Practical Aspects of Theoretical Reasoning

Gilbert Harman

Princeton University

Albert thinks about what route to take to get to Boston. He thinks that, while the direct western route is faster, the scenic eastern route is longer but more enjoyable with less traffic. He is in a bit of a hurry but could probably arrive on time going either way. He eventually reaches a decision.

The reasoning Albert goes through in settling on what route to take is *practical*. He is deciding what to do.

At about the same time, Albert's friend Betty tries to decide what route Albert will take. She thinks about what Albert has done before, what Albert likes in a route, and how much of a hurry Albert is in. Betty's reasoning is *theoretical*. She is trying to arrive at a belief about what Albert will do.

Practical reasoning in this more or less technical sense leads to (or modifies) intentions, plans, and decisions. Theoretical reasoning in the corresponding technical sense leads to (or modifies) beliefs and expectations. There is also the possibility that reasoning of either sort leaves things unchanged.

Any given instance of reasoning may combine both theoretical and practical reasoning. In deciding which route to take, Albert may have to reach theoretical conclusions about how long it will take to go by the eastern route. In thinking about

which route Albert will take, Betty may have to reason practically about whether to check her records about Albert's past trips to Boston.

Nevertheless, there is a difference between theoretical reasoning and practical reasoning and a corresponding difference between theoretical reasons and practical reasons. In particular, there is a distinction between theoretical reasons to believe something and practical reasons to believe something. For example, Samantha has theoretical reasons to believe that knowledge of the history of philosophy is not very useful in actually doing good philosophy today, reasons based on a careful study of the history of philosophy and of the best recent philosophical literature. On the other hand, she has practical reasons to believe that knowledge of the history of philosophy is very useful in actually doing philosophy today, because she wants to be hired by a philosophy department that has a policy of only hiring candidates who believe that a solid knowledge of the history of philosophy is very useful to anyone who tries to do philosophy today.

A purely theoretical reason to believe something is sometimes called an *epistemic* reason to believe it, in contrast with a *nonepistemic* practical reason (Foley, 1987).

There are interesting questions about how and to what extent practical reasons might be relevant to theoretical reasoning, strictly so-called. Practical reasons are certainly relevant to whether to undertake theoretical reasoning about a particular subject. Practical considerations may also be reflected in the role played by conservatism and simplicity in theoretical reasoning.

Preliminaries

Before discussing these issues, I need to discuss some preliminary points.

Reasoning and "logic"

There is a use of the term "logic" to mean the theory of reasoning or inquiry, as in Hegel's *Logic* (1812) or Mill's *Logic* (1869). But in contemporary philosophy the term "logic" is often used for a theory of implication and inconsistency, as in accounts of truth-functional logic, quantificational logic, and modal logic. Terminology is not important, but it is important not to confuse issues about reasoning and inquiry with issues about implication and inconsistency.

Reasoning or inquiry is a process by which you change (or don't change) your views. A theory of reasoning or inquiry is a descriptive or normative theory of that process. The theory of implication and consistency concerns abstract properties of propositions and abstract relations between propositions. That is not an especially normative subject and it does not have an especially psychological subject matter. We can meaningfully ask whether the theory of implication and consistency has any special relevance to the theory of reasoning and inquiry, a question that is often hidden from view by the ambiguity of the term "logic."

Reasoning or inquiry should not be identified with the construction of an argument or proof, although it may sometimes involve such construction. Even when reasoning does lead to the construction of an argument or proof, the process of reasoning does not normally begin by first considering the premises, then moving through intermediate steps, and finally ending with the conclusion. Anyone who has taken elementary

geometry knows that proofs or arguments are sometimes constructed backwards, from the conclusion through intermediate steps in reverse order to the premises. More often, you start in the middle and move both backwards and forwards in constructing an argument. Furthermore, when reasoning involves the construction of a proof or argument, the conclusion of your reasoning isn't typically the same as the conclusion of your argument. For example, in inference to the best explanation your conclusion may be the "premise" of an explanatory argument whose "conclusion" is something that you started out believing and that the argument serves to explain.

An argument or proof is an abstract structure of propositions, consisting of initial premises, intermediate steps, and final conclusion. A formal system of proof might state certain misnamed "rules of inference" and require that each step in an argument should either be a premise or should follow from previous steps in accordance with one of the so-called rules of inference. Such "rules" are about implication, not inference, and they are "rules" only in the sense that they are constraints on what structures count as formal arguments in that system. They are rules that proofs must satisfy, not rules for reasoners to follow.

A valid proof or argument shows some of the implications of the premises. Of course, the premises imply themselves, so a typical proof shows additional implications. Reasoning, on the other hand, does not just add further conclusions to things you already accept. It typically also involves giving up some of things previously accepted. If we describe what you initially accept as "premises," then we have to say that reasoning often involves abandoning premises and not just accepting further conclusions. (But it is best

not to use the term “premise” in discussing reasoning, because of this term's association with arguments and proofs.)

It is not easy to specify a special connection between reasoning and the theory of implication and consistency. For example, although sometimes the fact that your prior beliefs logically imply a conclusion may give you a reason to accept that conclusion, this does not hold in the general case. For one thing, you may not realize that the implication holds. For another, even if you do recognize the implication, the conclusion may be implausible, so that the implication may give you a reason to reject a prior belief rather than a reason to accept the conclusion (and it may not be true that there is a particular prior belief that one have a reason to reject). Even when the conclusion is not antecedently implausible, you will in the general case have no reason to be interested in whether it is true and so no reason to add it to your beliefs.

Distinguishing Theoretical from Practical Reasoning

I started this chapter with a contrast between Albert trying to decide which route to take (practical reasoning) and Betty trying to decide which route Albert is taking (theoretical reasoning). These examples suggest that to a first approximation theoretical reasoning is concerned with deciding what to believe and practical reasoning is concerned with deciding what to do. To a second approximation, we can say that theoretical reasoning is a process by which *in the first instance* you change your beliefs and expectations and that practical reasoning is a process by which *in the first instance* you change your choices, plans, and intentions. We have to say something like “in the first instance” because

changing what you plan to do can affect what you believe will happen and changing your beliefs may lead you to change your plans.

There are obvious similarities between theoretical and practical reasoning. In both cases you start with antecedent beliefs and intentions and reason in a way that makes changes in those beliefs and intentions typically by subtracting some and adding others. (In the limiting case reasoning leaves things as they were at the beginning, with no change.)

But there are also important differences between theoretical and practical reasoning. A very important difference has to do with wishful thinking, which is perfectly proper in practical reasoning in a way that it is not proper in theoretical reasoning. Albert's preference for the eastern route can give him a practical reason to take the eastern route rather than the western route. But Betty's preference for Albert to be taking the eastern route does not in the same way give her a theoretical reason to believe that he is taking the eastern route.

Another important difference between theoretical and practical reasoning has to do with the reasonableness of arbitrary choices. Suppose Albert is trying to decide whether to take the eastern route or the western route and he finds that nothing favors one route over the other. Then it is reasonable for him to decide arbitrarily to take one of the two routes. If it is urgent for him to get to Boston, it would be a mistake for him to suspend judgment in this case. On the other hand, if Betty is trying to decide which route Albert is taking and there is no particular reason to think he is going one way rather than the other, it is *not* reasonable for her to decide arbitrarily that he is taking the one route rather

than the other. In the theoretical case, Betty should suspend judgment. In the practical case, Albert's suspending decision can be deeply irrational.

The point about wishful thinking indicates a way in which a practical consideration deriving from your goals and desires is *not* properly relevant to your theoretical reasoning. But there are other ways in which practical considerations of this sort are properly relevant to your theoretical reasoning.

Practical Reasons to Reason Theoretically

Betty may have practical reasons to intercept Albert before Albert gets to Boston, so Betty may have practical reasons to figure out whether Albert is taking the eastern route or the western route. This illustrates one obvious way in which practical considerations can be relevant to theoretical reasoning, namely, by being relevant to what to reason theoretically about.

A related point is that reasoning uses resources like time and concentration (Simon, 1957; Gigerenzer et al., 1999). You have limited resources and reasoning about one issue keeps you from considering another. So, you have practical reasons to consider only certain questions rather than others. Practical considerations are also relevant to how much effort you should devote to investigating a given issue.

If Betty didn't have a reason to care which way Albert was going home, she would not have a reason to think about which way he was going home and she would not have a reason to reach any conclusion about which way he was going home.

For example, as already mentioned, the fact that Betty believes things that logically imply a given conclusion does not mean that she has sufficient reason to believe that conclusion. She may have no reason to be interested in whether that conclusion is true and every reason to be thinking about something else. Betty's beliefs logically imply infinitely many conclusions, most of absolutely no interest to her. She has to decide where to devote her resources. She should not clutter her mind with trivial consequences of her beliefs.

Suppose David points out to Betty that certain of her beliefs about roads in and near Princeton cannot all be true. Betty believes that Route 1 runs North-South, that Nassau Street runs East-West, and that Route 1 is parallel to Nassau Street. On discovering this conflict in her beliefs, Betty is not rationally required to drop everything to figure out which to abandon. She may have better things to do with her time. Maybe she should have lunch first. Maybe she simply has no reason to care that her beliefs are inconsistent in this way.[1]

Given resource limits, practical considerations are relevant to how much in the way of resources to devote to a given inquiry and to when to end an inquiry. The police have to decide which cases to investigate, how much effort to put into each investigation, which cases to keep open and which to close. A scientific researcher faces the same question of where to devote resources. So do the rest of us all the time.

To reach a conclusion is, among other things, to conclude an investigation. Practical reasons are relevant to reaching a conclusion, at least to the extent that they are relevant to whether to stop devoting resources to that investigation. This is not to say that practical

reasons can properly be used to decide between several competing theoretical conclusions. But practical reasons can properly be relevant to whether to end inquiry, for example on the grounds that further investigation is not likely to be worth the effort.

Conservation, Simplicity, and Coherence

Relevant factors in theoretical reasoning include conservatism, simplicity, and coherence. Roughly speaking, starting with an initial view, you try to retain as much as possible of that initial view (conservatism), to favor simpler over more complex hypotheses (simplicity), to reduce inconsistency (negative coherence), and to find explanations of things in which you are interested (positive coherence).[2]

Someone might ask what justifies a reliance on such factors as conservatism, simplicity, and coherence in our theoretical reasoning. Perhaps such reliance involves the sort of wishful thinking that we normally suppose is not theoretically reasonable. Maybe it's just that we want our present views to be correct and we don't want to have to change our minds. And maybe we want the general principles and theories we accept to be relatively simple because we have an aesthetic preference for simplicity or because it is easier for us to use simpler theories.

Perhaps reliance on conservatism, simplicity, and coherence can be justified as promoting our goals, in the way that believing in the usefulness of the history of philosophy might promote Samantha's goal of being hired by the Mooseton Philosophy Department. But then our reasoning would seem to be practical rather than theoretical, because the relevant considerations would be practical, not purely epistemic.

I now want to look in more detail at simplicity and conservatism in order to assess this sort of worry.

Simplicity

The first point then is that theoretical reasoning often favors simpler hypotheses over more complex ones. For example, suppose we have reason to believe that some quantity y is a function of a quantity x, and we are trying to figure out what the function is, given data about particular cases. We are trying to discover the function f such that $y = f(x)$.

If we have quite a bit of data and a linear function (of the form $y = f(x) = ax + b$ for constant a and b) fits the data, then even though there are also infinitely many more complex functions that also fit, we will be much more inclined to believe that the function is linear than that it is one of the more complex functions.

What explains this inclination toward for simpler hypotheses? It is not exactly that we assume or presuppose that the world is simple. Our inclination is to accept the simpler of two hypotheses that account equally well for the data. Data can and will lead us to reject all of the absolutely simplest hypotheses. Our use of simplicity can and will sometimes lead us to accept very complex hypotheses, having rejected all the simpler ones.

Roughly speaking, we reason as if we accepted a conditional probability function p such that, if hypothesis h_1 is simpler than h_2 and e is our evidence, then $p(h_1|e) > p(h_2|e)$.

I say that is only "roughly speaking" because in real life we trade off simplicity and data coverage. We allow for measurement error and noise in the data. Often, none of the hypotheses we consider fits the data perfectly. We weigh the extent to which a

hypothesis fails to fit the data (perhaps as measured by the sum of squared error) against its complexity (measured in some way or other), trying to minimize some function of these two quantities. This means that our present evidential data *e* will often in practice be actually inconsistent with the hypothesis we end up inferring from that data, so strictly speaking the relevant conditional probabilities will be zero and the wording in the preceding paragraph is inaccurate.

Still, we favor a simpler hypothesis over infinitely many more complicated hypothesis that do equally well or better at data coverage. We reason as if we believed that the simpler hypothesis is more likely to be correct in this case. But why should we believe this?

Actually, We do not exactly "believe" this. Our preference for simplicity is "built into" our system of reasoning—as it were, part of our initial probability distribution. It is a basic aspect of our epistemic probability.[3] . Our inferential practice treats simpler hypotheses as more epistemically probable than corresponding more complex hypotheses that account equally well for the data.

Inductive Bias

Once we realize that we are influenced by simplicity in this way, we can ask whether we should continue to allow ourselves to be influenced. We can ask why we should go along with this tendency in our reasoning practices.

One thing that seems relevant is that a reasoning system needs inductive bias if it is to reach any inductive conclusions at all.[4] A system without inductive bias cannot learn from experience. Now, perhaps certain entities can survive without learning, but

ordinary people cannot. Some bias that will enable learning has to be built into us somehow, perhaps as the result of evolution by natural selection.

Example

Go back to the example in which we have reason to believe that y is a function x and would like to know what the function is. We might proceed as follows.

Consider the set F of functions that can be expressed in standard mathematical notation. There are countably many such functions, which means they can be ordered in a way that is correlated with the natural numbers. Choose one such ordering, $f_1, f_2, \ldots, f_n, \ldots$ Using that ordering, we get some data and select the earliest function in the ordering that is compatible with the data. We then get more data and see whether the chosen function fits the additional data. If not, we again choose the earliest function on the list that is compatible with all the data we now have. If the right function is included in our list, F, this method will eventually arrive at that function, given enough time (Putnam, 1963).

There is the complication that we must allow for noise in the data. So, we must make some sort of trade-off, e.g., choosing that function f_n for which the sum of n and the mean squared-error on the data is least.

Our initial ordering of hypotheses might be based on how easy it is to use a hypothesis to answer questions in which we are interested. This would be one sort of simplicity ordering, where "simple" means simple to use.[5]

Ordering *sets* of hypotheses

Alternatively, we might rank all linear functions ahead of all quadratic functions, for example. In that case, we would not have a simple well-ordering of hypotheses, since infinitely many hypotheses would come between any linear hypothesis and any nonlinear quadratic hypothesis. Instead, we might use an ordering of cumulative *sets* of functions, $F_1, F_2, \ldots, F_n, \ldots$, where $F_1 \subset F_2 \subset F_3 \subset \ldots$

The idea is then first to consider the hypothesis in each of these sets that has the least squared error on the data. For each such hypothesis, we trade off the amount of the error for that hypothesis against the complexity of the earliest of the hypothesis classes to which it belongs.

We would presumably order these sets of functions so that later sets can accommodate more data points with no more error than earlier sets do. For example, given any two data points there is some linear hypothesis that fits them perfectly, so it is not interesting that some linear hypothesis fits two data points. However, it is not true that for any three data points there is a linear hypothesis that fits them perfectly. So, it is somewhat interesting if we have only three data points and some linear hypothesis fits them exactly or fairly closely. On the other hand, it is not interesting that some quadratic function captures three data points exactly, but perhaps mildly interesting that some quadratic function captures four data points.

Suppose more precisely that we choose classes of hypotheses $F_1 \subset F_2 \subset F_3 \subset \ldots$ in such a way that (1) for any N data points there is some hypothesis in F_{N+1} that exactly fits the

points, but (2) it is not generally the case that for any N data points there is some hypothesis in F_N that fits the points. Again we would have a simplicity ordering of sorts.

Using this measure of simplicity when we trade off simplicity against error might (at least under certain conditions) promise eventually to lead us to accept a hypothesis with a relatively low average error.[6]

Practical Reasons to Be Sensitive to Simplicity

Our inferential practices show a bias toward simpler hypotheses. We are considering the question whether we ought to continue to allow this bias.

Perhaps we can say that simpler theories are more likely to be true than more complicated theories that account for the same data. We can say this, anyway, if likelihood is epistemic likelihood, because considerations of simplicity are built into the procedures that determine what is likely in that sense.

On the other hand, there seems to be no direct noncircular argument that simpler theories are more likely to be true in the sense of objective likelihood. (But we will come back to this issue.)

We might be able to argue for an inductive bias that rests on one or the other of the two types of simplicity we have considered. The argument might be that that sort of inductive bias promises to help us eventually find the answer to questions we are interested in answering or that it promises to help us eventually to choose a hypothesis with as low an error rate as possible while having other practical advantages.

These are practical reasons to acquiesce in an inductive bias that favors simplicity, not theoretical reasons. But they do not reduce theoretical reasoning to practical

reasoning nor do they build wishful thinking into theoretical reasoning—at least not in the sense that *within* theoretical reasoning the desirability of a conclusion counts as a reason to believe it. There can be practical reasons for designing or acquiescing in a system that does not allow wishful thinking. A system defended in this way need not allow the internal use of practical considerations to decide between competing hypotheses.

Conservatism

Let us briefly look at an analogous issue concerning the role of conservatism in theoretical reasoning. Our reasoning is conservative in the sense that we start with our present view and try to improve it by getting rid of inconsistency and by increasing its coherence in ways that helps us answer questions in which we are interested (Rawls, 1971; Goodman, 1965; Peirce, 1931-58; Quine and Ullian, 1970; Dewey, 1938).

So, there is a further bias in reasoning beyond a simple inductive bias. This further bias favors beliefs that we already have over propositions that we do not already accept.

Again we can ask whether we should acquiesce in this further conservative bias? An alternative idea would restrict the conservative bias to certain "foundational" beliefs, such as beliefs about your most immediate experience and beliefs based on the recognition of self-evident truths (Descartes, 1637; Foley, 1987; Alston, 1989; Chisholm, 1982).

But there is a compelling argument for conservatism and against special foundationalism, namely, that special foundationalism leads inevitably to scepticism, and that again one will not be able to learn much of anything if one cannot rely on one's other

nonfoundational beliefs (Harman, forthcoming). Here is another practical reason to acquiesce in a certain way of doing theoretical reasoning.

Again, this sort of practical reason for an aspect of theoretical reasoning does not imply that theoretical reasoning itself is practical reasoning and it does not imply that theoretical reasoning involves wishful thinking because of its bias toward conservatism.

Nonpractical Interpretation

We can also interpret the reasons offered for simplicity and conservatism in a way does not treat them as practical reasons at all. Instead, we can think of these considerations as showing that a system without inductive bias in favor of simpler hypotheses or without a bias toward conservatism would not be able to reach conclusions that ought to be reached.

The point is that it would be deeply irrational to reason in a way that leads to skepticism. Any sceptical system simply gives the wrong results about what one should believe.

Someone might object that this line of thought is circular. I respond that it is not circular—it is reasoning! To think there is circularity here is to think of reasoning as the production of a logical argument with premises, intermediate steps, and a conclusion. For then it seems that the relevant premises must include the rationality of a system with inductive and conservative bias. But reasoning is not to be identified with the production of an argument. Reasoning is reasoned change in view in a way that improves coherence in a way that helps to answer questions in which you are interested.

Conclusion

My brief conclusion is that, although wishful thinking is not relevant in theoretical reasoning in the way that it is relevant in practical reasoning, certain aspects of theoretical reasoning can be given a practical defense. That defense does not mean that wishful thinking is allowed internally to theoretical reasoning. The defense can also be given a non-practical interpretation in terms of what conclusions ought to be reached.

References

Alston, W., (1989). *Epistemic Justification: Essays in the Theory of Knowledge*. Ithaca, NY: Cornell University Press.

Angluin, D.C., and Smith, (1983). Inductive inference: theory and methods. *Computing Surveys*, 15: 237-69.

Blum, L., and Blum, M., (1975). Toward a mathematical theory of inductive inference. *Information and Control*, 28: 125-55.

Blum, M., (1967). A machine-independent theory of the complexity of recursive functions. *Journal of the Association for Computing Machinery*, 14: 322-36.

Chisholm, R., (1982). *The Foundations of Knowing*. Minneapolis, Minnesota: University of Minnesota Press.

Descartes, R., (1637). *Discours de la méthode*.

Dewey, J., (1938). *Logic: The Theory of Inquiry*. New York: Holt.

Foley, Richard, (1987). *The Theory of Epistemic Rationality*. Cambridge, Massachusetts: Harvard University Press.

Gigerenzer, G. et al., (1999). *Simple Heuristics That Make Us Smart*. Oxford : Oxford University Press.

Gold, E. M., (1967). Language identification in the limit. *Information and Control*, 10:447-74.

Goodman, N., (1965). *Fact, Fiction, and Forecast*, 2nd edition. Indianapolis, Indiana: Bobbs-Merrill.

Harman, G., (1994). Simplicity as a pragmatic criterion for deciding what hypotheses to take seriously, in Douglas Stalker (ed.) *Grue: The New Riddle of Induction*. Peru, Ill.: Open Court, 1994), 153-71.

Harman, G., (forthcoming). Skepticism and foundations, to appear in *The Skeptics: Contemporary Essays*, edited by Steven Luper as a volume in the Ashgate Epistemology and Mind Series.

Hegel, G.W.F., (1812). *Wissenschaft der Logik*. Nürnberg: Schrag.

Kugel, P., (1977). Induction, pure and simple. *Information and Control*, 35: 276-336.

Ludlow, P., (1998). Simplicity and generative grammar. In R.Stainton and K.Murasugi (eds.), *Philosophy and Linguistics*. Boulder, Colorado: Westview Press.

Mach, E., (1960). *The Science of Mechanics*, 6th edition. Chicago: Open Court.

Mill, J.S., (1869). *A System of Logic, Ratiocinative and Inductive: Being a Connected View of the Principles of Evidence and the Methods of Scientific Investigation*, 7th edition. London: Longmans, Green, Reader, and Dyer.

Mitchell, T.M., (1997). *Machine Learning*. New York: McGraw-Hill.

Peirce, C.S., (1931-58). *Collected Papers of Charles Sanders Peirce*, 8 vols. Edited by C.Hartshorne, P.Weiss, and A.Burks. Cambridge: Harvard University Press.

Pollock, J., (1979). A plethora of epistemological theories. In G.Pappas (ed.), *Justification and Knowledge*. Dordrecht: Reidel: 93-114.

Putnam, H., (1963). "Degree of confirmation" and inductive logic. In A.Schillp (ed.), *The Philosophy of Rudolph Carnap*. LaSalle, Illinois: Open Court.

Quine, W.V., and Ullian, J.S., (1970). *The Web of Belief*. New York: Random House.

Quine, W.V.O., (1960). Carnap and logical truth. *Synthese* 12: 350-74.

Rawls, J., (1971). *A Theory of Justice*. Cambridge, Massachusetts: Harvard University Press.

Solomonoff, R.J., (1964). A formal theory of inductive inference. *Information and Control*, 7: 1-22, 224-54.

Simon, H., (1957). *Models of Man*. New York: Wiley.

Sober, E., (1975). *Simplicity*. Oxford: Oxford University Press.

Sober, E., (1988). *Reconstructing the Past*. Cambridge, Massachusetts: MIT Press.

Sober, E., (1990). Let's razor Occam's razor. In D.Knowles, ed., *Explanation and Its Limits*. Cambridge: Cambridge University Press.

Stalker, D., (1994). *Grue: The New Riddle of Induction*. Peru, Illinois: Open Court.

Turney, P.D., (1988). *Inductive inference and stability*. Ph.D. Dissertation, Department of Philosophy, University of Toronto.

Valient, L.G., (1979). The complexity of enumeration and reliability problems. *SIAM Journal of Computing*, 8: 410-21.

Notes

[1] David Lewis noticed that many people who live in Princeton believe these three things. When this is pointed out to them, most Princeton residents are amused but not motivated to correct their beliefs. (It turns out that Route 1 actually runs Northeast-Southwest when it is near Princeton, not North-South, and is anyway not really parallel to Nassau Street, which does run East-West. I have tried to explain this to my neighbors but they do not care.)

[2] Pollock (1979) distinguishes positive and negative coherence in this way. Simplicity might be included in positive coherence and not counted as a distinct factor.

[3] *Epistemic probability*, which has to do with what it is more or less reasonable for one to believe via theoretical reasoning, is to be distinguished from *objective probability*, which has to do with frequencies or propensities that one may or may not be aware of. Given a normal looking six sided die, one might reasonably suppose that any of the sides is equally likely to be topmost after the next roll of the die. That equal likelihood is a matter of epistemic probability. In fact, the die may be constructed in such a way that there is a strong propensity for the side labeled "4" to come up. The objective probability of 4 coming up may be 0.8 even though the epistemic probability of getting a 4 on the next roll is 1/6 or 0.17

[4] This obvious point is discussed in the context of machine learning in Mitchell 1997, Chapters 1-2. It is one moral of Goodman's (1965) "new riddle of induction," as discussed by philosophers. Stalker (1994) is a collection of essays on that riddle

[5] Ludlow (1998) defends this view of simplicity, citing earlier versions in Peirce (1931-1958) and Mach (1960). Harman (1994) argues for a similar view after surveying alternatives. Related computational approaches to simplicity are defended in Angluin and Smith (1983), Blum and Blum (1975), Blum (1967), Gold (1967), Kugel (1977), Solomonoff (1964), Turney (1988), and Valiant (1979). Sober (1975) argues for a "semantic" interpretation of simplicity. Later Sober (1988) and (1990) argues against the general relevance to inference of certain notions of simplicity.

[6] For further discussion, see e.g. Mitchell (1997, Chapter 7).