

Reasoning with Limited Resources and Assigning Probabilities to Arithmetical Statements

Haim Gaifman

There are three sections in this paper. The first is a philosophical discussion of the general problem of reasoning under limited deductive capacity. The second sketches a rigorous way of assigning probabilities to statements in pure arithmetic; motivated by the preceding discussion, it can nonetheless be read separately. The third is philosophical discussion that highlights the shifting contextual character of subjective probabilities and beliefs.

1. Limited Deductive Ability

Two kinds of obstacles stand in the way of judging and decision-making: lack of empirical knowledge and limited deductive ability. The first gives rise to disciplines that focus on various aspects of reasoning under uncertainty, such as subjective probability, belief revision and rational decision theory. The second underlies the research in computational complexity and related areas, such as cryptography; these are thriving research directions, yet they encompass an area considerably smaller than the disciplines of the first kind. This is particularly true when we consider the more philosophical type of investigation. The subjects that have to do with reasoning under uncertainty have given rise to philosophical research directions that have produced a voluminous literature. When we come to the second kind of epistemic obstacles, the harvest is very meager.

The effect of limited deductive ability cannot, of course, be ignored. In trying to give a systematic account of belief, one is confronted with the so-called logical omniscience problem. In the customary formal models, belief systems are construed as deductively closed theories; but actual agents are not logically omniscient. One might not believe something that is logically implied by one's beliefs, because one might fail to see the implication. An agent might even have beliefs that logically contradict each other. Mersenne had believed that $2^{67}-1$ is a prime number, which was proved false in 1903, cf. Bell [1951].¹ Together with Mersenne's other beliefs about multiplication and primality, that belief logically implies that $0=1$. Some philosophers have been willing to bite the bullet and ascribe that belief to Mersenne (a belief he was completely unaware of). But this evasive move, which consists in redefining 'beliefs' so as to impose by stipulation closure under logical implication, is of little value. Designed to save a philosophical conception by redefining the subject, it gives up on the problem of providing an account of beliefs as they figure in actual reasoning, beliefs we are aware of.

There is a deep reason for the difference between the handling of the two kinds of epistemic obstruction (the empirical and the deductive); it stems from the status of mathematics in human cognition. Mathematics, at least a certain elementary part of arithmetical reasoning, constitutes a hard core of coherent thought. We can imagine

¹ The factorization, discovered by Cole, is: $193.707.721 \times 761.838.257.287$.

empirical facts being other than what they are. Gore could have won the 2000 presidential elections; the coin could have landed on the other side. Even if some sort of non-logical necessity attaches to certain empirical outcomes, we can still conceive the possibility of their not taking place. But a state of affairs in which 67 is not prime, or in which $2^{67}-1$ is, is inconceivable. Standard constructions that are used in the analysis of uncertain knowledge—possible worlds (for epistemic possibility) or Boolean algebras (for subjective probability, or credence)—are useless when it comes to arithmetical beliefs. There is no possible world in which $2^{67}-1$ is a prime number; the event (in the Boolean space) that $2^{67}-1$ is prime is the 0-element. It is possible to have false arithmetical beliefs, since we may fail to comprehend a sufficiently long sequence of deductive steps; but the false belief cannot be realized in some possible world, for such a possible world would be conceptually incoherent.

Philosophers have therefore tended to regard the limitation of human deductive power as a noise factor that spoils the rational-agent picture of humans. Ignoring it, they have tended to take as a basis for philosophizing an admittedly idealized picture. Now there is no shortage of deductive errors and of false mathematical beliefs. Mersenne's is one of the most known in a rich history of mathematical errors, involving very prominent figures; cf. [De Millo, Lipton, Perlis, 1979], pp. 269-270. The explosion in the number of mathematical publications and research reports has been accompanied by a similar explosion in erroneous claims; on the whole, errors are noted by small groups of experts in the area, and many go unheeded. There is nothing philosophically interesting that can be said about such failures. I shall nonetheless argue that, far from a mere nuisance factor—an unfortunate frailty that can be ignored in the larger picture—the limitation of deductive capacity is constitutive to human cognition. Remove it and the human being is possessed with an infinite mind, which is able to see all kinds of connections within the corpus of physical data, a mind that can survey all possible patterns and all sorts of theories that can account for the phenomena. This is a creature whose conceptual space, whose goals, beliefs and desires we cannot begin to imagine. The limitation affects all domains, not only by restricting the extent of true belief, but also by determining what believing is.

Isaac Levi [1991], [1995], is one of the few philosophers who, taking deductive limitations seriously, has tried to incorporate this factor within a general philosophical account. The idealized agent appears in that account as a never-to-be-fulfilled goal. We, *qua* rational agents, are committed to the deductive consequences of our beliefs, though in practice we can never fully discharge the commitment in full. In Levi [1995] these limitations are described as an epistemic malady. The commitment imposes on us a duty to engage in therapy, which consists in the training of our logical abilities and in developing artificial aids (prosthetic devices, in Levi's terminology), from paper and pencil to computers. The commitment does *not* extend to the pursuit of non-mathematical truths in the empirical sciences,² a pursuit he characterizes as 'inquiry'. Following

² Levi speaks about logical, rather than mathematical, truth. My use of 'mathematical' is limited in this context to very elementary claims, e.g., that some given number is prime, or is perfect, or to claims in the area of finitary combinatorics that are derivable using the machinery of primitive recursive arithmetic. I do not intend to raise at this point issues relating to broader mathematical domains. I prefer the term

Peirce's terminology, he uses 'explicative' to characterize changes in belief that result from application of the deductive machinery, and 'ampliative'—for changes that accrue from an increase in empirical knowledge. Explicative changes take place in the course of fulfilling the standing commitment. But an Ampliative change represents a change of commitment, since we are now committed to a different theory. Levi thus ties the distinction to a characterization of rationality, and he faults Frege for failing to make it. Indeed, Frege saw logical inquiry as part of the general project of finding out the truth. Logical truth (which Frege thought includes arithmetical truth) is distinguished by its generality, but the project of discovering the truth is one.

I think Frege was right. Mathematical inquiry, including, for example, attempts to check the primality of particular numbers, is part of the general inquiry into truth. Rationality, let us agree, implies that if I know that something is deductively implied by what I believe, then I should believe it, or revise my other beliefs. The possibility of revising one's beliefs must be included in order to avoid the absurd outcome of being committed to believing that $0=1$ (the case when somebody's set of beliefs is inconsistent). As mentioned above, unrecognized contradictions in one's beliefs are not uncommon and it would be highly ad-hoc to deny in such cases that the agent believes what he sincerely professes. But this does not mean that a special commitment attaches to finding the deductive consequences of my beliefs, over and above the commitment to finding the truth. While super computers enhance our deductive power, this by itself does not make the building of super computers preferable to the building of super accelerators that enhance our knowledge in physics.

Now the fact that mathematical activity is truth-revealing should not blind us to the a priori nature, or to the necessity, of mathematical truth; to its being constitutive of our conceptual organization, without which thought loses its coherence. Wittgenstein thought it impossible to reconcile this aspect of mathematics with the status of facts. He thus proposed to view mathematical activity as some sort of legislation, rather than discovery. Wittgenstein's view does full justice to the role of mathematics as a generator of meaning, to its being a precondition for meaningful discourse. But it fails altogether as an account of mathematical discovery. This issue however is not this paper's concern; it is irrelevant to the present discussion, since Levi does not dispute the factual status of mathematical truths.

But if mathematics is in the business of discovering truths, then it plays in the same field with other scientific inquiries. The fact that an investigation is likely to extend our computational/deductive reach, and that rationality demands that we accept the deductive consequences of our beliefs (or revise them), does not by itself imply any special commitment to that investigation. By its very nature, no investigation comes free; every investigation is a resource consuming activity, where resources are at a premium. I do not mean lab equipment and working teams—items that figure largely in scientific projects—but the much more basic fact that pure thought takes time and effort. If the alleged special commitment that attaches to some type of research is not to be empty, then it can only mean that the research merits special considerations in allocating resources. Now

'mathematical' to 'logical' because it fits better my general views about the nature of mathematics and logic. In the present discussion nothing substantial hinges on this.

resource allocation is, in general, subject to diverse interrelated criteria. For the sake of the argument, let us put aside social, economic and political factors and consider only the advancement of understanding and knowledge. Then a research project can have special merit if it promises to throw new light on a whole area, or to lead to further developments, or to answer some central questions, and so on. All or any of these might count in favor of enhancing computational power, but they might equally count in favor of other truth-seeking activities. Increasing deductive power does not *per se* constitute a special reason.

Levi recognizes the inevitability of our deductive limitations, but he regards it as a doxastic defect, a sort of malady that puts us in constant need of remedies. By contrast, ignorance in empirical matters “reflects no such defect in our doxastic condition. Indeed, it is often a mark of a healthy mind to be able to acknowledge our ignorance”.³ I think that this picture is wrong. It is as if our inability to run at a speed of 40 mph (which some animals can achieve) were described as a defect in our physical condition. More importantly, just as we acknowledge our ignorance in physics, we can and should acknowledge our shortcomings in carrying deductions; this no less marks a healthy mind. The second acknowledgement is achieved on a meta-level, through reflection on our deductive practices; but the fact that it is an inference about inferences should not diminish its force. Moves from one level to another are quite common both in everyday and in formal reasoning. Levi’s position fails to make place for serious meta-evaluation that guides action. One can decide, for example, that carrying out a full deduction is too costly and opt for a policy that runs the risk of deductive error; a policy, that is, that may yield answers that would have been refuted, had we carried a costly deduction. One can moreover justify the policy through estimates showing that such errors are unlikely. A striking example is provided by the proliferation of probabilistic algorithms, which yield very reliable, but not fully proven answers, to purely mathematical questions. This, among other things, will be discussed in the next section.

Given the difference between empirical and purely deductive knowledge, a philosopher might overlook the extent to which the two are integrated. Mathematics is so intimately and so intricately related to various sciences, that it is all but impossible to factor the difficulty in these sciences into a mathematical and a purely empirical one. One may not realize the extent to which the challenge of developing a successful theory in physics is mathematical.

There is also a striking similarity between the effect of false beliefs in empirical matters and in arithmetic. Since in elementary logic any statement whatsoever follows from a contradiction, it might appear that Mersenne’s false belief about the primality of $2^{67} - 1$ (which together with other beliefs he held implies a contradiction) should infect his belief system on the whole. There is, of course, no such effect, even on arithmetical beliefs, and even on other beliefs concerning prime numbers. The reason is that the same deductive shortsightedness that makes the error possible isolates it to a large extent from the main body of other beliefs. The holding of it does not prevent one from developing other sound arithmetical beliefs about prime numbers. On the mathematical background of Mersenne’s time, believing that $2^{67} - 1$ is prime was like having a false belief about

³ Levi [1995] page 41.

some far away island. This is not due merely to the number's size, but to the scarcity of known deductive chains between that belief and others. When the known territories of the deductive web expand, the contradiction between that belief and others emerges; at that stage revision takes place. The process is not unlike the progression that accompanies an empirical investigation. This, let me emphasize again, is not to ignore the difference between the two. In empirical matters there is an independence (or at least an apparent independence), which makes it possible to hold a false belief about one item and a true belief about another, without encountering a contradiction. There is no such independence when it comes to arithmetic. But the length and complexity of the deductive chains create situations that mimic independence. In the next section I propose a way of modeling this mimicking, by construing it as a local phenomena. Finally, in the last section, I shall argue that the restriction on resources puts into question customary ways of ascribing beliefs to persons and calls for rethinking the very notion of one's belief system.

2. Assigning Probabilities to Arithmetical Statements

Sketch of the General Framework

The assignment of probabilities is the most common way of measuring uncertainty about empirical truth. In a probabilistic framework, the probabilities are assigned to members of some Boolean algebra—'events' in probabilistic parlance—that can be conceived as possible states of affairs, or sets of possible worlds, or propositions. The probability of A , $P(A)$, is a numeric value between 0 and 1. In the subjective interpretation, $P(A)$ represents one's degree of belief that A is the case. In an objective interpretation $P(A)$ is some physical parameter associated with A ; either (asymptotic) relative frequency of events of the type of A , or a propensity to produce such a relative frequency. If we think of events as propositions, the Boolean operations become those of sentential logic and the probability assignment is subject to well-known axioms. The story is standard and I shall skip further details.

The issue of assigning probabilities to arithmetical statements was broached in a groundbreaking paper by Hacking [1967]. For such an assignment the most suitable probability bearers are sentences. A logically omniscient agent will assign the same probability to logically equivalent sentences, or, more generally, to sentences whose equivalence is logically implied by the accepted axioms. In that case, the assignment reduces to an assignment defined over propositions, where logically equivalent sentences determine the same proposition. But since the whole point of the exercise is to handle lack of logical omniscience, we should take the probability bearers as sentences, or "what sentences express", something like Fregean thoughts (senses of sentences), which are not preserved under logical equivalence. Given the complexities and problems of Fregean senses, sentences are our best candidates. For the sake of unburdened notation, I shall use the same symbols for the language and for metalanguage and let context indicate the intended meaning. I take it for granted that the sentences that are subject to probabilistic assignment, in English or in some formal system, have translations in the

agent's language; the agent's belief in the truth or falsity of an English or a formalized sentence is construed as belief in the truth or falsity of the its translation.

Sentences *per se* do not form a Boolean algebra. We get the algebra only if we have some equivalence relation; the elements of the algebra are the equivalence classes, or entities that represent them. Thus, a straightforward application of the apparatus to sentences requires an equivalence relation, and a stipulation that equivalent sentences get the same probability. Such was the approach adopted in [Gaifman 1964], with logical equivalence serving as that relation. This leads of course to logical omniscience: Provable sentences get probability 1, refutable ones—probability 0; there is no place for expressing uncertainty by assigning an intermediate value. Primality claims, such as the claim believed by Mersenne, come under the sway of this zero-one law, because such claims are either provable or refutable: t is prime (where ' t ' is a numerical term) iff it does not have smaller proper divisors; the proof, or refutation, of ' t is prime' is obtained by a straightforward checking of all the smaller numbers (using some standard division algorithm); the computation that gives the answer can be translated into a proof of the claim, or of its negation, in the formal language of arithmetic.

The present proposal is to represent the agent's perspective by some finite, large enough set, \mathcal{P} , of sentences. These sentences are available to the agent in constructing proofs. Though finite, the set can be very large, so as to encompass the reasoning (including use of computers) that is applied in solving a given class of problems. The probability function is partial. In one version it is defined over \mathcal{P} ; in another, it is defined over the closure of \mathcal{P} under Boolean connectives. The probability axioms are supposed to hold, as long as we are within the probability's domain of definition.

The treatment of a single probability function generalizes naturally to cases where the agent's point of view is represented by a family of functions, or to cases of interval valued functions, where one's probabilities are governed by inequalities that need not determine a unique value. I shall not pause on these further extensions, whose relations to the standard case of a real valued probability are well known. I start by sketching the first version (where the probability is defined over \mathcal{P}), later I shall remark briefly on the second.

I assume for the sake of illustration that the deductive system is based on a set of axioms, \mathcal{A} , and on modus ponens as the single inference rule (I shall later remark on another possibility). This is not a realistic picture of either human or computerized theorem proving, but the idea can be applied to more sophisticated systems. \mathcal{A} contains logical axioms (which are sufficient for deriving, via modus ponens, all logical theorems) as well as some subject-specific axioms. $P(\cdot)$ has values between 0 and 1 and is defined over \mathcal{P} . Thus, use of ' $P(\gamma)$ ' presupposes that $\gamma \in \mathcal{P}$.

- (I) $P(\alpha) = 1$, for every α in \mathcal{A} , which is in \mathcal{P} .

- (II) If α is in \mathcal{P} , so is $\neg\alpha$ and $P(\neg\alpha) = 1 - P(\alpha)$.
- (III) If $P(\alpha \wedge \beta) = 0$, and α , β and $\alpha \vee \beta$ are in \mathcal{P} , then $P(\alpha \vee \beta) = P(\alpha) + P(\beta)$.
- (IV) If $P(\alpha \rightarrow \beta) = 1$, and α and β are in \mathcal{P} , then $P(\alpha) \leq P(\beta)$.

Call a \mathcal{P} -proof, a proof (from \mathcal{A}) that consists of members of \mathcal{P} . I shall also refer to such a proof (where \mathcal{P} is understood) as a *local proof*. It is not difficult to see that, since modus ponens is the single rule of inference, (I) and (IV) imply that $P(\alpha) = 1$, for all α that are locally provable (i.e., for which there are local proofs) and (II) implies that, for such α , $P(\neg\alpha) = 0$. Together with (IV) this yields:

- (V) For α and β in \mathcal{P} , if $\alpha \rightarrow \beta$ is locally provable, then $P(\alpha) \leq P(\beta)$.

Note that the local provability of $\alpha \rightarrow \beta$ and of $\beta \rightarrow \gamma$ need not imply that the local provability of $\alpha \rightarrow \gamma$, since the derivation (in sentential logic) of the last conditional from the first two might have to use sentences that are not in \mathcal{P} . Therefore we introduce a relation \preceq defined by: $\alpha \preceq \beta$ if there are sentences, $\alpha_0, \dots, \alpha_n$, such that $\alpha = \alpha_0$, $\beta = \alpha_n$ and $\alpha_i \rightarrow \alpha_{i+1}$ is locally provable, for all $i < n$ (the case $\alpha = \beta$ is obtained by putting $n = 0$). Define $\alpha \sim \beta$ as: $\alpha \preceq \beta$ and $\beta \preceq \alpha$. It is obvious that:

- $\alpha \preceq \beta$ implies $P(\alpha) \leq P(\beta)$.
- \sim is an equivalence relation.
- If $\alpha \sim \alpha'$ and $\beta \sim \beta'$, then $\alpha \preceq \beta$ iff $\alpha' \preceq \beta'$.
- If $\alpha \sim \alpha'$, then $P(\alpha) = P(\alpha')$.

Therefore the equivalence classes of \sim form a partial ordering and $P(\cdot)$ can be viewed as an assignment of values to these classes. If \mathcal{P} is appropriately chosen, there is a “large enough” subset of it over which the partial order structure is a Boolean algebra, where Boolean operations are representable by the logical connectives. That is to say, the Boolean union of two elements, represented by α and β , is represented by some sentence $\alpha' \vee \beta'$, such that $\alpha \sim \alpha'$ and $\beta \sim \beta'$.

I shall avoid technical details, but the idea is not difficult: By restricting deductions to those that lie within \mathcal{P} , we can get, over some subset of \mathcal{P} , an equivalence relation, \sim , which induces a Boolean algebra. Call it a *local algebra*. It represents a local coherent view that combines probabilistic with resource-bound deductive reasoning. There can be α and β for which $\alpha \rightarrow \beta$ follows logically from \mathcal{A} , though α is not $\preceq \beta$. In particular, equivalent sentences (given \mathcal{A}) need not be equivalent under \sim , and can get different probabilities. The belief that α implies β can be expressed by conditionalizing on the element of the local algebra that corresponds to $\alpha \rightarrow \beta$. It is possible that when \mathcal{P} is

enlarged to some \mathcal{P}^* , the new \preceq -relation is enlarged even when restricted to \mathcal{P} (because there are more local proofs connecting members of \mathcal{P}). It is a plausible requirement that, over the original algebra (of elements represented by members of \mathcal{P}), the new probability should be the conditionalization of the original probability on the conjunction of the sentences of \mathcal{P} that are \mathcal{P}^* -provable.

Note that \mathcal{P} need not be closed downwards; that is, it may contain a sentence without containing some, or all, of its sentential components. Recognizing that $\alpha \wedge \neg \alpha$ is a contradiction, the agent assigns to it probability 0, without necessarily getting into the structure of α . Sentences can be thus encapsulated and treated as atoms, for reasons of economy. I suspect, however, that the more important limiting effect is obtained by restricting the instantiations of quantified sentences and that a downward closed \mathcal{P} (i.e., one that contains the sentential components of its members) may yield a plausible modeling.⁴ Consider, for example, an agent who believes that $\forall x \alpha(x)$ is true, and who recognizes the validity of the scheme $\forall x \alpha(x) \rightarrow \alpha(t)$, but refrains from using it for a particular t , which is too long, or too complex. In the proposed setup $P(\cdot)$ assigns 1 to $\forall x \alpha(x)$ and is undefined for $\forall x \alpha(x) \rightarrow \alpha(t)$.

Here is a sketch of the second version. Given \mathcal{P} , define \preceq and r as above. Let \mathcal{P}^+ be the closure of \mathcal{P} under sentential (i.e., Boolean) connectives. There exists a smallest equivalence relation, r^+ , which extends r , such that the equivalence classes of the members of \mathcal{P}^+ form a Boolean algebra whose Boolean operations are represented by the sentential connectives⁵ (this Boolean algebra is, of course, finite). $P(\cdot)$ is required to induce a probability distribution over this algebra; that is, if $\alpha r^+ \beta$, then $P(\alpha) = P(\beta)$ and the usual conditions for (finitely additive) probabilities hold. This is of course much stronger than the bare conditions (I)–(IV). The local algebra of the second version is the algebra whose members are represented by the sentences of \mathcal{P}^+ . As in the first, the agent's beliefs must cohere with deductive capacities that are restricted by \mathcal{P} (since the underlying partial order, \preceq , is the one established by \mathcal{P} -proofs). But the probabilities are required to satisfy a more demanding criterion: the assignment must respect the Boolean consequences of the agent's beliefs concerning the members of \mathcal{P} . In the first version, since the local algebra is encompassed within \mathcal{P} , we have to choose a rich enough \mathcal{P} in order to get a sufficiently large algebra. In the second version, \mathcal{P} can be chosen in a more

⁴ In such cases a Götzen-type deductive system is most suitable; for the valid implications of sentential logic are deducible in it by proofs that use only the sentential components of the premises and the conclusions. We should require that if $\Gamma \rightarrow \Delta$ is a theorem in the sequent calculus and the members of Γ and Δ are in \mathcal{P} , and if $\gamma, \delta \in \mathcal{P}$, where γ and δ are, respectively, the conjunction and disjunction of the members of Γ and Δ , then $P(\gamma) \leq P(\delta)$.

⁵ This follows from the fact that Boolean algebras are axiomatizable by a set of equations. $\gamma r^+ \delta$ iff this follows by a finite number of applications of the following rules: (i) If $\alpha r \beta$ then $\alpha r^+ \beta$, (ii) rules to the effect that r^+ is an equivalence relation, (iv) rules to the effect that r^+ is a congruence relation: if $\alpha r^+ \alpha'$ and $\beta r^+ \beta'$ then $\neg \alpha r^+ \neg \alpha'$ and $\alpha * \beta r^+ \alpha' * \beta'$, where $*$ is any the binary connective, (v) the axioms of an equational axiomatization of Boolean algebras; e.g., $\alpha \wedge \alpha r^+ \alpha$, $\alpha \wedge \beta r^+ \beta \wedge \alpha$, $\neg(\neg \alpha) r^+ \alpha$, etc.

restrictive (hence more realistic) manner; but the agent is required to have more (but still limited) resources regarding the probabilities.

When it comes to modeling specific situations, like the one treated in the next subsection, we may want to include in \mathcal{P} theorems, which are recognized as such by the agent, without including in it all the sentences that are used in their proofs from standard axioms. We can do this by including them among the members of \mathcal{A} . The idea is that these theorems represent known results that are taken for granted by the agent. The agent can use them without having recourse to their proofs.

The restriction on resources is expressed in the framework by a restriction to some finite set of sentences. This implies also a restriction on proof length, since we can assume that proofs do not contain sentence repetitions. But a direct restriction on proof length (or on computation length) does not lead to a Boolean algebra, since the resulting implication relation is not transitive: we can have derivations within the desired bound from α to β and from β to γ , but no derivation within this bound from α to γ . If we try to remedy this by using the transitive closure then, of course, we remove any bound altogether. The present proposal bypasses the difficulty by focusing on some finite set of sentences and by restricting proofs to those that lie within this set. There is to be sure an idealization in assuming that the agent recognizes all \mathcal{P} -provable sentences of \mathcal{P} . In particular, the search for proofs even within a small corpus of sentences can be quite difficult. But a direct restriction on search capability does not lead to an equivalence relation over sentences, because of the seemingly inevitable failure of transitivity. As I shall indicate, the present proposal can provide a picture that illuminates existing practices

While my proposed system yields formal models for a philosophical view that is very much like Hacking's, it differs from his suggestion in essential respects. Hacking bases his suggestion on restricting the number of modus ponens applications;⁶ as just noted, this cannot lead to a Boolean algebra. He assumes (implicitly) that $P(\cdot)$ is total, and makes room for violations of the probability axioms. An agent X , may think that $\neg\alpha$ is possible and assign to it a positive value and yet, believing in the truth of both $\alpha \rightarrow \beta$ and α , assign to each the value 1; for, as Hacking explains, "he has not thought of putting them together". Hacking uses *possibility* as an additional primitive concept. His two basic conditions are: $P(\alpha) = 1$ if $\neg\alpha$ is impossible, and $P(\alpha \vee \beta) = P(\alpha) + P(\beta)$, if $\alpha \wedge \beta$ is impossible. 'Impossible' means impossible by the agent's lights. Hacking does not

⁶ Hacking recognizes sentence-length as a restricted resource, but argues that sentences of unlimited length pose no problem for subjective probability and that any restriction on length is highly artificial (*ibid.* p. 318). This is inaccurate. The significance of sentence length depends on the details of the deductive system. It is conceivable that restrictions on the number of steps will have little effect if the sentences are not restricted as well. For example, using the idea underlying Craig's theorem, one can show that any recursively-enumerable theory has a recursive set of axioms, from which every theorem follows by one application of modus ponens. I do not know if something of this nature holds if we also require that the axioms be given by a finite number of schemes. Note that 'sentence length' should be construed as the size of the code required for representing the sentence, via some fixed coding scheme. A length restriction therefore implies a restriction on the number of the primitive symbols, hence also a restriction on the number of sentences altogether.

provide a formal modeling of ‘possibility’, but suggests the following: Something is possible for X , if X has not proven its negation. While X recognizes the general validity of modus ponens (the truth of α and of $\alpha \rightarrow \beta$ implies that of β), he refrains, due to his limited resources from applying it in all cases. This leaves open the possibility that $P(\alpha \rightarrow \beta) = P(\alpha) = 1$, but $P(\beta) < 1$ (whereas, on the present proposal, $P(\beta)$ should be either 1 or undefined). Such immediate incompatibilities can be explained as the agent’s failing to put α and $\alpha \rightarrow \beta$ together; but I think that this failure of omission is better modeled by letting $P(\beta)$ be undefined.

A particular aspect of the logical omniscience problem has been addressed by Garber. [1983]. This is a problem in the philosophy of science that stems from situations where the increase of knowledge consists solely in coming to recognize the deductive implications of a known theory. Garber suggested a system consisting of: (i) a language, L , of sentential logic (i.e., without predicates or quantifiers) based on an infinite list of atomic sentences, a_i , and (ii) an extension, L^* , of L , obtained by adding atomic sentences that are written as: $A \vdash B$, where A and B are sentences of L . The idea is that the atoms a_i can be interpreted as sentences belonging to some unspecified first-order language (not included in the system) and that $A \vdash B$ says that, under this interpretation, A logically implies B . The agent’s state of knowledge is represented by a probability function $P(\cdot)$, defined in the usual way over the sentences of L^* , which satisfies the equality:

$$P(A \wedge (A \vdash B)) = P(A \wedge B \wedge (A \vdash B)).$$

Taking into account the intended meaning of ‘ $A \vdash B$ ’, the equality is natural (additional requirements of this nature are considered as well). Learning that A logically implies B , the agent conditionalizes on $A \vdash B$, replacing the original $P(\cdot)$ by $P(\cdot \mid A \vdash B)$. The system does not address the question of how this learning comes about (e.g., the agent may come to believe that A logically implies B on the word of some authority), and it does not tackle bounded resources, since it avoids any proof-theoretic or computability considerations. The agent is materially, or *de facto*, omniscient concerning all logical implications between the sentences of L^* (as sentences of the sentential calculus), since the probability amounts to a function defined over the standard infinite Boolean algebra, where tautologically equivalent sentences represent the same element. The agent is not *de jure* omniscient, in as much as the fact that A tautologically implies B need not imply

$P(A \vdash B) = 1$, though it implies $P(A \rightarrow B) = 1$. The setup provides for a schematic representation of situations in which additional knowledge of logical implications is expressed through Bayesian conditionalization. As remarked above, such conditionalization is also provided for in the present proposal: if we learn that α implies β , we conditionalize on $\alpha \rightarrow \beta$. But the main point of my proposed system is the modeling of the limit on resources.

Assigning Probabilities to Arithmetical Statements.

There are two different types of probabilistic assignment to mathematical statements. A mathematician, having deliberated a problem, may have a strong conviction that a certain sentence is true and be willing bet on it with odds, say 3:1. This can be interpreted, in

probabilistic terms, as a subjective probability that is ≥ 0.75 . Willingness to bet on the negation with odds 1:4 can be read as showing that the probability is ≥ 0.80 . For all we know Mersenne, or some of his friends, might have been willing to bet on the primality of $2^{67} - 1$ with odds 20:1, but might have agreed to very small odds, say 1:100, on the negation, which would have put his subjective probability in the interval $[0.95, 0.9901]$. A different case of the same type is the following. Being asked to judge a chess situation, a chess player glances at the board and says that she is prepared to bet 4:1 that it is a winning situation for black. Far from having the complete strategy in view, she is able to evaluate the situation to an extent that gives her that level of confidence. The methodology of eliciting probabilities by considering bets applies in the case of mathematics as it applies in general. Not that I find the methodology unproblematic, but its problems have little to do with the distinction between the types of statements.

We do not need to go here into the question how precise the probability values are and what they are based on—questions that arise with respect to many other assignments: one's probability that it will rain tomorrow, that the stock market will improve, that Bin Laden is alive, that the train will be late, , and so on. Subjective probability assignments to mathematical statements are not worse in this respect than subjective probabilities in general. Justifying the particular values is not at issue; we may not be able to give an account that explains the expert's values, except for gesturing at the expert's past experience, or at some analogy with other cases. The framework of subjective probability does not require that such an account be given. The probabilities may guide a person, without being anchored in objective frequencies, and they may be systematically employed through the methodology of getting the values from experts. Here, again, there should be no distinction between the empirical and the purely deductive. In principle, there can be experts who specialize in certain types of combinatorial problems, just as there are experts that provide probabilities for finding oil. In both domains we can, in principle, get reliable judgment that falls short of conclusive proof.

The second type of probabilistic assignment is based on standard statistical methodology. It can be applied across the board in the empirical domain and, it turns out, also in evaluating mathematical statements, or mathematically defined values. Monte Carlo methods have long been employed in applied mathematics, in order to derive very reliable estimates of mathematical magnitudes that arise in analysis (e.g., the values of some integrals). Better illustrations for my present purpose are provided in numerous probabilistic algorithms that answer arithmetical, or combinatorial questions. One common kind of probabilistic algorithms gives “yes”/“no” answers to mathematical questions of the form, $\alpha(t)$? where $\alpha(x)$ is a formula expressing a mathematical property of x (e.g., ‘ x is prime’) and where t is a varying input (e.g., a natural number in some standard notation). The answer—produced by running the algorithm a number of steps, which depends on the input—is true with high probability; the probability can be pushed up, as near to 1 as desired, at the cost of running the algorithm longer.

Rabin [1976], [1980] proposed a probabilistic method for finding large primes that is now commonly used; a variant was independently found by Solovay and Strassen [1977]. For the purpose of the present discussion, we need only the general form of such methods and need not go into the finer details. The algorithm is based on a

compositeness test, **Test**(a, \mathbf{x}), which yields, on input a, \mathbf{x} , a positive or negative answer. If the answer is positive the number \mathbf{x} is composite (not prime), and a is said to be a (compositeness) *witness* for \mathbf{x} . A negative answer leaves the question of primality open. Underlying the algorithm is a theorem that says that for every composite number, at least 3/4 of the numbers smaller than it are witnesses. Given \mathbf{x} , the probabilistic algorithm chooses randomly a number, a , smaller than \mathbf{x} , and runs **Test**(a, \mathbf{x}). This is repeated until either a witness is found, or the number of negative tests reaches a certain bound, say 10. If the bound is reached \mathbf{x} is said to be a probable prime. If \mathbf{x} is not prime then the probability of getting 10 negative tests is $(1/4)^{10}$. This, we shall see, implies that the probability that a probable prime is prime is extremely high. To push it higher, one increases the number of tests. The algorithm is used to find many large primes and its practical value derives from the use of such primes in cryptography. Its efficiency derives from the fact that performing many tests requires incomparably shorter time than the time required by non-probabilistic primality tests.

It is sometimes claimed that the probability that a probable prime is not prime is 4^{-k} where k is the number of tests. This is another instance of the common error of conflating likelihood ratios with probabilities. The right estimate is the following. For a fixed number \mathbf{x} , let h be the hypothesis that \mathbf{x} is prime; let e be the evidence that 10 tests were negative, and let $P(A | B)$ be the conditional probability of A , given B . Then $P(e | \neg h) = 4^{-10}$, $P(e | h) = 1$. The well-known equalities for conditional probabilities give:

$$\frac{P(h|e)}{P(\neg h|e)} = \frac{P(e | h)}{P(e | \neg h)} \cdot \frac{P(h)}{P(\neg h)} = 4^{10} \cdot \frac{P(h)}{P(\neg h)}$$

Let $r = \frac{P(h)}{P(\neg h)} = \frac{P(h)}{1 - P(h)}$; elementary algebra yields:

$$P(h | e) = r \cdot \frac{1}{1 + 4^{-10}r} \geq 1 - 4^{-10}r^{-1}$$

Thus, the probability that \mathbf{x} is prime, given the evidence, depends on the prior probability that it is prime. This dependence on prior probabilities is a standard problem in statistics. Assuming that \mathbf{x} is chosen randomly from a certain range of numbers, we can use well-known results about the distribution of primes. If, for example, \mathbf{x} is chosen from the numbers between 2^{100} and 2^{1000} , then, since the frequency of primes in that interval is $\geq 1/1000$, $P(h)$ is at least 10^{-3} , implying that $P(h | e) > 1 - 2^{-10}$. Actual algorithms for finding large primes incorporate a preliminary screening that weeds out numbers with small factors; this yields a much larger probability. The probability approaches 1 fast when the number of tests increases. With 50 tests, $P(h | e) > 1 - 2^{-90}$.

The case in which there is no clear-cut prior probability is a standard difficulty in Bayesian methodology. Non-Bayesian systems, such as the Neyman-Pearson approach, employ methods that do not rely on priors. I do not wish to enter here into the well-

known debates surrounding the various approaches. My local-algebra based proposal belongs to the general Bayesian framework. But it should be noted that the lines dividing the statistical schools cut across the distinction between the mathematical and the empirical, and their relative strengths and weaknesses do not derive from that distinction.

$P(h)$ (the prior probability of h) reflects the viewpoint of the agent, who may decide to derive this value from relative frequency, even in the absence of actual randomization. If the likelihood ratio is large enough to swamp differences between various priors, then the effect of choosing among them is minimized. This can be done, at the cost of increasing the number of tests. Note that the objective probabilities that underlie the algorithm, which endow it with robustness, do not derive from a random choice of \mathbf{x} , but from the random samplings of the numbers a that are used in the tests.

So far this is a standard piece of statistical reasoning. It makes no difference that h is a mathematical conjecture about \mathbf{x} , as long as we are not given a mathematical description of that number. In this situation ' \mathbf{x} ' figures as some non-mathematical specification, like 'the number tested on occasion Ξ '. Alternatively, we can treat it as an unstructured name about which we know only that it denotes some number. The situation changes if we add an equality of the form:

$$\mathbf{x} = t$$

where ' t ' is a mathematical term that constitutes a mathematical description of \mathbf{x} ; e.g.,

$$\mathbf{x} = 2^{400} - 593.^7$$

With \mathbf{x} thus specified, we can check directly, by a non-probabilistic computation, if it is prime. But we may choose to ignore the additional information that $\xi=t$ if we judge that the price of extracting from it the required answer is too high. Yet, we use this information in order to substitute ' \mathbf{x} ' by ' t ' in the sentence ' \mathbf{x} is prime'. If h' is the hypothesis that t is prime, we conclude that $P(h'|e) > 1-2^{-10}$. That is: given the evidence of the ten randomly chosen tests, the probability that $2^{400} - 593$ is prime is $> 1 - 2^{-10}$. Furthermore, the mathematical term ' t ' can be used in additional computations and the assumption that t is prime can serve as a basis for further mathematical conclusions, e.g., $2^t \equiv 2 \pmod{t}$, which are implied by the primality of t . Such conclusions have probabilities \geq the probability that t is prime. The probabilities can be seen as the values of an assignment, defined over a domain of sentences that represents the resources that a computer-aided agent brings to bear on various problems involving t . The assignment reflects a coherent view that integrates various pieces of information, statistical and purely deductive, under the limitation that restricts all deductions to those consisting of members of a certain set, \mathcal{P} .

Note that it is conceivable that by pure chance we end up with t whose primality can be decided by our restricted resources. (This would be analogous to cases where the name

⁷ The primality of this number was probabilistically established in 1975 by V. Pratt, using the then unpublished method of Rabin. It took no more than several minutes on a slow computer (oral communication).

of a person gives us easy clue about some personal attribute, say, ethnic origin.) \mathcal{P} can include some known theorems, treated as axioms, which enable us to recognize, via a local proof, whether t is prime or not. (In practice, the possibility can be ignored.) In this case the local probability function assigns a 0-1 value to $\text{prime}(t)$ —the sentence asserting that t is prime. The following then holds, where e is the evidence of the primality tests and where ξ is, as above, a non-informative name.

$$P(\text{prime}(t) \mid e) = P(\text{prime}(t) \mid e \wedge (\mathbf{x} = t)) \neq P(\text{prime}(\mathbf{x}) \mid e \wedge (\mathbf{x} = t)) \neq P(\text{prime}(\mathbf{x}) \mid e)$$

In principle, the non-equality $P(\text{prime}(\mathbf{x}) \mid e \wedge (\mathbf{x} = t)) \neq P(\text{prime}(\mathbf{x}) \mid e)$ can obtain also in the case where some connections established by local proofs (in which the term t figures) give $P(\text{prime}(t) \mid e)$ some value between 0 and 1.

Here is a brief sketch of \mathcal{P} . The language in which the sentences of \mathcal{P} are formed contains: (i) an arithmetical vocabulary sufficient for expressing the usual notions of first-order arithmetic, (ii) a statistical vocabulary for describing the randomly chosen numbers $< t$, which are used in the tests. The statistical vocabulary consists of symbols ' \mathbf{x}_1 ', ' \mathbf{x}_2 ', ..., ' \mathbf{x}_k ' denoting k independent random variables, uniformly distributed over the integers between 1 and t , whose values are used in the compositeness tests. If $\text{Test}(x, y)$ is the wff saying that x is not a witness for y , then the evidence e , of the negative results of 10 tests, is expressed as the conjunction of $\text{Test}(\mathbf{x}_i, t)$, $i = 1, \dots, 10$. These and other related sentences are in \mathcal{P} . The probabilities are obtained in the standard way from the joint distribution over the values of the \mathbf{x}_i 's. Note that we do not need to know the actual values of the \mathbf{x}_i 's that figure in the tests. Hence we do not have to include in \mathcal{P} all the sentences $\text{Test}(\underline{n}, t)$, where ' \underline{n} ' ranges over standard names of the numbers $< t$.

The purely arithmetical sentences of \mathcal{P} are adequate for carrying out various arithmetical deductions, but not for carrying out the deduction that decides the primality of t . The sentence that asserts that t does not have proper divisors is $\forall y[1 < y < t \rightarrow \neg(y \mid t)]$, where $y \mid x$ says that y divides x is in \mathcal{P} . This sentence is in \mathcal{P} , but only a limited number of its instantiations are. Few (relatively speaking) of the sentences $\neg(s \mid t)$, where ' s ' is a term denoting a number $< t$, are in \mathcal{P} : those in which s is small, or sufficiently simple.

This prevents one from proving inside \mathcal{P} that t is prime by brute-force checking, and makes place for a probabilistic assignment of a value strictly between 0 and 1. All in all the probability that is defined over the local algebra is a coherent view, which expresses the epistemic state of the agent, who chooses to base further reasoning on the probabilistic estimate rather than on a deductive proof. To what extent is this choice justified?

Justification

Justification derives from the limited resources, which amounts in this case to considerations of cost. The basic principle, argued for by Hacking, is that costs of

computation and costs of empirical inquiry are on a par. An example can make the point best. We are given a collection of 1000 stones, of which 20 are of type T1 and the rest of type T2. To all appearances the types are indistinguishable, since the difference derives from a minute change of molecular form, which can be only detected by a costly test. There is also a very cheap test, with the property that the answer “T2” is completely reliable, but there is an extremely small chance of error if it answers “T1”. We need a T1 stone. So we choose randomly from the 1000 stones and conduct the cheap test. We repeat this until we get the answer “T1”. Standard Bayesian analysis shows that the chance of error is sufficiently small to justify (costs being what they are) a decision to forgo the crucial test and use the stone as a T1 stone. In the end the justification appeals to expected utility, or some other methodology of assessing risks and benefits, which is based on given probabilities.

Now consider a parallel story in which the stones are inscribed with distinct natural numbers (in some standard notation) 20 of which are prime. These are very large numbers for which non-probabilistic testing with the computer we have is extremely expensive. We need one of the 20 primes; that is, we need to know its inscribed name. The cheap test consists in running the compositeness test on randomly chosen numbers up to k times (for a suitable k). The probability of error is the same as in the first story. If the Bayesian analysis and the resulting decision are justified in the first case, then I cannot see why they are not justified in the second. In both cases we have to find an object (a stone in one case, a number in the second) of a certain type. In each case the type of any given object is predetermined and can be found either by an expensive procedure, or by a cheap one that has a very small chance of error; the procedures involve chemical tests in the first case, computations in the second. The error probabilities are the same. In the first story we end with a particular stone (we hold it in our hands); in the second story we end with a particular number (we have mathematical name of it). Let h_1 be the hypothesis that the stone we are holding is of type T1. Let h_2 be the hypothesis that the number whose name we have is prime. Let e_1 and e_2 be the respective evidences. I suggest that if we admit a probabilistic judgment of the form: $P(h_1|e_1) > 1-\epsilon$, we should also admit the judgment $P(h_2|e_2) > 1-\epsilon$. The proposed framework extends this probabilistic assignment to a wider domain that contains a limited portion of the deductive apparatus of first-order arithmetic.

3. Limited Perspectives

Carnap’s so-called requirement of total evidence enjoins us to take into account all available evidence. This seemingly common sense dictum is flatly contravened in examples of the kind just discussed. The most relevant evidence that decides the question can be deliberately ignored, if processing it is too costly. The modified principle should therefore be: use evidence, as long as the use is cost effective.⁸ The qualification should

⁸ Finding whether a procedure is cost effective may, in itself, require computations, whose cost effectiveness has to be established. This might lead to an infinite regress, a worry voiced by Savage [1967]. In fact, we do not start from nowhere. Some ideas of cost, which can be crude or mistaken, serve as a starting point. Eventually it is possible to establish sound results about cost effectiveness and, having established them, we can go back and see that the obtaining of these results was worth the invested effort. And if it was not,... well, these things happen; one can make mistakes.

put rational reasoning in quite a different light from the customary idealization that informs our picture of a thinking agent. More is actually true: the ignoring of possibly relevant aspects is a necessary precondition for reasoning in general. It is constitutive to reasoning that it is local; it takes place during a limited period of time, with a restricted scope of awareness, subject to entrenched conceptual habits. Let me consider first, by way of example, the effect of the limited perspective on learning from experience. (Here I use ‘perspective’ in a general non-formal.)

The most general systematic account of learning from experience is given in the Bayesian framework. The agent’s belief state is represented by a subjective probability function; learning from experience is achieved by conditionalization: If $P(\cdot)$ is the function representing my state then, upon obtaining evidence E , my new state is given by the conditional probability $P(\cdot|E)$, where $P(A|E) = \frac{P(A \cap E)}{P(E)}$. I have switched to

customary notation, since these probabilities are assigned to events, or propositions. Davidson, [1976] p. 273, claimed that Bayesian decision theory is constitutive to rational behavior, just as length and mass measurement are to physics, or Tarski’s theory—to truth and meaning:

The theory in each case is so powerful and simple, and so constitutive of the concepts assumed by further satisfactory theory (physical and linguistic) that we must strain to fit our findings, or our interpretations, to preserve the theory.

Now Bayesian conditionalization applies neatly in numerous contexts, but it cannot be applied globally by creatures with restricted computational resources. The picture is therefore misleading. Here is a simple illustration.

A coin is tossed 50 times. Jack analyzes the sequence of 0’s (heads) and 1’s (tails) and concludes that the sequence is random, with nearly equal probabilities for each side. Then Jill informs him that the coin contains a tiny mechanism, controlled by a microchip, which shifts its gravity center so as to produce a pseudo random sequence—one that passes the standard randomness tests but is in fact defined by a mathematical algorithm.⁹ Jill’s claim is not impossible (imagine the coin to be very thick, or imagine that the experiment is conducted with a die, with three sides marked ‘0’ and three marked ‘1’.) Jill gives Jack the algorithm and the outcomes of additional 50 tosses fully accord with the calculated predictions. Let h be the hypothesis that the outcome sequence coincides with the sequences produced by the algorithm. Let the evidence be e . Given e , it would be irrational of Jack not to accord h extremely high probability. Had $P(h)$ —Jack’s prior subjective probability for h —been non-zero, conditionalization on e could have pushed it up. But if $P(h) = 0$, then, $P(h|e) = 0$, for every e for which $P(e) > 0$. Now Jack’s prior probability accommodates the possibility of independent tosses with non-

⁹ Here is such a pseudo random sequence: 0, 1, 0, 1, 1, 1, 0, 1, 1, 1, 0, 1, 1, 1, 1, 0, 0, 0, 0, 0, 1, 0, 0, 0, 1, 1, 1, 0, 0, 1, 0, 1, 1, 0, 1, 0, 1, 1, 1, 0, 0, 1, 1, 0, 0, 0, 1, and here is its continuation produced by additional output of the algorithm: 1, 0, 0, 1, 1, 1, 1, 0, 1, 0, 0, 0, 1, 1, 0, 1, 0, 1, 0, 1, 0, 1, 1, 0, 1, 0, 0, 1, 1, 0, 0, 0, 0, 1, 1, 1, 1, 0, 1, 0, 0, 1, 0, 1, 0, 1.

zero probabilities for '0' and '1'. Therefore $P(\cdot)$ assigns each finite sequence a non-zero value, hence $P(e) > 0$. If $P(h) = 0$, then Jack cannot learn from e what he should, unless he changes probabilities in a way that violates the Bayesian prescription. Jack has never considered the possibility of a coin that produces, deterministically, a mathematically defined sequence; hence, if the prior probability is to reflect anything like Jack's psychological state, $P(h) = 0$. But could he, in principle, have made place for such possibilities in advance?

Making place means that the prior accords positive probability to each hypothesis that says that the sequence of outcomes equals to the sequence defined by such and such a mathematical formula. Since the underlying language is countable, such priors exist. The trouble however is that the more hypotheses we want to accommodate, the more complex the accommodating prior should be. A claim of this nature was proven in [Gaifman-Snir 1982], where a first-order language containing arithmetic served to state the hypotheses as well as to define the priors, and where “complexity” was based on the arithmetic hierarchy.¹⁰ More relevant to the present concern are formulas that describe algorithms, and a complexity measure that is based on hardness-to-compute criteria. Some results in this direction can be established; for certain natural classes of algorithms, computing the values of a prior that accords non-zero probabilities to all members of the class (i.e., to every hypothesis that the given sequence is defined by that member) requires an algorithm more powerful than the class members. It is likely that the phenomenon is general: Non-dogmatism (the accordance of non-zero prior probabilities to larger classes of hypotheses) must be bought at the price of increased computational complexity.

These considerations show that large scale Bayesian updating that does not restrict our ability to learn from experience requires something like God's eye view. Limited deductive ability implies that all perspectives are local.

At the level of everyday reasoning this fact is quite familiar. To accommodate within one's thought an increasing number of possibilities is to slow our reasoning; eventually this becomes too high a price. Thus, efficiency requires that many possibilities be ignored. We become aware of them when the evidence hits us, and then we revise our beliefs abruptly, not by conditionalization.

There can be also abrupt changes that are not caused by new evidence. One may become aware of a pattern that puts the situation in a completely different light. Say Jill is presented with an apparently random binary sequence of length 50, with a relative frequency 0.3 of 1's.¹¹ Perhaps the data suggests also the possibility that the probability of 1 decreases as the sequence progresses. Suddenly she realizes that 1 occurs exactly at

¹⁰ Since the probability is real-valued, its definition in arithmetic means the definition of the ternary relation $a < P(\alpha) < b$, where a and b range over the rational numbers (given as pairs of natural numbers) and where α ranges over all finite binary sequences (i.e., events that specify the outcome of the first tosses; these events generate the whole field). The result states that a prior that accommodates all hypotheses belonging to a certain level of the arithmetical hierarchy must itself belong to a higher level.

[illegible]

the prime places (the 2nd, the 3rd, the 5th, the 7th, etc.). A hypothesis that is incompatible with randomness and was not within Jill's perspective, becomes extremely probable.

A belief state can change merely by having one's attention called to certain possibilities. Imagine that I teach an advanced graduate seminar with 8 participants (a fictitious but not unrealistic scenario). My rough impression is that during a semester an average student will miss—for various reasons: sickness, visiting relatives, a job interview, a paper to finish—two out of twelve classes. As I rush to class it will never occur to me that the seminar might not take place because, by sheer coincidence, no student will show. If I were asked, I will dismiss the possibility. Yet the scenario is not impossible. If the events are independent, and for each student the probability of missing class is $2/12$ (simplifying assumptions to be sure), the probability that no student will show is $(1/6)^8$, hence $> 1.700.000^{-1}$. Consider, by contrast, a one-prize lottery with 1.700.000 participants. The probability that Beth will win is $1.700.000^{-1}$; though it is very small, I will be very much aware of this possibility and will not dismiss it. There are some clear reasons for the difference. A lottery is a paradigmatic example of a class of equal possibilities. The symmetry that makes them all equal provides easy reasoning: Someone must win, why not Beth? By comparison, the probabilistic estimate in the first scenario is based on extrapolating the lottery-based notion in a sophisticated, far from obvious way. Yet someone with adequate grasp of probabilities, may, given time for reflection, admit the possibility that my seminar will not take place, on a par with the possibility that Beth will win the lottery. (As a measure of caution one may settle for a smaller lower bound, say $10.000.000^{-1}$; the point still remains). Our different reaction to the two scenarios is due to a conceptual organization in which the lottery scenario is self-evident, while the seminar scenario requires non-trivial analysis. That deep factors underlie this conceptual organization—just as deep factors underlie what is easy and what is difficult in mathematics—is not at issue here. The point is that there is no principled difference in the application of probabilistic reasoning in the two cases. For those who know already some probability, the difference amounts essentially to a difference in cost of reasoning.

Now reflect, dear reader, how ubiquitous is the ignoring of unlikely possibilities, and how necessary it is for the conducting of everyday life. I hardly think of crossing the street, taking a subway, mailing an important letter, as lotteries in which there is a small chance of failure. All my perspectives are local. On single occasions, if my attention is called to some known facts and time is not pressing, I might pause, reflect and accept some remote, hitherto ignored, possibility as real. In doing this I have already changed my perspective. In a philosophical mood and with time on one's hands, one might even entertain skeptical scenarios and—given a certain temperament—worry about them, as some philosophers did.

If this is right, then notions such as “X at time t takes A as a serious possibility”, or “X fully believes A at time t”—on which Levi, [1991] and [1995], bases his theory of belief—face the same difficulty that is faced by the subjective probability framework. What one takes to be possible is sensitive to the local perspective. Sometimes, for particular purposes, we can ignore these aspects. But, in general, ignoring them yields the wrong kind of picture.

References

Bell, E.T. [1951] “The Queen of Mathematics”, reprinted in *The World of Mathematics*, ed. J.R. Newman, Simon and Schuster 1956.

Davidson, D. [1976] “Hempel on Explaining Action”, *Erkenntnis* 10, pp. 239-253. Reprinted in Davidson’s *Essays on Action and Events* Oxford, 1989. Page numbers refer to the book.

De Milo, R. Lipton, and A. Perlis [1979] “Social Processes and Proofs of Theorems and Programs,” *Communication of the ACM*, vol. 22, No. 5. Reprinted in *New Directions in the Philosophy of Mathematics*, ed. T. Tymoczko, 1998, Princeton University Press. Page numbers refer to the book.

Gaifman, H. [1964] “Concerning measures in first order calculi” *Israeli Journal of Mathematics*, vol. 2, pp. 1 - 18.

Gaifman, H. and M. Snir [1982] “Probabilities over rich languages, testing and randomness” *Journal of Symbolic Logic*, vol. 43, No. 3, pp. 495 –548.

Garber, D. [1983] “Old Evidence and Logical Omniscience in Bayesian Confirmation Theory”, in *Testing Scientific Theories* ed. J. Earman, Minnesota Studies in the Philosophy of Language, vol. X pp. 99–131.

Hacking, I. [1967] “Slightly More Realistic Personal Probability” *Philosophy of Science*, 34, pp. 311–325.

Levi, I. [1991] *The Fixation of Belief and Its Undoing*, Cambridge University Press.

Levi, I. [1995] “The logic of full belief”, in Levi’s *The Covenant of Reason*, Cambridge University Press 1997, chapter 3, pp. 40-69. Page numbers refer to the book.

Rabin, M. [1976] “Probabilistic Algorithms” in *Algorithms and Complexity, Recent Results and New Directions*, ed. J. Traub, Academic Press, New York, pp. 21–39.

Rabin, M. [1980] “Probabilistic Algorithm for testing primality”, *Journal of Number Theory*, 12 pp. 128–138

Savage, L. [1967] “Difficulty in the Theory of Personal Probability” *Philosophy of Science*, 34.

Solovay, R. and V. Strassen [1977] “A fast Monte Carlo test for primality” *SIAM Journal of Computation*, 6, pp. 84–85; Erratum, *ibidem* 6 (1978) p. 118.