoning program, and pragmatic principles. *Psychological Review*, 98, 182–203.
- Carnap, Rudolf (1980). A basic system of inductive logic II. In R. C. Jeffrey (ed.), *Studies in Inductive Logic and Probability* (pp. 7–155). Berkeley CA: University of California Press.
- Collins, Peter J. (2017). *Rationality, Pragmatics, and Sources*. Doctoral dissertation. Birkbeck: University of London.
- Collins, Peter J., Krzyżanowska, Karolina H., Hartmann, Stephan, Wheeler, Gregory, & Hahn, Ulrike (2020). Conditionals and testimony. *Cognitive Psychology*, 122, 101329.
- Crupi, Vincenzo & Iacona, Andrea (2020). The evidential conditional. *Erkenntnis*, in press.
- de Finetti, Bruno (1995). The logic of probability. *Philosophical Studies*, 77, 181–190.
- De Neys, Wim (2012). Bias and conflict: A case for logical intuitions. *Perspectives on Psychological Science*, 7, 28–38.
- Douven, Igor (2012). Learning conditional information. *Mind & Language*, 27, 239–263.
- Douven, Igor (2016). *The Epistemology of Indicative Conditionals*. Cambridge: Cambridge University Press.
- Douven, Igor (2017). How to account for the oddness of missing-link conditionals. *Synthese*, 194, 1541–1554.
- Douven, Igor (2021). How explanation guides belief change. *Trends in Cognitive Sciences*, 25, 829–830.
- Douven, Igor (2022). *The Art of Abduction*. Cambridge MA: MIT Press, in press.
- Douven, Igor, Elqayam, Shira, Gärdenfors, Peter, & Mirabile, Patricia (2021). Conceptual spaces and the strength of similarity-based arguments. *Cognition*, in press.

¹⁹We are indebted to the editors for valuable comments on a previous version.

- Douven, Igor, Elqayam, Shira, & Krzyżanowska, Karolina H. (2022). Inferentialism: A manifesto. In S. Kaufmann & D. E. Over (eds.), *Conditionals: Logic, Linguistics, and Psychology*. London: Palgrave Macmillan, in press.
- Douven, Igor, Elqayam, Shira, & Mirabile, Patricia (2021). Inference strength predicts the probability of conditionals better than conditional probability does. *Journal of Memory and Language*, in press.
- Douven, Igor, Elqayam, Shira, Singmann, Henrik, & van Wijnbergen-Huitink, Janneke (2018). Conditionals and inferential connections: A hypothetical inferential theory. *Cognitive Psychology*, *101*, 50–81.
- Douven, Igor, Elqayam, Shira, Singmann, Henrik, & van Wijnbergen-Huitink, Janneke (2020). Conditionals and inferential connections: Toward a new semantics. *Thinking & Reasoning*, *26*, 311–351.
- Douven, Igor & Verbrugge, Sara (2010). The Adams family. *Cognition*, *117*, 302–318.
- Douven, Igor & Verbrugge, Sara (2013). The probabilities of conditionals revisited. *Cognitive Science*, *37*, 711–730.
- Elqayam, Shira (2006). The collapse illusion effect: A pragmatic–semantic illusion of truth and paradox. *Thinking & Reasoning*, *12*, 144–180.
- Elqayam, Shira & Evans, Jonathan St. B. T. (2011). Subtracting “ought” from “is”: Descriptivism versus normativism in the study of human thinking. *Behavioral and Brain Sciences*, *34*, 233–248.
- Elqayam, Shira, Handley, Simon J., Evans, Jonathan St. B. T., & Bacon, Alison M. (2008). On some limits of hypothetical thinking. *Quarterly Journal of Experimental Psychology*, *61*, 784–808.
- Elqayam, Shira & Over, David E. (2013). New paradigm psychology of reasoning. *Thinking & Reasoning*, *19*, 249–265.
- Evans, Jonathan St. B. T. & Over, David E. (2004). *If*. Oxford: Oxford University Press.
- Fitelson, Branden (2003). A probabilistic theory of coherence. *Analysis*, *63*, 194–199.
- Fugard, Andrew, Pfeifer, Niki, Mayerhofer, Bastian, & Kleiter, Gernot (2011). How people interpret conditionals: Shifts toward the conditional event. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *37*, 635–648.
- Galison, Peter (1997). *Image and Logic*. Chicago: University of Chicago Press.
- Gärdenfors, Peter (2000). *Conceptual Spaces*. Cambridge MA: MIT Press.
- Ghasemi, Omid, Handley, Simon J., Howarth, Stephanie, Newman, Ian R., & Thompson, Valerie A. (preprint). Logical intuition is not really about logic. Retrieved from <http://dx.doi.org/10.31219/osf.io/wtupm>.
- Griggs, Richard A. & Cox, James R. (1982). The elusive thematic materials effect in the Wason selection task. *British Journal of Psychology*, *73*, 407–420.
- Hahn, Ulrike & Oaksford, Mike (2007). The rationality of informal argumentation: A Bayesian approach to reasoning fallacies. *Psychological Review*, *114*, 704–732.
- Handley, Simon J., Newstead, Stephen E., & Trippas, Dries (2011). Logic, beliefs, and instruction: A test of the default interventionist account of belief bias. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *37*, 28–43.
- Harris, Adam J. L. & Hahn, Ulrike (2009). Bayesian rationality in evaluating multiple testimonies: Incorporating the role of coherence. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *35*, 1366–1373.

- Hartmann, Stephan & Hahn, Ulrike (2020). A new approach to testimonial conditionals. *Proceedings of the 42nd Annual Conference of the Cognitive Science Society*, 981–986.
- Jern, Alan, Derrow-Pinion, Austin, & Piergiovanni, AJ (2021). A computational framework for understanding the roles of simplicity and rational support in people's behavior explanations. *Cognition*, 210, 104606, <https://doi.org/10.1016/j.cognition.2021.104606>.
- Klein, Peter & Warfield, Ted A. (1994). What price coherence? *Analysis*, 54, 129–132.
- Kneale, William & Kneale, Martha (1962). *The Development of Logic*. Oxford: Oxford University Press.
- Koscholke, Jakob & Jekel, Mark (2017). Probabilistic coherence measures: A psychological study of coherence assessment. *Synthese*, 194, 1303–1322.
- Kripke, Saul (1975). Outline of a theory of truth. *Journal of Philosophy*, 72, 690–716.
- Krzyżanowska, Karolina H. (2015). *Between “If” and “Then.”* Doctoral dissertation. University of Groningen.
- Krzyżanowska, Karolina H., Collins, Peter J., & Hahn, Ulrike (2017). Between a conditional's antecedent and its consequent: Discourse coherence vs. probabilistic relevance. *Cognition*, 164, 199–205.
- Krzyżanowska, Karolina H., Collins, Peter J., & Hahn, Ulrike (2021). True clauses and false connections. *Journal of Memory and Language*, in press.
- Krzyżanowska, Karolina H., Wenmackers, Sylvia, & Douven, Igor (2014). Rethinking Gibbard's river-boat argument. *Studia Logica*, 102, 771–792.
- Krzyżanowska, Karolina H., Collins, Peter J., & Hahn, Ulrike (2020). . . . *that P is relevant for Q*: Indicative conditionals and learning from testimony. *Proceedings of the 42nd Annual Conference of the Cognitive Science Society*, 987–993.
- Kuipers, Theo A. F. (1988). Inductive analogy by similarity and proximity. In D. H. Helman (ed.), *Analogical Reasoning* (pp. 299–313). Dordrecht: Kluwer.
- Kyburg, Henry, Teng, Choh M., & Wheeler, Gregory (2007). Conditionals and consequences. *Journal of Applied Logic*, 5, 638–650.
- Lewis, David K. (1976). Probabilities of conditionals and conditional probabilities. *Philosophical Review*, 85, 297–315.
- Lipton, Peter (2004). *Inference to the Best Explanation* (2nd ed.). London: Routledge.
- Lombrozo, Tania (2016). Explanatory preferences shape learning and inference. *Trends in Cognitive Sciences*, 20, 748–759.
- Mackie, John L. (1973). *Truth, Probability and Paradox: Studies in Philosophical Logic*. Oxford: Oxford University Press.
- Marr, David (1982). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. San Francisco: Freeman.
- McGee, Vann (1985). A counterexample to modus ponens. *Journal of Philosophy*, 82, 462–471.
- McMullin, Ernan (1992). *The Inference that Makes Science*. Milwaukee: Marquette University Press.
- McDermott, Michael (1996). On the truth conditions of certain “if”-sentences. *Philosophical Review*, 105, 1–37.
- Mercier, Hugo & Sperber, Dan (2011). Why do humans reason? Arguments for an argumentative theory. *Behavioral and Brain Sciences*, 34, 57–74.
- Mill, John S. (1843/1872). *A System of Logic: Ratiocinative and Inductive* (8th ed.). London: Long-

- mans, Green, Reader, & Dyer.
- Mirabile, Patricia & Douven, Igor (2020). Abductive conditionals as a test case for inferentialism. *Cognition*, 200, 104232, doi:10.1016/j.cognition.2020.104232.
- Niiniluoto, Ilkka (1988). Analogy and similarity in scientific reasoning. In D. H. Helman (ed.), *Analogical Reasoning* (pp. 271–298). Dordrecht: Kluwer.
- Oaksford, Mike & Chater, Nick (2003). Conditional probability and the cognitive science of conditional reasoning. *Mind & Language*, 18, 359–379.
- Oaksford, Mike & Chater, Nick (2007). *Bayesian Rationality: The Probabilistic Approach to Human Reasoning*. Oxford: Oxford University Press.
- Oaksford, Mike & Chater, Nick (2010). Causation and conditionals in the cognitive science of human reasoning. *The Open Psychology Journal*, 3, 105–118.
- Oaksford, Mike & Chater, Nick (2013). Dynamic inference and everyday conditional reasoning in the new paradigm. *Thinking & Reasoning*, 19, 346–379.
- Oaksford, Mike & Chater, Nick (2014). Probabilistic single function dual process theory and logic programming as approaches to non-monotonicity in human vs. artificial reasoning. *Thinking & Reasoning*, 20, 269–295.
- Oaksford, Mike & Chater, Nick (2017). Causal models and conditional reasoning. In M. R. Waldmann (ed.), *The Oxford Handbook of Causal Reasoning* (pp. 327–346). Oxford: Oxford University Press.
- Oaksford, Mike & Chater, Nick (2020). Integrating causal Bayes nets and inferentialism in conditional inference. In S. Elqayam, I. Douven, J. St. B. T. Evans, & N. Cruz (eds.), *Logic and Uncertainty in the Human Mind* (pp. 116–132). London: Routledge.
- Oberauer, Klaus, Weidenfeld, Andrea, & Fischer, Katrin (2007). What makes us believe a conditional? The roles of covariation and causality. *Thinking & Reasoning*, 13, 340–369.
- Olsson, Erik (2002). What is the problem of coherence and truth? *Journal of Philosophy*, 94, 246–272.
- Osta-Vélez, Matias & Gärdenfors, Peter (2020). Category-based induction in conceptual spaces. *Journal of Mathematical Psychology*, 96, 102357, doi: 10.1016/j.jmp.2020.102357.
- Over, David E. (2009). New paradigm psychology of reasoning. *Thinking & Reasoning*, 15, 431–438.
- Over, David E., Douven, Igor, & Verbrugge, Sara (2013). Scope ambiguities and conditionals. *Thinking & Reasoning*, 19, 284–307.
- Over, David E. & Evans, Jonathan St. B. T. (2003). The probability of conditionals: The psychological evidence. *Mind & Language*, 18, 340–358.
- Over, David E., Hadjichristidis, Constantinos, Evans, Jonathan St. B. T., Handley, Simon J., & Sloman, Steven A. (2007). The probability of causal conditionals. *Cognitive Psychology*, 54, 62–97.
- Paris, Jeffrey B. & Vencovská, Alena (2017). Combining analogical support in pure inductive logic. *Erkenntnis*, 82, 401–419.
- Ramsey, Frank P. (1929/1990). General propositions and causality. In his *Philosophical Papers*, edited by D. H. Mellor (pp. 145–163). Cambridge: Cambridge University Press.
- Rostworowski, Wojciech, Pietrulewicz, Natalia, & Będkowski, Marcin (2021). Conditionals and specific links: An experimental study. *Synthese*, in press.
- Ryle, Gilbert (1950). “If,” “so,” and “because.” In M. Black (ed.), *Philosophical Analysis* (pp. 323–340). Ithaca NY: Cornell University Press.

- Schippers, Michael (2014). Probabilistic measures of coherence: From adequacy constraints towards pluralism. *Synthese*, 191, 3821–3845.
- Schippers, Michael & Schurz, Gerhard (2017). Genuine coherence as mutual confirmation between content elements. *Studia Logica*, 105, 299–329.
- Schurz, Gerhard & Hertwig, Ralph (2019). Cognitive success: A consequentialist account of rationality and cognition. *Topics in Cognitive Science*, 11, 7–36.
- Shogenji, Tomoji (1999). Is coherence truth-conducive? *Analysis*, 59, 338–345.
- Skorupski, John (1989). *John Stuart Mill*. London: Routledge.
- Skovgaard-Olsen, Niels, Singmann, Henrik, & Klauer, Klaus C. (2016). The relevance effect and conditionals. *Cognition*, 150, 26–36.
- Spellman, Barbara A. & Holyoak, Keith J. (1992). If Saddam is Hitler then who is George Bush? Analogical mapping between systems of social roles. *Journal of Personality and Social Psychology*, 62, 913–933.
- Spohn, Wolfgang (2015). Conditionals: A unifying ranking-theoretic perspective. *Philosophers' Imprint*, 15, <http://hdl.handle.net/2027/spo.3521354.0015.001>.
- Stalnaker, Robert (1975). Indicative conditionals. *Philosophia*, 5, 269–286.
- Stenning, Keith & van Lambalgen, Michiel (2008). *Human Reasoning and Cognitive Science*. Cambridge MA: MIT Press.
- Tarski, Alfred (1944). The semantic conception of truth and the foundations of semantics. *Philosophy and Phenomenological Research*, 4, 341–376.
- Tversky, Amos & Kahneman, Daniel (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review*, 90, 293–315.
- van Fraassen, Bas C. (1989). *Laws and Symmetry*. Oxford: Oxford University Press.
- van Rooij, Robert & Schulz, Katrin (2019). Conditionals, causality and conditional probability. *Journal of Logic, Language and Information*, 28, 55–71.
- Walker, Caren M., Lombrozo, Tania, Williams, Joseph J., Rafferty, Anna N., & Gopnik, Alison (2017). Explaining constrains causal learning in childhood. *Child Development*, 88, 229–246.
- Wojtowicz, Zachary & DeDeo, Simon (2020). From probability to consilience: How explanatory values implement Bayesian reasoning. *Trends in Cognitive Sciences*, 24, 981–993.