Introduction: Philosophical Foundations

JONATHAN E. ADLER

1. Reasoned Transitions

Reasoning is a transition in thought, where some beliefs (or thoughts) provide the ground or reason for coming to another. From Jim's beliefs that

(1) Either Bill receives an A or a B on the final.

and

(2) Bill does not receive an A.

he infers that

(3) Bill receives a B.

Assuming that Jim bases his inference on the deductive relation of (1) and (2) to (3), his conclusion is warranted, since the argument is valid. (1) and (2) implies (3), since it is not possible, as contradictory, for (1) and (2) to be true and (3) false. More formally, ϕ is a logical consequence of Γ if and only if there is no interpretation (model) in which all sentences of Γ are true but ϕ is false (Tarski 1983).

Although in reaching (3) Jim comes to a new belief, its information is already entailed by (1) and (2). Unlike deduction, an inductively good argument provides for new beliefs whose information is not already entailed by the beliefs from which it is inferred:

(4) Bill brought his back pack to class every day of the semester.

So, [probably] (5) Bill will bring it to the next class.

The falsity of the conclusion (5) is compatible with the truth of the premises (4). The premises only render the truth of the conclusion more probable (than in their absence). Although this is a good inductive argument, the premises can be true and the conclusion false, so the argument is invalid.

Deductive validity is monotonic: A valid argument cannot be converted into an invalid argument by adding additional premises. But an inductively good argument is nonmonotonic: new premises alone can generate an argument that is not good. If I add to the argument from (4) to (5), a premise that

(4.1) Bill's back pack was stolen,

the conclusion no longer follows.

In either argument, there is a reasoned transition in thought. The person who draws the inference, takes the premises as his reasons to believe the conclusion (or, in the second case, to believe it probable.) By contrast, the transition in thought from the belief that

(6) Joe's cousin drives a BMW.

to

(7) I better call Fred.

is not reasoning because, let's suppose, (6) is merely a cue or stimulus or prompt for the thought that (7) to arise. (6) could not serve as the reason for accepting (believing) (7) as true, as (1) and (2) could for (3). (Another technical use of 'accepting' is for momentary purposes, as, say, when one accepts a supposition for a proof Stalnaker 1987).

Grice (2001) draws the connection between reasons and reasoning by noting that if reason is the faculty which "equips us to recognize and operate with *reasons*" then we should also think of it as the faculty which "empowers us to engage in *reasoning*." Elaborating, he writes

> if reasoning should be characterizable as the occurrence or production of a chain of inferences, and if such chains consist in (sequentially) arriving at conclusions which are derivable from some initial set of

> premises . . . . of which, therefore, these premises are . . . reasons, the connection between the two ideas is not accidental. (5)

Grice's 'not accidental' is, presumably, a cautious expression for a conceptual dependence of reasoning on reasons.

Minimally, to have a reason is to have a favorable consideration. However, a reason to do something as in 'my reason to go to the ice cream store is get a sundae' serves to motivate action, whereas a reason to believe does not serve in a motivational role. You can be indifferent to the grade Bill receives, but not, presumably, to the ice cream sundae. Of course, reasons or evidence are typically uncovered through investigation, as when trying to determine the grade Bill receives. But then the motive to investigate obtains independent of the reason to believe. However, in other cases, and much more typically, we acquire evidence that a statement is true and then we come to believe that statement, like it or not. If you overhear Jim affirm (3) that Bill receives a B, then special circumstances aside (e.g., you do not trust Jim), you will come to the corresponding belief, even if you are indifferent to Bill's grade. There is no gap between judging that there are sufficient reasons to believe *p* true and judging (accepting) that *p* is true, nor between judging that *p* is true and believing it.

What is a favorable consideration? Is (1) and (2) a reason to believe (3) because they constitute a mental state or because they constitute facts which serves as the content of that state? If my wanting the sundae is my reason to go to the store, the mental state is the reason. If, instead, what I want to be the case – the fact that I buy apples – is my reason, it supports the truth of the belief's content that I go to the market. ("Belief" suffers a similar ambiguity. Does it refer to an attitude (believing) or to the content of that attitude? We assume that when disambiguation is needed, context will prove adequate.)

Is the reason (as a proposition) a consideration to hold a certain attitude – believing or desiring – or is it a consideration favoring the truth of the content of that attitude (Parfit 2001)? The mother who learns that her son survived a fire in school will be relieved by coming to the belief that he survived, which is then a reason – a consideration in favor – of her taking the attitude of believing it. But that value or utility to her of holding the belief is not a reason that renders it true that her son did survive. In general, it seems that the utility of believing a statement, since it is never a reason for the statement's truth, can never serve as a proper reason to believe. Arguably, though, even if utilities can not bear on what to believe, they may enter with the question of whether to hold a belief rather than not to hold any (Nozick 1993 Ch. III).

2. Belief and Truth

Induction and deduction supply reasons to believe, since each seeks to preserve the truth of its premises, while extending them to new truths acquired as beliefs. Beliefs are the product of reasoning since belief aims at truth. The end result in belief explains why reasoning matters so profoundly to us. We care to get correct conclusions both intrinsically, since that is what having a belief claims, but also, and more obviously, extrinsically. Belief guides actions, and actions are expected to succeed (reach their goal) only if the beliefs that guide them are true. If you want an ice cream sundae immediately and you believe that the only near-by place to purchase one is on the corner of Broadway and 110th St., then you are expected to succeed (to satisfy your desire for the sundae) only if your belief as to its location is true.

Both forms of reasoning or inference aim to discern what is the case, and so aim, figuratively, for the mind to fit the world (e.g., that I come to believe that a sundae is produced at the store just because it is). By contrast, to desire the ice cream sundae, which specifies one's goal in action (to acquire and to eat the ice cream sundae), is to desire the world to conform to the mind. Beliefs and desires have opposite "directions of fit" (Anscombe 1957, Searle 1983).

The fundamental notion of belief is that of "believing that", a characteristic propositional attitude. If Jim believes that Mary is in Alaska, Jim believes the proposition Mary is in Alaska to be true. Propositions are the contents of sentences or statements as expressed on an occasion. The sentence 'I like Krispy Kreme donuts' cannot be true or false as it stands, since the 'I' has no definite reference. But, on an occasion of use, the fixed meaning of 'I' (and similarly for other indexicals like 'you' or 'now') will have their reference determined; the reference of 'I' is the speaker on that occasion (Kaplan 1989). When values for all indexical and similarly context-sensitive terms in an assertion are fixed, the statement expresses an abstract entity of a corresponding form, a proposition. (The prominent features of context are speaker, hearer, location, and time.)

What is it we are claiming of a proposition when we attribute to it truth or falsity? There does not seem to be any difference between

(8) John believes that the proposition that the nearest ice cream store is on Broadway and 110th St. is true.

and

(9) John believes that the nearest ice cream store is on Broadway and 110th St.

Both seem to say the same thing – to be true or false under the same circumstances – suggesting the generalization:

(T) The proposition that **p** is true if and only if **p**.

The left-hand side of (T) ("The proposition that **p** is true") speaks about a proposition. The right-hand side speaks about the world or a fact of the world namely, that the nearest ice cream store is on Broadway and 110th St. The circular appearance does not run deep.

If (T) is correct, there is no further problem about understanding truth than understanding the corresponding proposition. If you understand the proposition that the library is open on Saturday, no special difficulty attends to your understanding the proposition – that the library is open on Saturday – is true (Tarski 1983; Horwich 1990). However, the (T) equivalence does not tell you how to determine or verify or discover whether a proposition is true.

3. Theoretical and Practical Reasoning

Reasoning to how one should act can involve inductive and deductive transitions, but its aim or purpose is distinctive from reasoning whose endpoint is belief:

(10) I want an ice cream sundae.

(11) The closest ice cream store is on Broadway.

(12) There are no barriers to my going there.

So, (13) I should now go to the ice cream store on Broadway.

[Alternatively, (13) I shall/intend to now go . . .]

(10)–(12) constitute good reasons for concluding (believing) that (13) is true. But the ultimate purpose of this reasoning is not to figure out what is the case. Reasoning whose endpoint is belief is referred to as *theoretical reasoning*. Rather, this reasoning (10)–(13) aims to figure out how one should act or *practical reasoning*. The goal is to figure out what one [I] should do (Millgram 2001). As indicated by the alternative reading, (13) should be viewed not just as a judgment as to what is best for me to do, but the actual intention to so act.

Theoretical reasoning aims to answer whether **p** is the case, not whether I ought to believe it, whereas practical reasoning is concerned to determine what I ought to do. The structure of theoretical reasoning is obscured if its conclusions are taken to be of the form 'I ought to believe **p**.' What it is best to do is that act which is better than all the alternatives, on the available reasons. But what one can or should believe is only what is genuinely worthy of belief, not what is currently better than the alternatives. (Think here of the difference between poker, where the best hand wins, and rummy, where only the right or proper hand can win Adler 2002).

The end or goal to which practical reasoning is directed is characteristically set by what one wants or desires, expressed in premise (10). (Not that any desire or want specifies a real end or goal – something that you aim to pursue or that even supplies a reason or motive to pursue. You may have a desire to humiliate yourself, which you neither value nor with which you identify.) Practical reasoning aims at figuring out how to go about satisfying a desire, if opportunity permits. When one's wants or desires set a genuine end or goal, motivation to act according to the conclusion's directive is built in. It is unremarkable self-interest to attempt to satisfy one's own ends.

Can one be motivated to act other than internally (from one's wants or desires)? The Humean 'internalist' answers 'no', whereas the Kantian and other 'externalists' answer 'yes.' (Williams this volume). Externalists hold that one can be motivated purely by recognition of a reason (belief) that a rule or principle or duty applies. So, for example, can a child be motivated to visit his grandmother without any desire to do so nor any threat of punishment? Can his recognition that visiting his grandmother is the right thing to do give him a reason to act accordingly, even if he has no internal – desire – motive to do so? Can reason alone, as a source of judgments of truth and falsity, be a source of reasons (motivation) to act?

With slight differences, the internalist answers "no" to these questions, holding that reason is inert. Reason (belief) is only able to guide one to those actions that are likely to

satisfy the motivation that lies elsewhere (in one's desires or wants).

Are one's wants or desires the endpoint to fix one's goal or aims to which practical reasoning is directed? Internalists typically deny that one's ends (or goals) can be rationally altered except on the basis of further desires or wants. One methodological weakness in this instrumentalism is that one's desires are often too unspecific to fix any end. (Richardson 1994; Millgram this volume). Hardly anything is fixed by one's desire to be happily married (to whom?) or to get a good job (which one?). One's ends must be specified to serve as a guide to action, and the specification requires input from beliefs.

Similarly, one's plans need constant updating and modification as they begin to be executed (Bratman 1987). When you learn of a traffic jam further up on the highway, you turn off to the service road. In this way, you fill in your plans, not just modify them. One's plans direct one toward one's goal, but they do so in an open-ended way, leaving room to fill in details and for modifications, as more information is learned.

In theoretical reasoning, motivation is not an ingredient, which is another way to mark its "inertness." Once you judge a conclusion true, based on the reasoning, you thereby believe accordingly, idiosyncratic psychological barriers aside (e.g., distraction). Belief is in one way passive and not subject to choice: Think of all the beliefs you pick up on the way to your morning commuter train to which you are completely indifferent, for example, that your new neighbor is wearing a green jacket today. No motivation is necessary for belief to respond to a convincing argument. It is a heard contradiction, discussed further below, to affirm a statement of the form "*p* is true, but I don't believe it" ("Moore's Paradox").

The objective of theoretical reasoning is to relieve doubt or to satisfy curiosity or to diminish puzzlement by achieving corresponding beliefs, whereas the objective of practical reasoning is to secure the means to realize one's ends. Because practical reasoning is directed toward action it is overtly constrained by time and resources – its objective is to discover which option is best, all available things considered. But the objective of theoretical reasoning is not merely to discover what proposition (option) is best supported by one's available evidence, but what is correct (true). Consequently, to draw a conclusion in theoretical reasoning requires the claim not just that one's evidence is the total relevant evidence available, but that the evidence is representative, rather than a skewed sample. It follows further that our limits in gathering and assessing evidence contours theoretical reasoning, but, like utilities, our limits cannot play an overt role in drawing a conclusion. You should not – and, perhaps, cannot – believe a conclusion because it is best supported by the evidence so far and that you do not have more time to examine further evidence. (The problems raised here are for assimilating theoretical to practical reasoning. For a discussion of the assimilation of practical to theoretical reasoning, see Velleman 2000.)

4. Theoretical Reasoning: Limits, Closure, and Belief-Revision

Our limits restrict the resources and time to devote to empirical search, testing, and inquiry, as well as to the inferences worth carrying out. The valid and sound argument from "Trump is rich" to "Trump is rich or cousin Harry is in Jamaica" yields no new worthwhile information. Endless such trivial consequences (e.g., *p*, *p* or *q*, *p* or *q* or *r*, . . . ; *p*, *p* & *p*, *p* & *p* & *p* . . .)[1] can be so generated, which will just "clutter" one's memory as explicit beliefs (Harman 1986; Sperber and Wilson 1986). Also, if one "loses" or forgets the origination of the disjunctive belief in the belief that Trump is rich, one will mislead oneself on attending to it that one has special reason to believe Harry is in Jamaica or that it bears a significant connection to the other alternative that Trump is rich.

Theoretical reasoning involves revising beliefs we already hold. Rules of standard logic or implication, however, do not (Harman 1986). Jane believes that if she attends Yale, she'll become an atheist. She believes that she will attend Yale. If she reasons by the impeccable rule of Modus Ponens (MP: *p* and if *p*, *q* implies *q*), she concludes that she will become an atheist. But although she now has a reason to believe that conclusion, logic does not decide that she will or should believe it. Once Jane becomes aware of that conclusion, she also becomes aware of other beliefs, which deny that she will ever be an atheist. Instead of drawing the MP conclusion, Jane ceases to believe the conditional, which served as her main premise. Reasoning that results in modification of beliefs of one's own may be dubbed "self-reductios" (ad absurdum).

Examples such as the previous one show how from deduction we can learn something new about the content of our beliefs, even though, in a figurative way of speaking, deduction only

renders explicit information already in the premises. Briefly, our beliefs are not closed under deduction. (Similar, but distinguishable, worries attend to the requirement that our beliefs at a given time be consistent. The worries are different for failures of closure or consistency that the agent does not recognize and those cases in which the agent does recognize the failure. The latter cases generate more forceful conceptual friction with the concept of belief. For recent treatment of these logical requirements on belief, see Christensen 2004.) One's beliefs are closed just in case if one believes *p* and *p* implies *q*, one believes *q*. None of us mortals have bodies of beliefs that are deductively, let alone inductively, closed. There are complex tautologies or logical equivalents to what we believe, which we will not believe and may even disbelieve. The failure of deductive closure for belief is a facet again of our limits, including our limited grasp of our own beliefs, our lack of omniscience, and our "inability" to perceive the future. If Socrates believes that no one does wrong knowingly, does Socrates believe that Richard Nixon did no wrong knowingly? (For examples and critical reflection's see Stalnaker 1987: Ch. 5.) Implications or deductions from one's beliefs can yield surprising conclusions. Well before the discovery of penicillin by Fleming, biologists knew that molds cause clear spots in bacteria cultures, and they knew that a clear spot indicates no bacterial growth. Yet, they did not come to realize, or even to hypothesize, that molds release an antibacterial agent. The observations did not render salient the disparate beliefs and place focus on them together (Cherniak 1986).

So, putting aside closure imposed as an idealization for specialized purposes (Hintikka 1962), one can believe, or even know, *p*, and *p* imply *q*, without believing or knowing *q*. Similarly, it can be the case that a = b and that one believes (and even knows) that a is F, without one believing that b is F, though the embedded argument is valid. So, for example, Lois Lane may know that

(14) Superman flies.

In the tale, it is true that

(15) Superman is Clark Kent.

But Lois Lane does not know (and actually believes false) that

(16) Clark Kent flies.

The fault lies with a lack of knowledge of the middle step – (15). Knowledge or belief is "opaque" – in "S believes that *a* is F" the position of "a" is not purely referential. (Opacity intrudes on what counts as a reason: If Lois Lane wants to marry only a man who flies, does she have reason to marry Clark Kent?) Consequently, substitution of arbitrary coreferential terms is not truth-preserving. Within the scope of Lois Lane's beliefs or knowledge, "Superman" in (14) does not simply refer to an object (Superman), but to that object as understood by Lois Lane. (The problem originates with Frege [1970]. An alternative account holds that the substitution does go through. The assumption as to how the person [Lois Lane] thinks of the name is only pragmatic. A parallel worry applies to the distinction between attributive and referential meanings of a term [Kripke 1977]. To appreciate this alternative reading substitute for the names a pure pointing device like "this" or "that.")

A much-discussed example takes the problem of closure a step further, because it holds that knowledge is not closed even when the person knows the "middle" step – the relevant implication (Dretske 1970; Nozick 1981). Assume that Tony is looking at an animal behind a cage marked "zebra," which looks like a zebra. Barring any weird circumstances, we would say that Tony knows that

(14) The animal I am looking at is a zebra.

Let's now grant that Tony also knows the implication that

(15) If the animal I am looking at is a zebra, then it is not a mule cleverly disguised to look like a zebra.

(14) and (15) imply

(16) The animal I am looking at is not a mule cleverly disguised to look like a zebra.

Still, we are reluctant to attribute knowledge of (16) to Tony. Those who oppose closure reason that Tony has never checked that the animal he is looking at is not such a cleverly disguised mule. Tony is simply looking at the animal from outside the cage. These theorists reject:

> (Epistemic Closure EC) If X knows that *p* and X knows that *p* implies *q*, then X knows that *q*.

Arguably, this is the principle licensing Descartes' famed sceptical argument: If you know that you are in your office and you know

that if you are in your office, you are not just *dreaming* it, then you know you are not just dreaming it. But you do not know that you are not dreaming it. So you do not know you are in your office.

The rejection of EC fits the previous zebra example, answers Descartes' sceptical argument, and it is explained as due to our not checking on all the implications of propositions that we know. The rejection also follows from analyzing knowledge as involving satisfaction of the following subjunctive or counterfactual conditional:

> (Tracking Knowledge TK) Were *p* false, S would not believe *p*.

The most likely (or nearest) way for it to be false that you are in your office is for you to be somewhere else, like your kitchen. If so, you would clearly recognize where you were in the other room. Consequently, you would satisfy (TK), because you would not believe that you were in your office. (TK) does not then support (EC). (Contextualists, whose views we return to below, hold that you do know that you are not dreaming, when you are in an ordinary setting. However, when Descartes or a skeptic mentions the possibility that you are dreaming, they alter the context or standards for knowing. Only then you do not know that you are not dreaming. But, in that case and compatible with EC, you do not know that you are in your office either.)

Despite these advantages, the dominant view is that EC cannot be rejected, since deductive implication preserves truth. What better way to know the truth of a proposition but by deducing it from a proposition one does know? (For overview, see Luper 2006). Without pursuing this line, it's worth noting that sometimes (non-trivial) deductions seem not be a way to advance knowledge. From the evidence of (17),

(17) The Smiths are making an extravagant wedding for their daughter.,

(18) is concluded:

(18) The Smiths are wealthy.

From (18), (19) follows:

(19) In making the extravagant wedding, the Smiths are not just appearing to be wealthy.

Assume that you are in a discussion with someone who disputes whether the Smiths are really wealthy. Although (18) implies (19), it seems to *beg the question* in this context to use (18) as a reason to believe (19). (17) can only provide evidence for (18) if (19) is assumed or presupposed. But (19) is in dispute. If it is presupposed in treating (17) as evidence for (18), then the warrant or support that (17) lends to (18) does not *transmit* to the conclusion (19) (Wright 2000).

5. Belief-Revision, Holism, and the Quine–Duhem Thesis

If a corpus of beliefs is not closed for the reasons suggested, it is likely to be inconsistent. If there are serious implications of one's beliefs that one fails to believe, one is likely to acquire the contrary of some of those beliefs without recognizing the incompatibility. Here's a very ordinary illustration of Lewis's:

> I used to think that Nassau Street ran roughly east-west; that the railroad nearby ran roughly north-south; and that the two were roughly parallel. (1982, 436)

Once these beliefs are brought together with the evident tacit belief, Lewis recognizes that the set of beliefs {Nassau Street ran roughly east-west; the railroad nearby ran roughly north-south; Nassau Street and the railroad nearby are roughly parallel; if one path is east-west and another is north-south, they are not parallel} is inconsistent. They cannot be simultaneously true. Once Lewis recognizes the inconsistency, he can no longer hold on to all these beliefs ("I used to think . . . "). The question that he now confronts – the question of belief-revision – is how he should restore consistency.

Rejecting any one or more of the members of the inconsistent set will restore consistency. Quine (1980) argued that selection to restore consistency depends on extralogical considerations. He made these claims in developing his criticism of the "dogma of empiricism" that statements (hypotheses) can be tested in isolation. Instead, he put forth what is referred to as the "Quine–Duhem Thesis" (Duhem 1954) that hypotheses are never tested in isolation. Hypotheses (or theories) do not entail any observational predictions by themselves. To derive predictions that serve to test a hypothesis, assumptions are required that crucial terms are not empty and that conditions are normal. Newton's enormously successful theory of gravity and mechanics erred in its (pre-1846) prediction of the orbit of Uranus. But Uranus's deviation was treated as an anomaly, rather than

a falsification, because the theory made substantial assumptions about the operative gravitational forces. In the discovery of Neptune, some of those assumptions were abandoned, rather than the Newtonian theory itself. In general, when a well-regarded hypothesis fails, we do not immediately conclude that the hypothesis is false, as the traditional view implies, rather than that some of the conditions assumed normal – the auxilliary assumptions – failed.

The hypothetico-deductive model incorporating the Quine–Duhem thesis is represented schematically as

H and Auxilliary Assumptions (AA) imply O.

If O, then (H and AA) are confirmed.

If not O, then (H or AA) fails.

The latter is the crucial result because consistency does not demand the falsity of H.

The problem of what extralogical principles to apply to belief-revision has generated numerous investigations and constructions of logics of belief revision (Hanson 2006; for an introduction to a computational approach to belief reasoning and revision, see Pollock and Cruz 1999: Ch. 7). A central proposal is that belief revision should be conservative. One revises one's beliefs so that rejection or modification is minimal. You all-out believe that Skinner wrote *Walden II* and that Chaucer wrote the *Canterbury Tales*. But if you discovered that one of these is wrong, you would sooner surrender one rather than both and the latter, rather than the former, which is attested to by a greater variety of good sources. To surrender the belief about Skinner's authorship would require surrendering – nonconservatively – much more information than surrendering the latter.

But conservatism cannot stand alone. Lewis cannot just decide to give up the belief that Nassau Street ran roughly east-west and keep the one that the railroad nearby ran roughly north-south, although that surrenders only one, rather than more, among equally contentful, incompatible, beliefs. Merely his deciding would not be a sufficient reason that the belief retained is true.

Other principles of belief revision include prominently simplicity and coherence. The more a belief coheres, fits, or is explanatorily connected, with others, the more resistant it should be to modification. But some conflictual beliefs may be surrendered (in strain with conservatism) to increase coherence. You believe that ten-year-old Jim is in good health, that he will be in the tennis tournament tomorrow, and that he will meet you for lunch at noon. You learn that he is not at school today and that his mother is not at her office. You infer that Jim is sick. That best explains the latter two beliefs – unifies them in an explanatory nexus. But as a consequence you surrender other beliefs about Jim (e.g., that he will be at tennis practice later in the afternoon).

Coherence is an internal requirement on one's beliefs. But our belief corpus improves by external input, especially through our senses. The improvement is not just in the addition of new beliefs picked up as we navigate our environment. Perception and other sensory mechanisms provide for ongoing self-correction, which is a hallmark of scientific method (Sellars 1963). If you believe that Lisa is in Alaska and you see her car at the local diner, you surrender – and so correct – your belief. Normally, beliefs operate as a filter on perceptual judgment. If Lisa drives a blue Ford and that was all that was in your perceptual field as you approached the diner, then if you noticed the car, you would not think of it as hers, given that you believe that she is in Alaska.

However, there must be limits to this filtering role, otherwise our beliefs would not be subject to correction. Once you see the car more closely, and observe the familiar dent on the hood, you are compelled to notice that it is Lisa's blue Ford. Your belief is revised. The new perceptual information nullifies your prior belief, evidence that perception can succeed as a self-corrective on reasoning only if it has some independence of operation from belief and reasoning (from one's "central systems," Fodor 1983).

If the formation of perceptual beliefs always had to be first checked (for veracity) by way of one's corpus of beliefs, it would be subject to the "dogmatism paradox." If you know that Lisa is in Alaska, why should you even acknowledge as putative undermining evidence that it is her car at the Brooklyn diner? Shouldn't you rather judge that, say, her husband must have taken her car, because, if you know that she is in Alaska, shouldn't you know that putative evidence against it must be mistaken or misleading? [Think of the tautology: p–>((q–>~p) – >~q).]

However, there is something deviant about the conditional "If there is evidence against my knowledge (that Lisa is in Alaska), then that evidence is mistaken or misleading." It does not

seem to be open to modus ponens, just as the following is not:

> If my wife cheated on me, I would never know. (Harman 1973; Ginet 1980; Stalnaker 1987)

Were I to discover that my wife cheated on me, I would reject the conditional (or its antecedent), rather than conclude that I would never know that she cheated on me. Similarly, the previous conditional of the dogmatism paradox is to be rejected when the undermining evidence is obtained, rather than rejecting the evidence as misleading or false.

6. Deductive Rules and Deviant Logics

Usually, of course, conclusions drawn from one's beliefs simply form new beliefs. But often conclusions are drawn that cast doubt back on the premises (beliefs) or the inferential transitions from which those conclusions are drawn.

In drawing out conclusions, what rules should be used? Although in the case of induction and especially deduction a core of rules and results are well established, disputes abound about their scope and other putative rules are flat-out contested. After considering some of these rules and disputes, we briefly turn to how the rules are to be justified and selected.

Standard or classical logic is first-order quantification or predicate logic – the logic of the truth-functions ("and," "or," "not," "if, then") and the quantifiers ("For every . . . ," "For at least one . . . ") The "first-order" implies that the variables of quantification take as values objects or individuals, not names, predicates, propositions, or properties. First-order logic (including identity) is sound: no proof (a syntactic notion) will take one from truths to falsehoods. Every proof corresponds to a valid argument, a semantic notion. But also and more distinctively, first-order logic is complete (every logical truth or valid argument is provable). Once second-order quantification is admitted, particularly to embrace set theory, the extended logic is no longer complete.

Unlike additions to standard logic, as in adding axioms for necessity and possibility to logic, deviant logics deny some basic logical law. Quine's "holism" opened a door to defend deviation from classical logic, which Quine (1970) attempted to quickly shut. The holistic assumption that justification for logical laws, like the justification of empirical claims, is sensitive to the whole body of beliefs provides an opening for arguing that a logical law is to be rejected because removing it from one's corpus of beliefs increases coherence. But although Quine's opened this door as a theoretical possibility, he argued that the standard logical laws are too useful or indispensable to the progress of science for abandonment.

Additionally, Quine argued that when you try to deny a logical law like the law of noncontradiction (i.e., $\sim(p \ \& \sim p)$, your "&" and "$\sim$" no longer translate "and" and "not," as intended. These operators are fully specified by the truth-tables and the implied laws. They are defined implicitly by their roles in these assignments and laws. (However, we know that unless restraints are imposed on implicit definitions, specifically, that they introduce no new theorems ["conservative"], the introduction of new connectives can generate crazy rules, ones from which anything can be deduced; Prior 1960; Belnap 1962). If someone infers from an utterance of "A or B" to "A," we can be sure his "A or B" is not the English disjunction. The deviant logician's predicament is that "when he tries to deny the doctrine he only changes the subject" (Quine 1970), p. 81.

Quine's argument opposes the plausible claim that the denial of an operator in some inferential roles is compatible with its playing the appropriate roles in other inferences, sufficiently so as to remain a viable candidate for capturing the basic meaning. Even the denial that all contradictions are false, allows for preserving a good deal of classical logic with suitable adjustments. In classical logic, every sentence follows from a contradiction. However, this trivialization can be excised by denying the rules from p to $p \vee q$ (weakening); and from $p \vee q$ and $\sim p$ to q (disjunctive syllogism)). In the former case, the denial is independently motivated by the lack of relevance in subject matter of the premise to the disjunct q. With a number of other adjustments a major portion of basic logic remains in tact, and the resulting logical system can be sound and complete (Priest 1998).

Although most of the alleged examples of contradictions that are candidates for the ascription of truth are rarified, one of them is a resource from which a number of deviant logics draw strength, and it resonates with everyday reasoning. The resource is in the phenomena of vagueness. The vagueness of a term is that while it sorts objects into those to which the term applies and those to which it does not, it leaves undecided many objects. When a teenager's room has no clothes on the floor and dirty dishes

have been removed, but it has not been dusted or swept, it is indeterminate, let's suppose, whether it is clean or not. A defender of the view that there can be true contradictions might say that the room is both clean and not clean. Others may deny the law of excluded middle: that either the room is clean or it is not clean.

A contextualist confronted by an assertion such as "John's room is clean" will respond that for everyday purposes, it is enough that his clothes are off the floor and that the dirty dishes have been removed. But if John develops asthma, you will not count the room as clean until a careful dusting is complete. The proposition expressed by the assertion is false in that context. Outside of any contextual specification, there is no assigning it a truth-value (or assigning it additional truth-values) (Lewis 1983). Contextualism explains how one utterance of "John's room is clean" can be true and another false, when there is no change in John's room. Because contextual variations include variations in the importance of the matter, contextualism also makes sense of why when you inquire as to the truth of a hypothesis, you are bound to investigate harder in one context (where the costs of error are greater) than in another.

Contextualism, however, is a minority view in how to handle the sceptical implications of vagueness that originates from the "sorites paradox." In its historical form, the sorites is presented as the "paradox of the heap" (Sainsbury 1988; this volume). If you have a heap of sand, and you subtract one grain, you still have a heap (one grain cannot make a significant difference). The judgment suggests a principle: if you now subtract one grain from the previous heap, you still have a heap. But repeated applications leaves you with a couple of grains that you are committed to taking to be a heap, even though they obviously do not make a heap.

Intuitionists respond that when there are too few grains in a pile to clearly be a heap and too many grains to clearly not be a heap, it is not true that either that pile is a heap or that it is not a heap, contrary to the law of excluded middle. And if it not clearly either, the failure to not be a heap does not imply that it is a heap, contrary to the law of double-negation.

The sorites is derivable via the impeccable principle of mathematical induction: R_o is the base case with a certain property P (e.g., That large collection of grains of sand is a heap). And if R_n has P, there is some fraction of it (e.g., one grain of sand) such that if we decrease R_n by that amount, then what results – Rn-1 – still is P (e.g., that 1-less grain collection is a heap). Then any lesser R_i has P (e.g., every collection of grains of sand less than the base case is a heap including a one-grain collection). Alternatively, the paradox can be presented simply as a string of MP arguments, each one yielding a decrement from the previous. Either way, the nonnoticeable difference for any sized heap – the decrement between R_n and R_{n-1} – becomes a marked difference after enough applications.

The sorites paradox is particularly wrenching because it seems to arise merely from the vagueness of terms, which holds of most terms. Most, if not all (nonartificial) terms leave undecided an unlimited range of cases for example, "blue," "happy," "short," "table," "flat," "rich," "child." At what moment does childhood end? The exception would be contrived cases where an exact specification is provided: We could define a U.S. adult citizen as rich* just in case his total wealth is $484,234.04 or higher. However, to attempt exact replacements for our vague terms would fail to preserve their value or usefulness. (On utility considerations for vagueness, see Parikh 1994.) The contrived precision would require a sharp break in judgments, where a gradation of responses is appropriate (between, e.g., a person who is rich and one who is very well off financially). The vagueness of a term reflects its "tolerance" for certain tiny alterations, which can mount up to significant alterations.

But is this insensitivity in the term or is it merely because of the limits of our discriminatory powers, which are foisted on to the term? "Epistemicists" favor the latter, which allows them to avoid offending against standard logic. They hold that there is an exact boundary for vague terms, but it is unknowable (Williamson 1994; Sorensen 2001). (The previous question suggests another: Is it reality – heaps themselves – that is vague or how we describe it, e.g., some collections of grains of sand are described as "heaps"?)

Among the numerous attempts to solve the problem the dominant view preserves almost all of standard logic, allowing only for truth-value gaps. A "supervaluationist" observes that within the indeterminate cases, we are free, as far as a consistent assignment of truth-values, to decide them as serves our purposes, as contextualists claim, too (Fine 1975). Some will treat a U.S. citizen with total wealth of $325,683.03 as rich and others not (for purposes of assigning, say, an estate tax), because this amount is clearly between the definitely not rich and the definitely rich. When we so decide cases we

provide "sharpenings." However, on every sharpening "That citizen is either rich or not rich" will be true, so supervaluationists can accept the logical law of excluded middle (although neither disjunct may be definitely true). Consequently, on a supervaluationist view, a conditional sentence – the second step in the sorites paradox (e.g., if that citizen is rich with $325,683.03 then he is rich with $325,683.02) will not come out true for every sharpening of it. There is a sharpening under which the antecedent is true, but the consequent false, so that the conditional is false. Because supervaluationism does not reject logical laws and it does not require powers in our language that supercede our own, it has the advantage of providing for a conservative response to the sorites.

But is the supervaluationist right, to return to the earlier example, that '"either that room is clean or not" is definitely true, when the room is clearly a borderline case? Worries like this incline others, although far fewer, to take the route of treating the initial reaction that neither alternative is true at face value. We can say only that John's room is clean to a certain degree, or to a higher degree than others. The error on this probabilistic approach is to contrive to derive absolute judgments from matters of degree.

7. Ordinary Language Challenges to Logic and the Conversationalist Response

A very different source of doubts about logical reasoning as standard first-order logic derives from alleged deviations from ordinary language. Numerous patterns of inference of ordinary language, as well as straightforward readings of complex statements, prima facie do not obey the rules governing deductive logic or the logical operators. Some examples:

(20) John goes drinking and John gets arrested.

(21) John gets arrested and John goes drinking.

If "and" is the "&" of formal logic, it is symmetric, so (20) and (21) should be equivalent. Yet, they do not seem to mean the same, (21) does not follow from (20) (or conversely). Another:

(22) John will order either pasta or steak, but he orders pasta.

So (23) John does not order steak.

The inference seems valid, but fails on the truth-table analysis of "v" (inclusive "or"). Finally for an example using a conditional to which we devote the next section:

(24) If you tutor me in logic, I'll pay you $50.

So (25) if you don't tutor me, I won't.

The conclusion seems to follow. However, the straightforward translation of it into logic yields a fallacious form, one that appears valid, but that isn't.

One reaction to such discrepancies is: so much the worse for ordinary language. It requires formal regimentation to be a satisfactory medium of reasoned argument. The opposed reaction is: so much the worse for logic's claim to provide a systematic analysis of ordinary reasoning.

The most profound reaction is that of H. P. Grice's (1989). His account of conversational reasoning opposes both previous ones. Grice's claim is that the logic of ordinary language is already that of formal logic. However, we impose, without recognition of the imposition, assumptions or expectations on ordinary language because we treat the sentences as assertions or other contributions to a conversation. In (22)–(23), we assume that John could not order both pasta and steak, given our knowledge about eating. The inference from (20) to (21) (or their equivalents) stands. However, the speaker exploits the listing of conjuncts, which is mutual with the hearer, as implicating an ordering (in time). Grice's account explains these deviations without positing an ambiguity in the relevant logical constant (e.g., "and," "and then").

Conversation or social communicational exchanges are facilitated by shared or mutual assumptions that are not part of what is said or its logical implications, but which nevertheless are invited, given the common goals of the cooperative exchange. The fundamental "maxim" (and so expectation or presumption) is that the speaker intends to cooperate to advance the purposes of the conversational exchange. The Cooperative Principle (cp) includes subsidiary maxims. The speaker intends his contribution to be informative, warranted, relevant, and well formed (for brevity, style, politeness, and comprehension). This package of maxims under the cp Grice thought to be justified as principles for rational cooperative arrangements for beneficial ends (of transferring information).

What we mean to communicate typically goes beyond what is said, although calculated on the basis of what is said: