Stefan **Bold**, Benedikt **Löwe**,
Thoralf **Räsch**, Johan **van Benthem** (*eds.*)
Foundations of the Formal Sciences V
Infinite Games

Simplicity, Truth, and the Unending Game of Science

Kevin T. Kelly*

Carnegie Mellon University
Department of Philosophy
Baker Hall 135
Pittsburgh, PA 15213-3890, USA
kk3n@andrew.cmu.edu

ABSTRACT. This paper presents a new explanation of how preferring the simplest theory compatible with experience assists one in finding the true answer to a scientific question when the answers are theories or models. Inquiry is portrayed as an unending game between science and nature in which the scientist aims to converge to the true theory on the basis of accumulating information. Simplicity is a topological invariant reflecting sequences of theory choices that nature can force an arbitrary, convergent scientist to produce. It is demonstrated that among the methods that converge to the truth in an empirical problem, the ones that do so with a minimum number of reversals of opinion prior to convergence are *exactly* the ones that prefer simple theories. The approach explains not only simplicity tastes in model selection, but aspects of theory testing and the unwillingness of natural science to break symmetries without a reason.

1 Introduction

In natural science, one typically faces a situation in which several (or even infinitely many) available theories are compatible with experience. Standard

*I would like to thank Seth Casana, John Taylor, Joseph Ramsey, Richard Scheines, Oliver Schulte, Pieter Adriaans, Teddy Seidenfeld, and Balazs Gyenis for discussions and comments. Special thanks are due to the organizers of the Fifth International Conference on Foundations of the Formal Sciences for such an interesting interdisciplinary conference devoted to infinite game theory. The centrality of determinacy to the main results of this paper reflects the influence of some of the excellent presentations by other authors in this volume.

practice is to choose the *simplest* theory among them and to cite "Ockham's razor" as the excuse (Figure 1). "Simplicity" is understood in a variety

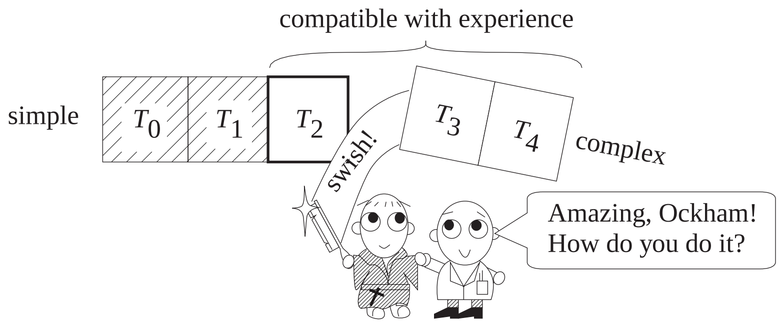

Figure 1. Ockham to the rescue.

of ways in different contexts. For example, simpler theories are supposed to posit fewer entities or causes (Ockham's original formulation), to have fewer adjustable parameters, to be more "unified" and "elegant", to posit more uniformity or symmetry in nature, to provide stronger explanations, or to be more strongly cross-tested by the available data. But in what sense is Ockham's razor truly an excuse? For if you already know that the simplest theory compatible with experience is true, you don't need any help from Ockham (Figure 2). And if you don't, then the true theory

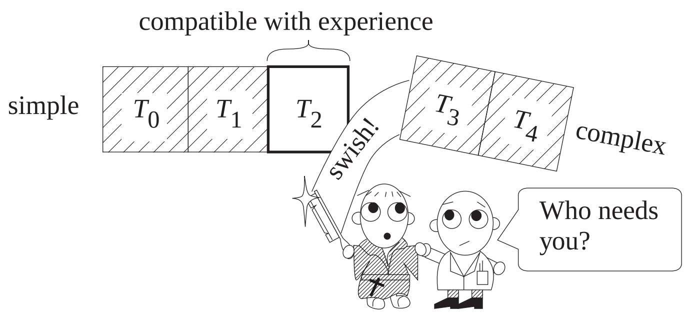

Figure 2. A Platonic Dilemma: Case I.

might be complex, so it is unclear how Ockham helps you find it (Figure 3). Indeed, how could a *fixed* bias toward simplicity indicate the possibly

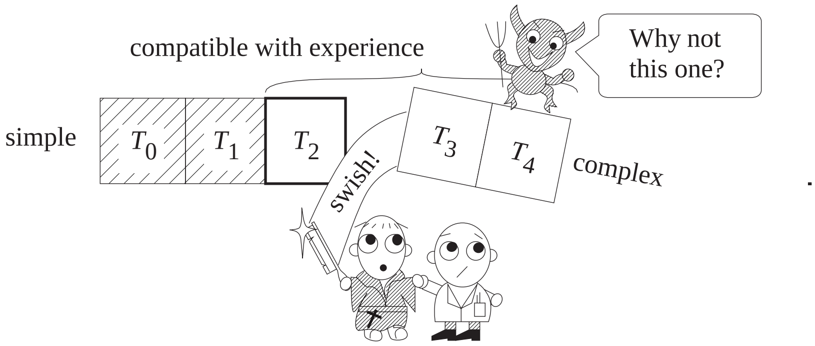

Figure 3. A Platonic Dilemma: Case II.

complex truth any better than a broken thermometer that always reads "zero" can indicate the temperature? You don't have to be a card-carrying skeptic to wonder what the tacit connection between simplicity and truth-finding could possibly be.

This essay explains the connection between simplicity and truth by modelling inquiry as an unending process, in which the scientist's aim is to converge to the truth in a way that minimizes, in a worst-case sense, reversals of opinion prior to convergence to the truth. Scientific methods may then be analyzed formally as strategies in an infinite game of perfect information, which brings to the subject powerful mathematical tools such as Donald Martin's Borel determinacy theorem [Mar$_1$75]. The proposed, long-run, strategic perspective on inquiry may appear abstract and remote from the day-to-day nuances of concrete scientific practice. Nonetheless, it is very general and singles out Ockham's razor as the best possible strategy to follow at every stage of inquiry, so its import for short-run practice is both sharp and concrete. Furthermore, the following review of standard attempts to link simplicity with theoretical truth in the short run reveals that they are all either irrelevant or based upon wishful thinking or circular arguments. A relevant, non-circular, long-run explanation may be better than no explanation at all.

2 Some Traditional Explanations of Ockham's Razor

Gottfried W. Leibniz explained the elusive connection between simplicity and truth by means of a direct appeal to the grace of God (Figure 4; [Lei14, §§ 55–59]) Since God is omnipotent and infinitely kind (to scientists, at least), it follows that the actual world is the most elegant (*i.e.*, simple)

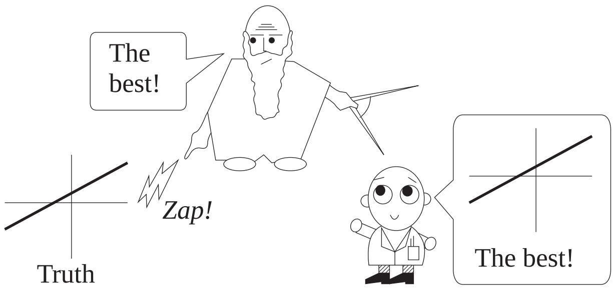

Figure 4. Gottfried W. Leibniz' Theological Explanation.

universe that could possibly produce such a rich array of effects. Hence, simplicity doesn't track the truth the way a thermometer tracks temperature; truth, by the grace of God, tracks simplicity. This explanation merely underscores the desperate nature of the question.

Immanuel Kant confronted the issue in his *Kritik der Urtheilskraft*.[1] According to Kant, the faculty of judgment must prescribe or presuppose that the diverse laws of nature may be unified under a small set of causes if nature is to be intelligible at all. But theories that involve a few extra causes are also intelligible, so intelligibility falls far short of explaining why one should prefer theories with fewer causes or entities over those that involve more.

Some latter-day philosophers have emphasized that simple theories have various "virtues", most notably, that simpler or more unified theories are more thoroughly tested or confirmed by a given evidence set (*e.g.*, [Pop68, p. 251–281], [Gly80], or [$\mathrm{Fri}_1$83, p. 236–250]). For if a theory has many free parameters (ways of being true), then new evidence simply "sets" the parameters and there is no risk of the theory, itself, being refuted altogether. But a simple theory does carry the risk of being refuted. It seems only fair to pin a medal of valor on the simple theory for surviving its self-imposed ordeal. The question, however, is truth, not valor, and the true theory might not be simple, in which case it wouldn't be valorous. To assume otherwise amounts to wishful thinking — the epistemic sin of concluding that the truth is as pleasing (intelligible, severely testable, explanatory, unified, uniform, symmetrical) as you would like it to be. Rudolf Carnap sought uniformity of nature in logic itself [$\mathrm{Car}_0$50, p. 562–567]. This "logic" amounts, however, to nothing more than the imposition of greater prior probabilities on more

[1] *Cf.* the *Akademie-Ausgabe* [NatWin08, p. 185].

uniform worlds, where uniformity is judged with respect to an arbitrarily selected collection of predicates. The argument goes like this (Figure 5). Suppose there are but two predicates, "green" and "blue" and that every-

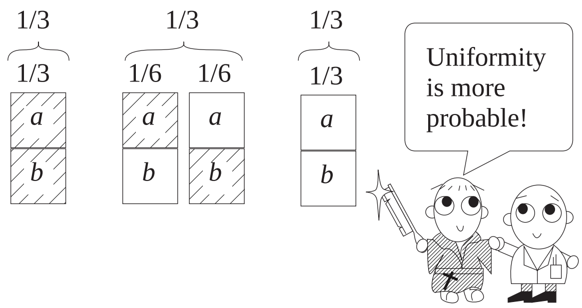

Figure 5. Rudolf Carnap's "Logical" Explanation.

thing is either green or blue. Suppose there are two observable objects, a and b. Two worlds are *isomorphic* just in case a one-to-one substitution of names takes you from one world to the other in a way that preserves the basic predicates in your language. Hence the uniform world in which a and b are both green is in its own isomorphism class, as is the uniform world in which a and b are both blue. The two non-uniform worlds in which a and b have different colors can each be reached from the other by a one-to-one, color-preserving substitution of names, so they end up in the same isomorphism class. Now Carnap invokes the principle of indifference to put equal probabilities of one third on each of these three automorphism classes and invokes it again to split the one third probability on the non-uniform class over the two non-uniform worlds. The resulting probability distribution is then biased so that uniform worlds get probability one third and non-uniform worlds get probability one sixth. So uniform worlds are more probable than non-uniform worlds (by a factor of two in this tiny example, but the advantage increases as observable individuals are added).

Nelson Goodman objected that whatever is logical ought to be preserved under translation and that Carnap's uniformity bias based on linguistic syntax isn't [Goo83, p. 59–83]. For uniformly green and uniformly blue experience are uniform. But one can translate green and blue into "grue" and "bleen", where "grue" means "green if a and blue if b" and "bleen" means "blue if a and green if b" (Figure 6). Then in the grue/bleen language, the worlds that used to be non-uniform are now uniformly grue or uniformly bleen, respectively and the worlds that used to be uniform are non-uniform,

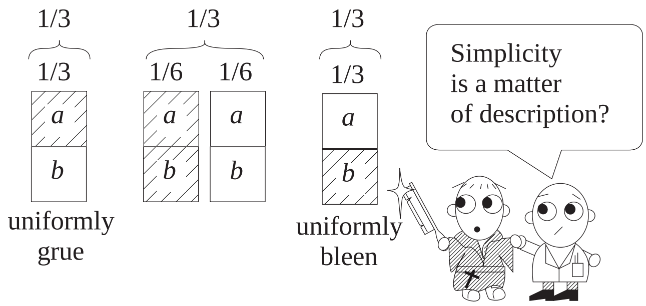

Figure 6. Nelson Goodman's "Grue" Argument.

for "green" means "grue if a and bleen if b" and "blue" means "bleen if a and grue if b". Since logical inferences are based entirely on syntax and syntactically the situation between green/blue and grue/bleen is entirely symmetrical, uniformity cannot be a feature of logical syntax. The moral is that Carnap's story makes uniformity of nature a mere matter of description. But a guide to truth could not be a mere matter of description, since truth doesn't depend upon how it is described.

3 Statistical Explanations

So much for philosophy. Surely, the growing army of contemporary statisticians, machine learning researchers, and industrial "data miners" must have a better explanation based on rigorous, mathematical reasoning. Let's check. A major player in the scientific methodology business today is *Bayesian* methodology. The basic idea is to allow personal biases to enter into statistical inferences, where personal bias is represented as a "prior" probability measure over possibilities. The prior probability of hypothesis H is then combined with experience E_t available at t via *Bayes' theorem* to produce an updated probability of H at t', which represents your updated opinion concerning H:

$$P_{t+1}(H) = P_t(H \mid E_t) = \frac{P_t(H) \cdot P_t(E_t \mid H)}{P_t(E_t).}$$

It is clear from the formula that your prior opinion $P_t(H)$ is a factor in your posterior opinion $P_{t+1}(H)$, so that the simplest theory compatible with the new data ends up being most probable in the updated probabilities. Ockham's razor is just a systematic bias toward simpler theories. So to explain

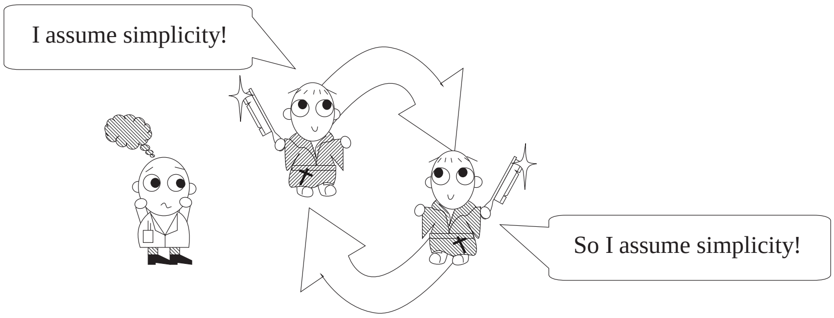

Figure 7. The Circular Bayesian Explanation.

its efficacy by appealing to a prior bias toward simplicity is patently circular (Figure 7).

Bayesians also have a more subtle story. Yes, it begs the question simply to impose a prior bias toward simple theories, so let's be "fair" and impose *equal* prior probabilities on competing theories, be they simple or complex. Now suppose, for concreteness, that we have just two theories, simple theory S and complex theory $C(\theta)$ with free parameter θ which (again for concreteness) can be set to any value from 1 to k (Figure 8). Suppose, further, that

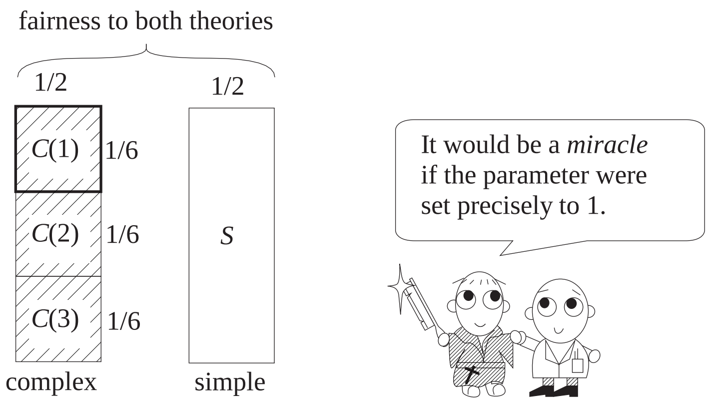

Figure 8. The Miracle Explanation.

S consistently entails E_t, as does $C(1)$, but that for all other parameter values i, $C(i)$ is refuted by E_t. Thus, $P_t(E_t \mid S) = P_t(E_t \mid C(1)) = 1$

but for all i distinct from 1, $P_t(E_t \mid C(i)) = 0$. Suppose, again, that you have no prior idea which parameter value of $C(i)$ would be the case if $C(\theta)$ were true (that's what it means for the parameter to be "free"). So $P_t(\theta \mid C(i))$ is uniform.[2] Turning the crank on Bayes' theorem, one obtains $P_t(S \mid E_t)/P_t(C \mid E_t) = k$. So even though the complex theory could save the data just as well as the simple one, the simple theory that does so without any *ad hoc* fiddling ends up being "confirmed" much more sharply by the same data E_t (*e.g.*, [Ros83, p. 74–75]). Surely that explains how severe testability is a mark of truth, for doesn't the more testable theory end up more probable after a fair contest?

One must exercise caution when Bayesians speak of fairness, however, for probabilistic "fairness" between "blue" and "non-blue" implies a marked bias toward "blue" in a choice among "blue, yellow, red". That is all the more true in the present case: "fairness" between S and C induces a strong bias for S with respect to $C(1), \ldots, C(k)$. One could just as well insist upon "fairness" at the level of parameter settings rather than at the level of theories (Figure 9). In that case, one would have to impose equal probabilities

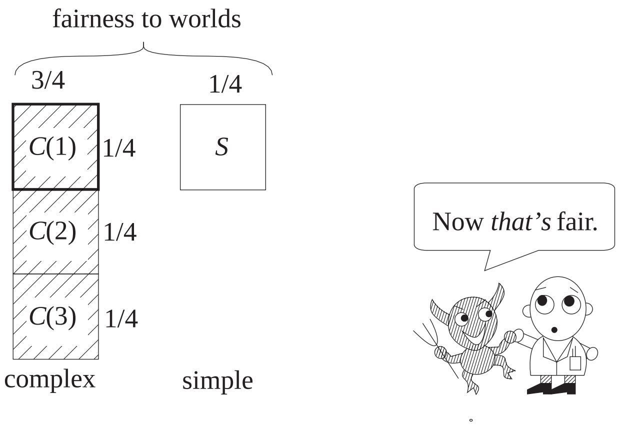

Figure 9. The Miracle Reversed.

of $1/(k+1)$ over the $k+1$ possibilities $\{S, C(1), \ldots, C(k)\}$. Now $C(0)$ and, hence, C, will remain forever at least as a probable as S in light of evidence

[2]This is a discrete version of the typical restrictions on prior probability in Bayesian model selection (*cf.* [Was04, p. 220–221]). If the parameters are continuous, each parameter setting receives zero prior probability, but the result is the same because the likelihood of the more complex theory must be integrated over a higher-dimensional space than that of the simpler theory.

agreeing with S. Classical statisticians explain Ockham's razor in terms of "overfitting" (*cf.* [Was04, p. 218–225] for a textbook review). "Overfitting" occurs when you want to estimate a sampling distribution by setting the free parameters in some statistical model. In that case, the expected squared predictive error of the estimated model will be higher if the model employed is too complex (*e.g.*, [For$_1$Sob94]). This is a kind of objective, short-run connection between simplicity and truth-finding, but it doesn't really address the question at hand, which is how Ockham's razor helps you find the true *theory*, which is quite another matter from which theory or model to use to estimate the underlying sampling distribution. The quickest way to see why is this: suppose that God were to tell you that the true model has fifty free parameters. On a small sample, the overfitting argument would still urge you to use a much smaller model for estimation and prediction purposes (Figure 10). So the argument couldn't be concerned with finding the true theory.

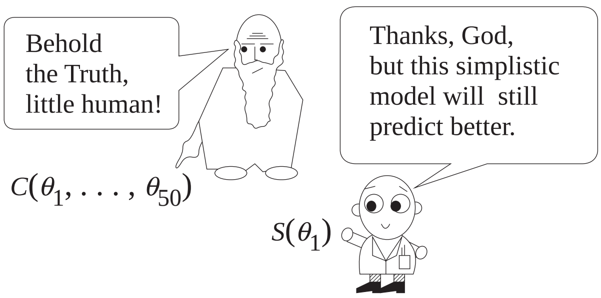

Figure 10. The "Overfitting" Explanation.

More subtly, the sense of approximate error employed in the overfitting literature is wrong for theory selection. Getting "close" to the underlying sampling distribution might not get you "close" to the form of the true model, since distributions arbitrarily close to the true distribution could be generated by models arbitrarily far from the true model.[3] Thus, distance from the theoretical truth is typically discontinuous with distance from the true sampling distribution, so minimizing the latter distance may fail to get you close in the former, as in the case of God informing you that the true model is very complex. Another point about overfitting is that even to the

[3] This is particularly true when the features of the model have counterfactual import beyond prediction of the actual sampling distribution, as in causal inference [GlySch$_1$Spi00, p. 47–53].

extent that it does explain the role of simplicity in statistical prediction, it is tied, essentially, to inference problems in which the data are stochastically generated, leaving one to wonder why simplicity should have any role in deterministic inference problems, where it still feels like a good idea.

Finally, there are theorems of the sort that some method equipped with a prior bias toward simplicity is guaranteed to *converge* in some sense to the true model as experience (or sample size) increases (*e.g.*, [LohZhe95]). That would indeed link Ockham's razor with truth-finding *if* it could be shown that *other* possible biases don't converge to the truth. But they do. The logic of convergence results is not that Ockham's advice points at or indicates the truth, but that it is "washed out" or "swamped", eventually, by accumulating experience, even if the advice is so misleading as to throw you off the track for a long time (Figure 11). But alternative biases would

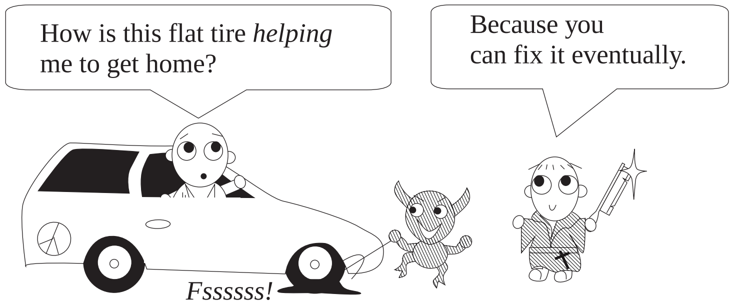

Figure 11. The Convergence Explanation.

also be washed out by experience eventually[4] so that's hardly a ringing endorsement of Ockham's razor. What is required is an argument that Ockham's razor is, in some sense, the *best possible* bias for finding the true theory.

4 Post Mortem

To recapitulate, the standard explanations of the mysterious relationship between simplicity and theoretical truth are either circular, wishful, or irrelevant. Still, they provide useful information about how the relationship can't be explained. For indication of the truth is too strong a connection to establish without begging the question at the outset, as Leibniz and the Bayesians do. On the other hand, mere convergence in the limit is too weak

[4] *Cf.* [Hal74, p. 212, Theorem A]. Also, see the critical discussion of Bayesian convergence theorems in [Kel96, p. 302–330].

to single out simplicity as the right bias. The crux of the puzzle, then, is to come up with a notion of "helping to find the truth" that is strong enough to single out simplicity as the right bias to have but that is not so strong as to demand a question-begging appeal to Ockham's razor at the outset in order to establish it. Thus, the appropriate notion of "help" must be stronger than convergence in the limit and weaker than indication in the short run.

The account developed below steers between these two extremes by considering a refined concept of convergence, namely, convergence with a minimum number of reversals of opinion prior to arrival at the goal (Figure 12). This is stronger than mere convergence in the limit, which says nothing

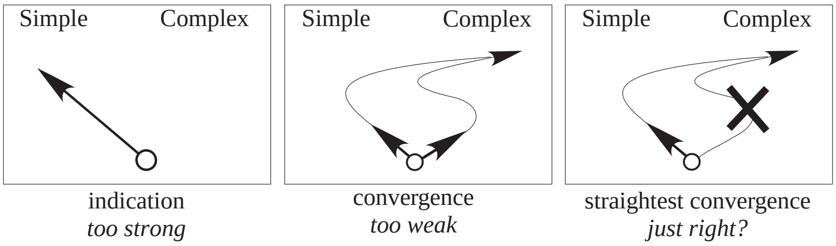

Figure 12. Three Kinds of "Help".

about minimizing reversals of opinion along the path to the truth, and is weaker than indication, which allows for no reversals of opinion whatever. It will be demonstrated that an ongoing bias toward simplicity minimizes kinks in your course to the truth in a certain precise sense. But first, I illustrate the general flavor of the approach by showing that something similar happens almost every time you ask for directions.

5 Asking for Directions

Suppose you are headed home on a road trip and get lost in a small town. In frustration, you stop to ask a local resident how to get home (Figure 13). Before you can even say where you are headed, he gives you the usual sort of advice: directions to the nearby freeway entrance ramp, which happens to be a few blocks back toward where you just came from. Now suppose that, in a fit of hubris, you disregard the resident's advice in favor of some intuitive feeling that your home is straight ahead (Figure 14). That ends up being a bad idea (Figure 15). You leave town on a small rural route that winds its wild way over the mountains. At some point, you concede the error of your ways and turn around to follow the resident's directions to the

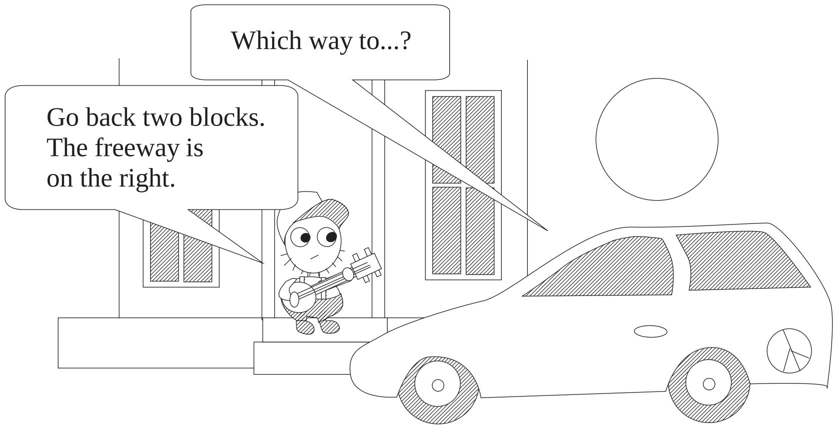

Figure 13. Asking For Directions.

Figure 14. Hubris!

freeway. The freeway then follows as straight a route home as is practicable through mountainous terrain. As you speed your way homeward, you have ample time to regret: *if you hadn't ignored the local resident's advice, you wouldn't have added that useless, initial U-turn to your otherwise optimal journey home.* Let's take stock of a few striking features of this mundane tale. First, the local resident's advice was indeed *helpful*, since it would have put you on the straightest possible path home. Second, *by disregarding the advice, you incurred an extra U-turn or kink in your route.* What is particularly vexing about the initial U-turn is that it occurs even before you properly begin your journey. It's a sort of navigational "original sin"

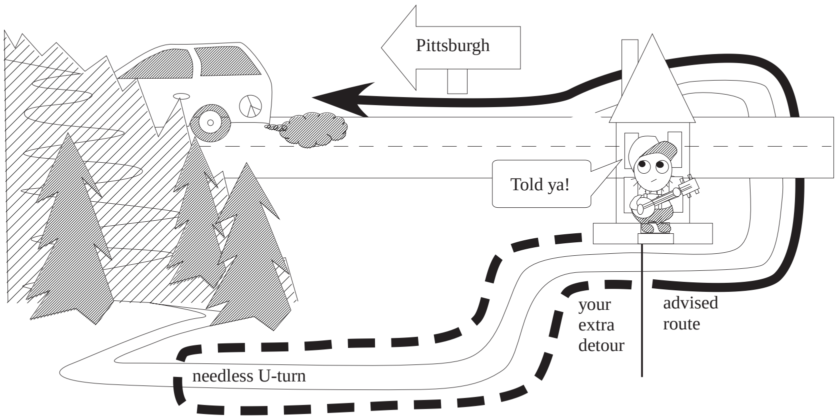

Figure 15. The U-turn Argument.

that you can never be absolved of. Third, the resident *didn't need to know where you were going in advance* in order to give you helpful advice. Any stranger asking for directions in a small, isolated town would do best to get on the freeway. Hence, the resident's ability to provide useful information without knowing where your home is *doesn't require an occult or circular explanation*. Suppose, on the other hand, that the resident could give you a compass course home before knowing where you are headed. That *would* require either a circular or an occult explanation (an Ouija board or divining rod). Fourth, even the freeway is not perfectly straight, so the resident's advice provides *no guarantee against future course reversals*, even if it is the best possible advice. Finally, the resident's advice is the best possible advice *even though it points you away from your goal* at the outset. If help required that you be aimed in the right direction, then the resident would have to give you a compass heading home, which wouldn't be possible unless he already knew where your goal was or had an Ouija board or divining rod.

So the typical situation in which you ask for directions home from a distant, small town has all the fundamental features that an adequate explanation of the truth-finding efficacy of Ockham's razor must have. Perhaps Ockham also provides fixed advice that puts you on the best possible route to the truth without pointing you at the truth and without guarantees against future course reversals along the way. It remains, then, to explain what the freeway to the truth is and how Ockham's advice leads you to it.

6 The Freeway to the Truth

Even lexicography suggests an essential connection between freeways and truth-finding, for both changes in course and changes in opinion are called changes in *attitude*. According to this analogy, Ockham's advice should somehow minimize changes of opinion prior to convergence to the right answer.[5] Let's consider how the story goes in the case of a very simple truth-finding problem.

Suppose that there is an emitter of discrete, readily detectable particles at arbitrary intervals and that you know that it can emit at most finitely many particles altogether (Figure 16). The question is how many particles

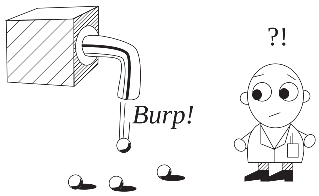

Figure 16. Counting Particles (Effects).

it will ever emit. What makes the problem interesting is that an arbitrarily long interval without new particles can easily be mistaken for total absence of future particles. This problem has more general significance than might be apparent at first, for think of the particles as detectable *effects* that are arbitrarily hard to detect as parameters in the true theory are tuned toward zero. For example, in curve fitting, the curvature of a quadratic curve may be so slight that it requires a huge number of data to notice that the curve is non-linear.[6] So the theory that the curve is quadratic but

[5]The idea of counting mind-changes already appears in [Put65]. Since then, the idea has been studied extensively by computer scientists interested in computational learning (*cf.* [Jai+99] for a review). The focus, however, is on categorizing the complexities of problems rather than on singling out Ockham's razor as an optimal method. Oliver Schulte and I began looking at retraction minimization as a way to severely constrain one's choice of hypothesis in the short run in 1996 (*cf.* [Sch$_9$99a, Sch$_9$99b]). Schulte has also applied the idea to the inference of conservation laws in particle physics [Sch$_9$01]. The ideas in this essay build upon and substantially simplify and generalize the initial approach taken in [Kel02, GlyKel04, Kel04]). While the present manuscript was in press, the approach was developed further in [Kel07]. Some differences between the two approaches are mentioned in subsequent footnotes.

[6]It is assumed that he data are increasingly precise but inexact; else three points would settle the question [Pop68, p. 131–131]. The same point holds if the data are noisy. In that case, tuning the parameters toward zero makes the effects statistically undetectable at small sample sizes (*cf.* [GlyKel04, Kel04] for an application of the preceding ideas to

not linear predicts the eventual detection of effects that would never appear under the linear theory. Similarly, the curvature of a cubic curve may be so slight that it is arbitrarily hard to distinguish from a quadratic curve. The point generalizes to typical model selection settings regardless of the interpretation of the parameters. So deciding among models or theories with different free parameters is quite similar, after all, to counting particles.

The traditional formulation of "Ockham's razor" is to *not multiply entities without necessity.* It is "necessary" (on pain of outright inconsistency) to assume as many particles as you have seen, but it is not necessary to assume more, so that *if* you conclude anything, you should conclude exactly as many particles as you have seen so far (Figure 17). The most aggressive

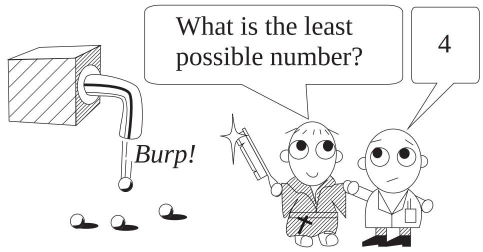

Figure 17. Ockham in Action.

Ockham method is the *counting method* that concludes that every particle has been seen at every stage. More realistic Ockham methods refuse to commit themselves to any answer at all until a long time has passed with no novel effects. Ockham's razor, itself, says nothing about how long this "confidence-building" time should last and the following argument for Ockham's razor doesn't imply anything about how long it should be either; it simply requires you to adopt some Ockham method, whether the method waits or not. That is as it should be, since even believers in short-run evidential support (*e.g.*, Rudolf Carnap and the Bayesians) allow for arbitrary individual differences concerning the time required for confidence buildup.

Other intuitive formulations of empirical simplicity conspire with the view that the Ockham answer should be the exact count. First, the Ockham theory that there are no more particles than you have seen is the most *uniform* theory compatible with experience, for it posits a uniformly particle-free future. Second, the Ockham theory is the most *testable* theory compatible with experience, for if it is false, you will see another particle and it will

stochastic problems.

be decisively refuted. Any theory that anticipates more particles than have been seen might be false, because there are fewer particles than anticipated, in which case it will never be refuted decisively, since the anticipated particles might always appear later. Third, the Ockham theory is most explanatory, since the theory that posits extra particles fails to explain the *times* at which those particle appear. The theory that there are no more particles fails to posit extra, unexplained times of appearance. Fourth, the Ockham theory is most *symmetrical*, since the particle-free future is preserved under permutation of times, whereas a future punctuated by new particle appearances would be altered by such permutations. Fifth, the Ockham theory has the *fewest free parameters*, because each time of appearance of a new particle is a free parameter in a theory that posits extra particles. So in spite of its apparent triviality, the problem of counting things that are emitted from a box does illustrate a wide range of intuitive aspects of empirical simplicity. That isn't so surprising in light of the analogy between particles and empirical effects tied to free parameters.

If you follow an Ockham solution to the particle-counting problem, then you change your mind in light of increasing data at most once per particle. If the true count is k, then you change your mind at most k times. By way of comparison, suppose that you have a hankering to violate Ockham's razor by producing a different answer (Figure 18). You might reason as follows.

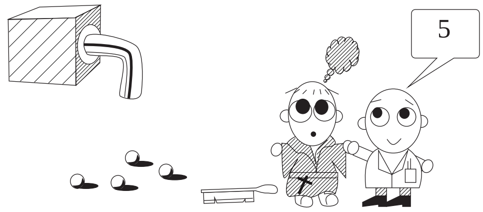

Figure 18. Ockham Violation!

The particle emitter has overturned every successive Ockham answer in the past (*i.e.*, "zero", "one", "two", and "three"), so you expect it will overturn the current Ockham answer "four" as well. So by induction on Ockham's unbroken losing streak in the past, you anticipate failure again and guess "five" (or some greater number of your choosing) rather than the Ockham answer "four". Philosophers of science call this the "negative induction from the history of science" [Lau81]. Why side with Ockham, rather than with the negative induction against him?

Efficiency is future-directed. Slush funds or debts may have been accumulated in the past, but efficiency optimization in the present concerns future costs incurred by future acts over which you have some control. So think of inquiry as starting from scratch at each moment. Accordingly, the *subproblem* entered at a given stage of inquiry consists of the restriction of possibilities to those consistent with current experience and only mind-changes incurred after entering the subproblem are counted in that subproblem.[7] Consider the subproblem entered when you first say "five", upon having seen only four particles. There is no deadline by which the fifth particle you anticipate has to show up, so you may have to wait a long time for it, even if you are right. You wait and wait (Figure 19). Your graduate

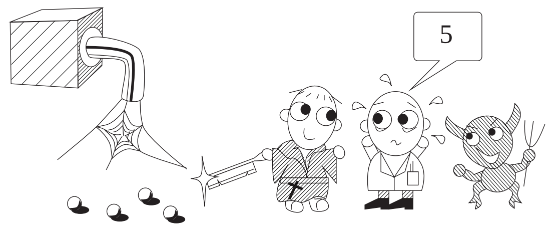

Figure 19. The Pressure Builds.

students exhaustively examine the particle emitter for possible malfunctions. Colleagues start talking about the accumulation of "null results" and discuss the "anomalous" failure of the anticipated marble to appear. True, the posited particle could appear (to your everlasting fame) at any time, so your theory isn't strictly refuted. Nonetheless, you feel increasing pressure to switch to the four-particle theory as the anomaly evolves into a full-blown "crisis". This increasing pressure comes not from the "weight of the evidence", as philosophers are wont to say, but from your strategic aim to converge to the truth, regardless of what it happens to be. For if you never change your mind from "five" to "four" and the fifth particle never appears, you will converge for eternity to "five" when the truth is "four". So at some time of your choosing, you must (on pain of converging to the wrong answer) cave in to the pressure from nature's strategic threat and switch back to the (Ockham) theory that the machine will produce just four particles (Figure 20). Won't that make for interesting gossip in the

[7]One might object that a sub-problem should hold past information and past theory choices fixed and sum the total cost from the outset of inquiry. That approach is developed in detail in [Kel07].

Particle Counting Association, where you are feted as the sole defender of the five particle theory?[8]

To summarize, in the subproblem entered when you first say "five", nature can force you to change your mind at least once (from "five" to "four"), in the manner just described, without presenting a fourth particle. The same is not true of Ockham, who enters the same subproblem saying "four" (or nothing at all) and who never changes his mind until the next particle appears. Thereafter, Ockham changes his mind exactly one time per extra particle. But you can be forced by nature to change your mind at *least* once per extra particle (on pain of not converging to the truth) in the same manner already described; for a long period during which there are exactly i particles forces you to say "i" on pain of not converging to the truth, after which nature can present the $i+1$st particle, *etc.* (Figure 21).

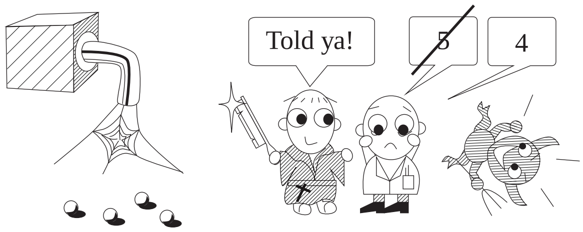

Figure 20. The Agony of Retreat.

Hence, *if a solution violates Ockham's razor in the particle counting problem, then in the subproblem entered at the time of the violation, whatever sequence of outputs the Ockham solution produces, the violator can be forced by nature to produce a sequence including at least the same mind-changes plus another one (the initial U-turn from "five" to "four").* You should have listened to Ockham!

[8]I am alluding, of course, to Thomas Kuhn's 1962 celebrated historical theory of the structure of scientific revolutions [Kuh$_1$62]. Null experiments generate anomalies which evolve after careful consideration into crises that ultimately result in paradigm change. Kuhn concludes, hastily, that the change is an unlawful matter of politics that has little to do with finding the truth. I respond that it is a necessary consequence of the logic of *efficient* convergence to the truth after a violation of Ockham's razor, as will become clear in what follows. Many of the celebrated scientific revolutions in physics have been the results of Ockham violations (*e.g.*, Ptolemy vs. Copernicus, Fresnel vs. Newton, Newton vs. Einstein, and creationism vs. Darwin). In each of these cases, a theory positing extra free parameters (with attendant empirical effects) was chosen first and a simpler theory was thought of later and came to replace the former, often after an accumulation of null experiments.

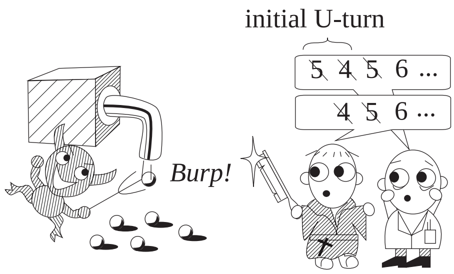

Figure 21. Ockham Avoids Your Initial U-turn.

The same argument works if you violate Ockham's razor in the other possible way, by saying "three" when four particles have been seen. For nature can refuse to present more particles until you change your mind to "four" on pain of never converging to the right answer if the right answer is "four". But in the same subproblem, Ockham would already have said "four" if he said anything at all and, in either case, you can be forced into an extra mind-change in each answer. So the U-turn argument also explains the need for maintaining consistency with the data.

So there is, after all, a close analogy between particle counting and getting on the freeway. Your initial mind change from "five" to "four" is analogous to your initial U-turn back to the local resident's house *en route* to the highway. Thereafter, no matter what the true answer is, you can be forced to change your mind at least once for each succssive particle, whereas Ockham changes his mind at most once per successive particle. These mind-changes are analogous to the unavoidable curves and bends in the freeway. So no matter what the truth is, you start with a U-turn Ockham avoids and can be forced into every mind-change Ockham performs thereafter. As in the freeway example, you have botched the job before you even properly get started. In both stories, the advice is the best possible. Nonetheless, it does not impose a bound on future course reversals; nor does it point you toward your goal by some occult, unexplained mechanism.

A striking feature of the explanation is that it is entirely game-theoretic. There is no primitive notion of "support" or "confirmation" by data of the sort that characterizes much of the philosophical literature on induction and theory choice (Figure 22).[9] Nor are there prior probabilities that foster the illusion of "support" of general conclusions by a few observations. The

[9]In this respect, my approach is a generalization and justification of the "anti-inductivism" of Karl Popper [Pop68, p. 27–30].

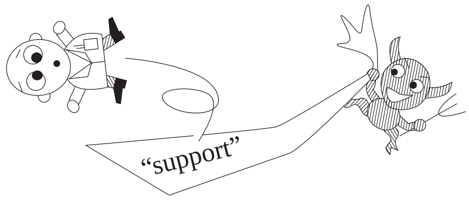

Figure 22. Pulling the Rug Out.

phenomenology of "support" by evidence emerges entirely from the aim of winning this truth-finding game against nature. Furthermore, the game is essentially infinite. For if there were an *a priori* bound on the time by which the next particle would arrive if it arrives at all, then you could simply "outwait" nature and avoid changing your mind altogether. So the argument is situated squarely within the theory of infinite, zero-sum games, which is the topic of this volume.

Here is why the reference to subproblems is essential to the U-turn argument. Suppose that you are asleep when you see the first particle and that when you see the second particle you wake up and guess "three", expecting that you will also sleep through the third (Figure 23). Thereafter, you

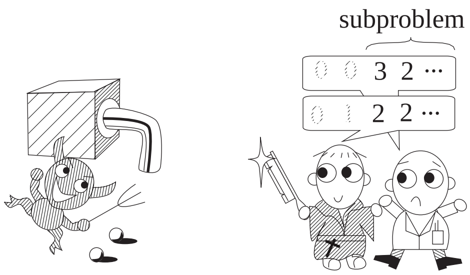

Figure 23. Ockham Still Wins in Subproblem.

always agree with Ockham. If the third particle doesn't appear right away, you can be forced to change your mind to "two", but that's only your second retraction — Ockham wouldn't have done better. Now that you have "caught up" with Ockham, you match him no matter how many particles you see in the future. But that is only because you "saved" a retraction in the past by sleeping through the first particle. That is like hoarding a "slush fund" to hide future mismanagement from the public. In the subproblem entered when you say "three", the slush fund is emptied and you have to demonstrate your efficiency from scratch. In that subproblem, your first reversal of opinion back to "two" gets added to all your later mind-changes and you never catch up, so Ockham wins. The moral: an arbitrary Ockham solution beats you in the subproblem in which you violate Ockham's razor, but the Ockham solution does as well as you in every subproblem, so the Ockham solution is better. In the best case for the violator, the anticipated fifth particle might appear immediately after the violation, before the method even has a chance to become queasy about over-estimating (Figure 24). In that case, the violator's output sequence in the subprob-

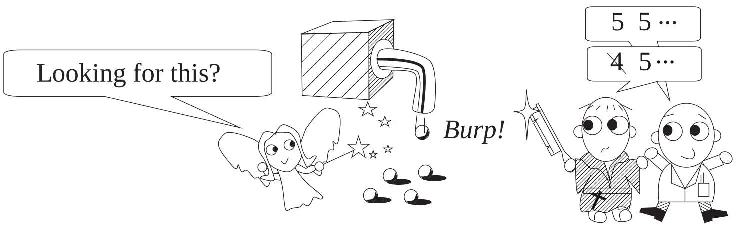

Figure 24. Best Case Fairy to the Rescue.

lem entered at the violation begins with "five", "five", whereas Ockham's output sequence in the same subproblem begins with "four", "five", which is worse. Hence, the Ockham method doesn't weakly dominate the violator's mind-changes in the subproblem in question. But that merely explains why the U-turn argument, which does establish the superiority of an arbitrary Ockham solution over an arbitrary non-Ockham solution, is not a weak dominance argument. The U-turn argument essentially involves a worst-case dimension lacking in weak dominance, for nature can *force* the non-Ockham solution from "five" back to "four" (on pain of convergence to the wrong answer) by withholding particles long enough and can then reveal another particle to make it say "five", "four", "five", which is never produced by any Ockham method and which properly extends the Ockham sequence "four", "five".[10] Nor, for that matter, is the U-turn argument a

[10]Indeed, Ockham methods are weakly dominated by methods that hang on to their original count for an arbitrarily long period of time (Teddy Seidenfeld, personal communication). That isn't so bad, after all, because there are compelling reasons not to under-count (*e.g.*, the undercount couldn't possibly be true). The crux of Ockham's razor is to motivate not over-counting, and over-counters do not dominate the retractions of Ockham methods in this way. More to the point, the alternative theory developed in [Kel07] is not subject to this objection, because tardiness of retractions is also penalized,

standard worst-case or "minimax" argument, for there is no fixed bound on mind-changes for any solution to the counting problem (nature can force an arbitrary solution through any number of mind-changes).

7 A General Conception of Scientific Problems

A *scientific problem* specifies a set Ω of possible worlds the scientist must succeed in together with a question $\mathcal{Q}$ which partitions Ω into mutually exclusive *potential answers*. The aim is to find the true answer for w no matter which world w in Ω you happen to live in. If $\mathcal{Q}$ is a binary partition, one thinks of a decision or test problem for one of the two cells vs. the other. If it is a fixed range of alternatives extensionally laid out in advance, one speaks of *theory choice*. If it is an infinite partition latently specified by some criterion determining the kind of theory that would count as success, the situation might be described as *discovering* the truth.

The most characteristic thing about empirical science is that you don't get to see w in its entirety. Instead, you get some incomplete evidence or information about w, represented by some subset of Ω containing w. The set of all possible information states you might find yourself in is modelled as the collection of open sets $\mathcal{V}$ in a topological space over Ω. A *scientific problem* is just a triple $(\Omega, \mathcal{V}, \mathcal{Q})$, where $(\Omega, \mathcal{V})$ is a topological space and $\mathcal{Q}$ partitions Ω (Figure 25). The idea is that, although the scientist never

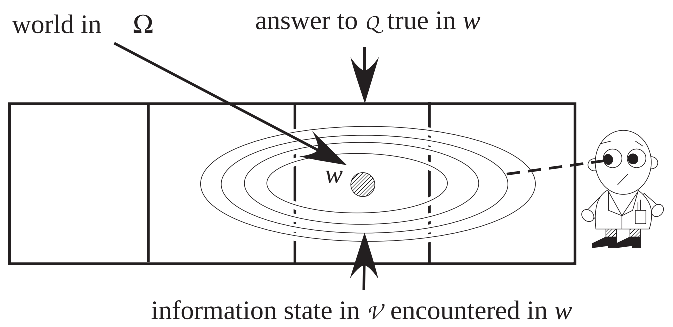

Figure 25. A Scientific Problem.

gets to see the actual world w, itself, he does get to see ever smaller open neighborhoods of w.

so the delayers do not end up ahead. The present theory has the advantage of greater mathematical elegance, however.

The use of topology to model information states is not a mere stipulation, for information concerns verifiable effects and topology is perhaps best understood as the mathematical theory of verifiability.[11] The point is seen most directly as follows. Identify each proposition with the set of possible worlds or circumstances in which it would be true, so propositions may be modelled as subsets of the set Ω of possible worlds. Say that a proposition is *verifiable* if and only if there exists a method or procedure that examines experience and that eventually illuminates a light if the proposition is true and that never illuminates the light otherwise. For example, illuminating the light when a particle appears yields a verification procedure for the proposition that at least one particle will appear.

The contradiction is the empty set of worlds (it can't possibly be true) (Figure 26). It is verifiable by the trivial verification procedure that never

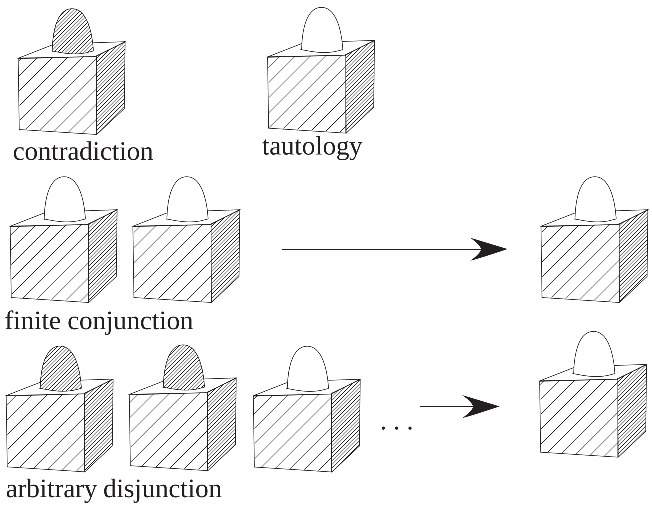

Figure 26. Verifiable Propositions are Open Sets.

illuminates its light. Similarly, the tautologous proposition consists of the whole set Ω of worlds and is verifiable by the trivial procedure that turns on its light *a priori*. Suppose that two verifiable propositions A, B are given. Their conjunction $A \cap B$ is verifiable by the procedure that turns on its light if and only if the respective verification procedures for A and for B have both turned on their lights. Finally, suppose a collection $\mathcal{D} \subset \Omega$

[11]Topology is also used to model partial information states in denotational semantics [$\mathrm{Sco}_1$82].

of verifiable propositions is given. Their disjunction $\bigcup \mathcal{D}$ is verifiable by the procedure that turns on its light just in case the procedure for some proposition $A \in \mathcal{D}$ turns on its light (you will see that light eventually as long as each respective procedure is only a finite distance away). Hence, the verifiable propositions $\mathcal{V}$ over Ω constitute the open sets of a topological space $(\Omega, \mathcal{V})$. So every theorem about open sets in a topological space is also true of ideal empirical verifiability. One of the most characteristic features of topology is that open sets are closed under arbitrary union but only under finite intersection. That is also explainable in terms of verifiability. Suppose you are given an infinite collection $\mathcal{C}$ of verifiable propositions. Is there a verification procedure for $\bigcap \mathcal{C}$? Not always. For the respective verification procedures for the elements of $\mathcal{C}$ may all turn on their lights, but at different times, so that there is no time by which you can be sure that it is safe to to turn on your light for $\bigcap \mathcal{C}$ (Figure 27). That is an instance

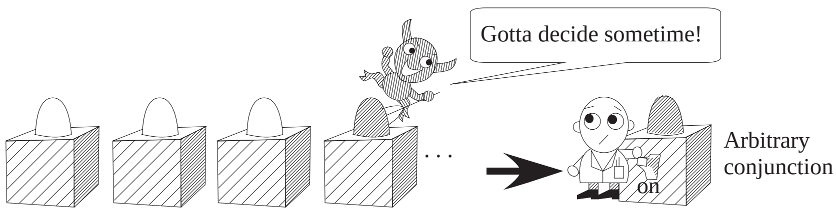

Figure 27. The Demon of Arbitrary Conjunction.

of the classical problem of induction: no matter how many lights you have seen go on, the next light might never do so. So not only are the axioms of topology satisfied by empirical verifiability; the characteristic asymmetry in the axioms reflects the problem of induction.

In a given topological space $(\Omega, \mathcal{V})$, the problem of induction arises in a world $w \in \Omega$ with respect to proposition H just in case every information state (open proposition) true in w is compatible both with H and with $\neg H$ (Figure 28).[12] In standard, topological parlance, the problem of induction arises with respect to H in w just in case w is a *boundary point* of H. So the demons of induction live in the boundaries of propositions one would like to know the truth values of. In a world that is an *interior point* of H, one eventually receives information verifying H (since an interior point of H has a neighborhood contained in H). Hence, not every verified proposition is verifiable, since a verified proposition merely has non-empty interior, whereas a verifiable proposition is open. But if a non-verifiable proposi-

[12]Let $\neg H$ denote $\Omega \backslash H$.

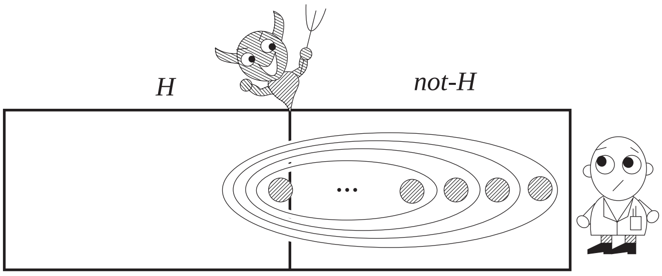

Figure 28. Demons Live in Boundaries.

tion is verified, some open proposition entailing it is verified, so information states can still be identified with open sets.

Less abstractly, recall the particle-counting problem. A possible world determines how many particles emerge from the machine for eternity and when each such particle emerges. Thus, one may model possible worlds as ω-sequences of bits, where 1 in position n indicates appearance of a new particle at stage n and 0 indicates that no new particle appears at n. Consider the situation in which you have seen the finite bit string $(b_0, \ldots, b_{n-1})$. The corresponding information state is the set of all ω-sequences of bits that extend the finite bit string observed so far. Call this proposition the *fan* with handle $(b_0, \ldots, b_{n-1})$, since all the worlds satisfying the fan agree up to n and then "fan out" in all possible ways from n onward (Figure 29). Any disjunction of verifiable events is verifiable (see above), so any union of fans is also verifiable and, hence, open (just wait for the handle of one of the fans to appear before turning on the light). The resulting space over arbitrary, ω-sequences of bits is very familiar in topology, where it is known as as the *Cantor* space. In the particle-counting problem, it is assumed that at most finitely many particles will appear, so one must restrict Cantor space down to the ω-sequences that converge to 0.

Consider the proposition that exactly two particles will be observed for eternity. This proposition is impossible to verify (no matter what you see, another particle may appear later). Hence, its interior is empty and every element is a boundary point, where the problem of induction arises. In this space, the boundary points are particularly suggestive of the problem of induction (Figure 30). For example, consider the world $(1, 1, 0, \ldots)$ where the dots indicate an infinite tail of zeros. No matter how far you travel down this sequence (*i.e.*, no matter what information you receive in this world), there exist worlds in which more than two particles appear later

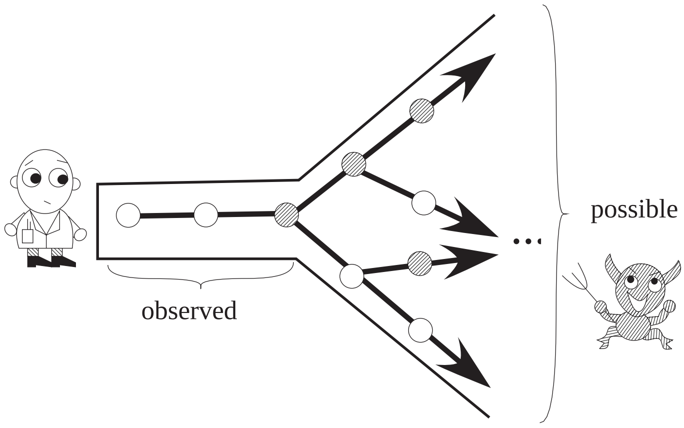

Figure 29. A Fan of Sequential Worlds.

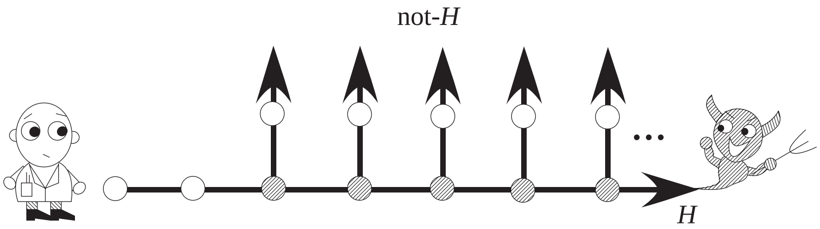

Figure 30. Boundary Points in Cantor Space.

than you have observed so far. So nature is in a position to drag you down the sequence $(1, 1, 0, \ldots)$ until you cave in and say "two" and is still free to show you another particle, as in the U-turn argument. The U-turn argument hinges, therefor, upon the topologically invariant structure of boundary points between answers to a question.

8 The Unending Game of Science

Each scientific problem determines an infinite, zero-sum game of perfect information (*cf.* [Kec95, p. 137–148] or [Kel96, p. 121–137]) between the *scientist*, who responds to each information state by selecting an answer (or by refusing to choose), and the impish *inductive demon*, who responds to the scientist's current guess history with a new information state. The demon is not a malicious force in nature; he merely personifies the difficulty of the challenge the scientist poses for himself by addressing a given scientific problem.

In this *truth-finding game*, the demon and the scientist take turns, starting with the scientist (Figure 31). Together, the demon and the scientist

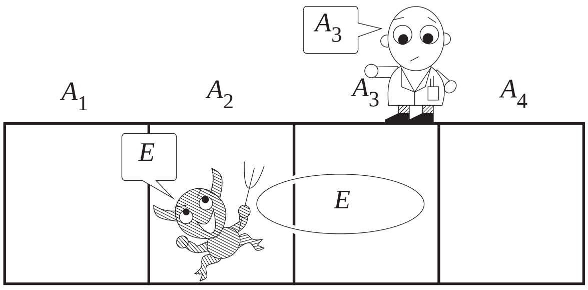

Figure 31. The Players.

produce a pair of ω-sequences, an *information sequence* produced by the demon and an *answer sequence* produced by the scientist. Life would be too easy for the demon if he were allowed to withhold some crucial information for eternity, so the scientist is the victor by default if the demon fails to present complete, true information about some world in Ω in the limit.[13] In other words, an information sequence $\{E_i : i \in \omega\}$ for the problem should be a downward-nested sequence of open sets whose intersection is non-empty and contained in some answer A. Then say that the information sequence is *for* A.

The scientist wins the *convergence game* by default if the demon fails to present an information sequence for some world in the problem and by merit if his outputs stabilize eventually, to the answer true in some world w the demon presents true information for. In other words, if the demon presents a legitimate information sequence for w, there must exist a stage in the play sequence after which the scientist's answer is correct of w. A winning strategy for the scientist in the convergence game is called a *solution* to the underlying empirical problem. For example, the obvious counting strategy solves the particle-counting problem. A problem is *solvable* just in case it has a solution.[14]

[13]One might reply that if it is impossible for the demon to fulfil his duty, the scientist loses since even the total experience received in the limit of inquiry doesn't settle the question. The game could be set up to reflect either viewpoint.

[14]It is interesting to inquire into the topological nature of solvability, since solvability

9 Comparing Mind-Changes

Consider two possible sequences of answers, σ and τ. Say that σ *maps into* τ (written $\sigma \leq \tau$) just in case there is an answer and order preserving mapping (not necessarily one-to-one) from positions in σ to positions in τ, where suspension of judgement is a wild-card in the former sequence that matches any answer in the latter (Figure 32).[15] Since the mapping

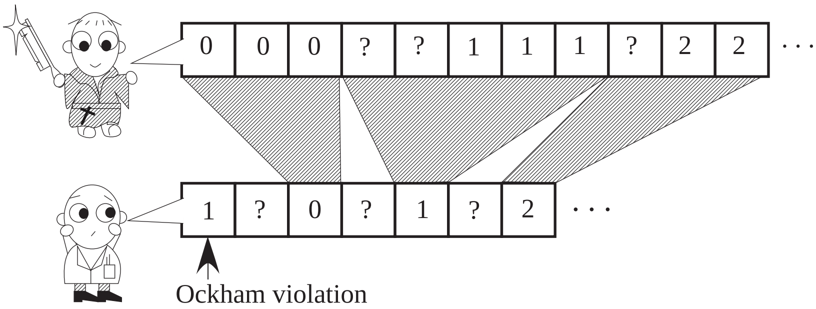

Figure 32. Top Output Sequence Better Than Bottom.

preserves answers and order, it also preserves mind-changes (not counting mere suspension as a mind-change). So when σ maps into τ, one may say that σ is *as good as* τ so far as mind-changes are concerned. Say that σ maps *properly* into τ (written $\sigma < \tau$) if, in addition, the latter fails to map into the former, as in Figure 32. Then σ is *better than* τ.

One can also say of two *sets* of output sequences that the former is *as good as* the latter just in case each element of the former is as good as some element of the latter (Figure 33) and is *better than* the latter if, in addition, the latter is not as good as the former.[16] The former set is *strongly better* than the latter just in case each of the former's elements is better than

is a topological invariant and must, therefore, be grounded in a problem's topological structure. For example, if the space is separable and the question is a countable partition, then solvability is equivalent to each cell being $\mathbf{\Delta}^0_2$ Borel (*cf.* [Kel96, p. 228, Corollary 9.10]). Such facts are not strictly necessary for understanding Ockham's razor, and are therefore omitted from this essay.

[15] In fact, an Ockham method's output sequences map *injectively* into output sequences the demon can force out of an arbitrary method. In [Kel07], methods are compared in terms of Pareto-dominanace with respect to number and timing of retractions, where a retraction occurs whenever an informative answer is dropped. The streamlined account of costs just presented does not penalize gratuitous question marks or tardiness of retractions, since question marks are wild cards and the mappings employed are many-one.

[16] This is not the same as weak dominance, since the existential quantifier allows for a worst-case pairing of output sequences by the mapping.

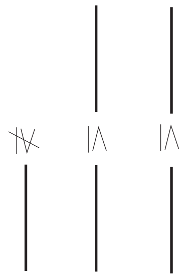

Figure 33. Top Set Better Than Bottom.

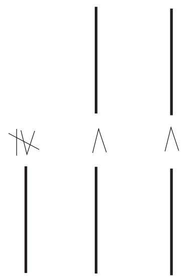

Figure 34. Top Set Strongly Better Than Bottom.

some element of the latter that is not as good as any element of the former (Figure 34).[17] Extend the symbols $\leq$ and $<$ to sets of output sequences accordingly.

The set of output sequences of a solution to a problem is the set of all output sequences σ such there exists some information sequence for some answer along which the method produces σ. Then one can say of two methods that the former is *as good*, *better*, or *strongly better* than the latter just in case their respective sets of output sequences bear the corresponding relation. Finally, say that a solution is *efficient* in a problem just in case it as good as any other solution in each subproblem. Again, the idea is

[17]The requirement that the sequence mapped to is not as good as any of the former method's output sequences precludes cases in which a method is strongly better than itself. For example, if there are only two answers in the particle problem, "even" and "odd", then each output sequence of the obvious method that answers according to whether the number of observed particles is even or odd is better than some other output sequence of the same method (*e.g.*, $(E, O, E, O, \ldots) < (O, E, O, E, O, \ldots)$).

that inefficiency is forward-looking and should not be offset by foibles or stockpiles of credit (slush funds) earned in the past.

By way of illustration, the counting solution is efficient in the particle-counting problem, as is any Ockham solution to this problem (remember that Ockham solutions can suspend belief for artibrary periods of time). That is because the demon can force any solution through any ascending sequence of answers and Ockham methods produce only ascending sequences of answers. Furthermore, any non-Ockham solution is worse than any Ockham solution in the subproblem entered when the violation occurs. Indeed, it was shown that the violator is strongly worse than any Ockham solution in that subproblem, because the demon can force the violator into any ascending sequence after the U-turn back to the Ockham answer. Hence, the counting problem has the remarkable property that *its efficient solutions are exactly its Ockham solutions.* That is surely a result worth pressing as far as possible! But first, Ockham's razor must be defined with corresponding generality.

10 What Simplicity Isn't

The concept of simplicity appears, at first, to be a hodge-podge of considerations, including uniformity of nature, theoretical unity, symmetry, testability, explanatory power, and minimization of entities, causes, and free parameters. But in spite of these manifold aspects, it remains possible that simplicity is a deep, unified, concept with multiple manifestations, depending on the particular structure of the problem addressed. It is suggestive in this regard that the trivial particle-counting problem already illustrates all of the intuitive aspects of simplicity just mentioned and that they seem to cluster around the nested problems of induction posed by the repeated possibility that a new particle might appear.

It is easy, at least, to say what simplicity *couldn't* be. It couldn't be anything fixed that does not depend on the structure of the problem. For it is a commonplace in the analysis of formal procedures that different algorithmic approaches are efficient at solving different problems. So if simplicity did not depend, somehow, on the structure of the particular problem addressed, Ockham's razor couldn't possibly be necessary for efficient convergence to the truth in a wide range of distinct problems possessing different structures.

That is the trouble with concepts of simplicity like notational brevity [LiVit97, p. 317–337], uniformity of worlds [$\mathrm{Car}_0$50, p. 562–567], prior probabilistic biases, and historical "entrenchment" [Goo83, p. 90–100]. Left to themselves, none of these ideas conforms to the essential structural interplay between a problem's question and its underlying informational topology, so none of them could contribute objectively to truth-finding efficiency over a range of different problems. All of them could be made to do so by selecting notation that reflects the relevant structure of the problem addressed [Mit97, p. 174]. But then the essence of simplicity is captured by the rules for selecting appropriate notation, rather than by brevity, uniformity, or the like.

11 Simplicity and Ockham's Razor Defined

The task of defining simplicity is facilitated by knowing in advance how Ockham's razor is justified. We can, therefore, "solve backwards" for simplicity, by generalizing the features of particle counting that give rise to the the U-turn argument. The key to the U-turn argument is the demon's ability to *force* a given sequence of mind-changes from an arbitrary solution. In the particle-counting problem, the demon can present information from the zero-particle world until the scientist caves in and concludes that there will be zero particles (on pain of not converging to the true answer) (Figure 35). Then the demon can present a particle followed by no further particles

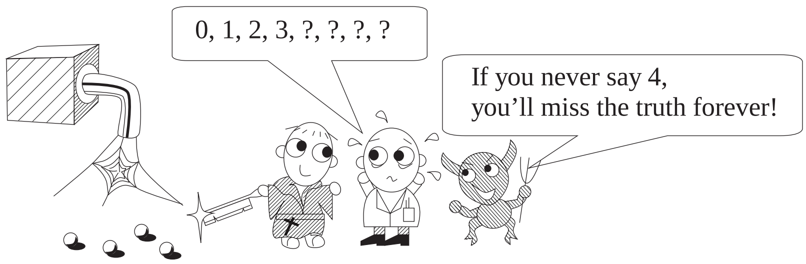

Figure 35. Demon Forcing a Sequence of Answers.

until the scientist concludes "one particle", again on pain of not converging to the true answer, and so forth. This can't go on forever, though, because the demon must present data from some world in Ω, and all such worlds present at most finitely many particles. Hence, for each finite ascending sequence σ of answers, the demon can force an arbitrary solution to the particle-counting problem into an output sequence that σ maps into. But the demon has no strategy for dragging an arbitrary solution through any *non-ascending* sequence, say, $(1, 0)$. For the obvious counting method will wait to see the first particle before concluding "one" and, thereafter, the demon can no longer trick it into thinking that there are no particles, since the particle has already been presented. That is a fundamental asymmetry in the problem.

More generally, if σ is a finite, non-repetitive sequence of answers, then the *σ-avoidance game* for a problem is won by the scientist just in case the demon fails to present an appropriate information sequence or the scientist wins the truth-finding game *and* fails to produce a sequence of conjectures as bad as σ. The demon wins if he presents appropriate information that makes the scientist lose the truth-finding game *or* that somehow lures the scientist into producing an output sequence as bad as σ. When the demon has a winning strategy in the σ-avoidance game, one may say that the demon can *force* σ from an arbitrary solution to the problem. For example, it was shown that the demon has a winning strategy in the $(0, 1, 2, \ldots, n)$-avoidance game in the particle-counting problem, since every method can be induced to produce that output sequence (or a sequence that is at least as bad). Then say that σ is *demonic* in a problem just in case the demon can force it in the problem.

The demonic sequences in a problem reflect a deep relationship between the question $\mathcal{Q}$ and the underlying topology $\mathcal{V}$. The ability of the demon to force demonic sequence $(0, 1, 2, \ldots, n)$ implies that there is a zero particle world that is a limit point of one particle worlds each of which is a limit point of two particle worlds and so forth. So demonic sequences represent iterated problems of induction within the overall problem (Figure 36). According to

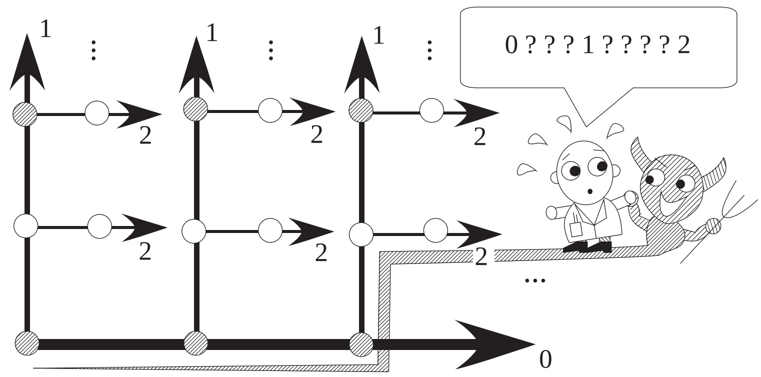

Figure 36. Demonic Sequence in the Particle-Counting Problem.

intuition, simpler answers are associated with the most deeply embedded problems of induction, for starting from 0, the demon can drag a solution through every ascending sequence, but after presenting some particles, he can never drag the counting solution back to 0. That suggests a natural definition of empirical simplicity. If A is a potential answer, then say that the *A-sequences* are the demonic sequences starting with A. Say that answer A is *as simple as* B just in case the set of B-sequences is as good as the set of A-sequences and that A is *simpler than* B just in case the set of B-sequences is better than the set of A-sequences. This definition agrees with intuition in the counting problem and, hence, in parameter-freeing problems of the usual sort, such as curve-fitting.

The proposed explication of simplicity has a striking, intuitive advantage over the familiar idea that simplicity has something to do with dimensionality of an answer in a continuous parameter space. For if you were to learn that the true parameter values are rational, then the topological dimension of each answer would drop to zero, flattening all simplicity distinctions. But since the rational-valued subspace of the parameter space preserves demonic sequences, it also preserves simplicity in the proposed sense.

Now say, quite naturally, that an answer A is *Ockham* in a problem just in case A has an information sequence and is as simple in the problem as any other answer. This is a way of saying that A is the simplest answer compatible with experience, where compatibility with experience includes the assumption that complete evidence can be presented at all; else there is no use debating whether simplicity helps one find the truth. Ockham's razor is then: *never say an answer unless it is Ockham for the current subproblem.* Finally, a solution is *Ockham* if it solves the problem and always heeds Ockham's razor. "The" Ockham answer is typically unique, but not always. If there is no Ockham answer, an Ockham method must suspend judgment. If there is more than one, an Ockham method may produce any one of them. It may sound odd to allow for an arbitrary choice among Ockham answers,[18] but keep in mind that two hypotheses could be maximally simple (no other answer's demonic sequences are *worse*) without being Ockham (every other answer's demonic sequences map in). The truly objectionable choices turn out to be among maximally simple answers that are not Ockham, as will be explained below.

Here is a handy re-formulation of the Ockham answer concept, where $*$ denotes concatenation. The proofs of all propositions are presented in the Appendix.

Proposition 1 (Ockham characterization). If the problem is solvable, A is Ockham if and only if for every demonic sequence σ for the problem, $A * \sigma$ is demonic for the problem.

Proof. Suppose that A is Ockham. Let σ be an arbitrary, demonic sequence. If σ is empty, then trivially $A * \sigma$ is demonic, since there is an information sequence for some world in A and a solution must converge to A on this

[18] The theory presented in [Kel07] does not have this property.

sequence. So suppose that $\sigma = B * \gamma$. Since A is Ockham, there exists demonic $A * \tau$ such that $B * \gamma \leq A * \tau$, in virtue of some mapping φ. If $\varphi(0) = 0$, then $B = A$, so:

$$A * \tau \geq A * \gamma \geq A * A * \gamma = A * (B * \gamma) = A * \sigma.$$

So since $A * \tau$ is demonic, so is $A * \sigma$. If $\varphi(0) > 0$, then since $B * \gamma \leq A * \tau$ and $B \neq A$, it follows that $\sigma = B * \gamma \leq \tau$. Hence: $A * \tau \geq A * \sigma$. So since $A * \tau$ is demonic, so is $A * \sigma$.

Conversely, suppose that for each demonic sequence σ, $A * \sigma$ is demonic. The empty sequence $()$ is trivially demonic, so $A * () = (A)$ is demonic, so some $\delta \geq (A)$ is forcible, so *there is some information sequence whose intersection is* A. For suppose otherwise. Since the problem is solvable, let a solution be given. Since there is no information sequence for A, the solution remains a solution if its output is changed to "?" each time it produces A. Hence, no sequence as bad as (A) is forcible. Contradiction. Finally, let B be an answer and let $B * \sigma$ be demonic. Then by assumption, $A * B * \sigma$ is demonic. So the A sequences are as bad as the B sequences. Thus, in light of the italicized claim, A is Ockham. q.e.d.

12 Efficiency, Games and Determinacy

Lifting the U-turn argument to the general version of Ockham's razor just defined requires a short digression into the nature of efficient solutions. A method is *as good* as a set of sequences just in case the method's set of output sequences is as good as the given set, and similarly for the other ordering relations defined above. Then it is immediate that the demonic sequences are as good as an arbitrary, efficient solution to the subproblem, since each solution can be forced to produce each demonic sequence. It is far less trivial whether an efficient solution must be as good as the set of demonic sequences. This is where Ockham's razor interfaces with recent developments in descriptive set theory (*cf.* [Kec95, p. 137–146] for a succinct development of the following material).

Say that a game is *determined* just in case one player or the other has a winning strategy and that a scientific problem is *forcing-determinate* just in case for each finite answer sequence σ, the σ-avoidance game is determined. Forcing-determinacy turns out to be a surprisingly mild restriction. Say that a problem is *typical* just in case the set $\mathcal{Q}$ of possible answers is countable and the set of possible information streams for worlds in Ω is Borel. Then the following proposition is an easy consequence of Donald Martin's Borel determinacy theorem.[19]

[19] *Cf.* [$\mathrm{Mar}_1$75].

Proposition 2. Typical, solvable problems are forcing-determinate.

Proof. Let $(\Omega, \mathcal{V}, \mathcal{Q})$ be a solvable problem, let ε be an ω-sequence of open sets, let γ be an ω-sequence of answers and let σ be a finite sequence of answers. The pair (ε, γ) wins for the demon in the σ-avoidance game in $(\Omega, \mathcal{V}, \mathcal{Q})$ if and only if (i) ε is an information sequence for some answer in $\mathcal{Q}$ and either (ii) there exists answer A in $\mathcal{Q}$ such that ε is for A and γ does not converge to A or (iii) γ is as bad as σ. Condition (i) is Borel by assumption. Condition (iii) is open and, hence, Borel. In light of condition (i) and the fact that the problem has a solution M, condition (ii) reduces to: there exists an A in $\mathcal{Q}$ such that M converges to A along ε and γ does not converge to A. Convergence is Borel and $\mathcal{Q}$ is countable, so the overall winning condition for the demon is Borel. Apply Martin's (1975) Borel determinacy theorem, which states that all such games with Borel winning conditions are determined. q.e.d.

The following results all concern solutions and, hence, are vacuously true if the problem in question is unsolvable. Therefore:

Proposition 3. Each of the following propositions remains true if "forcing-determinate" is replaced with "typical".

Proof. Immediate. q.e.d.

Now it is easy to show that:

Proposition 4 (Efficiency Characterization). Let the problem be forcing-determinate. An arbitrary solution is efficient if and only if it is as good as the set of demonic sequences in each subproblem.

Proof. Let an efficient solution M to a forcing-determinate problem be given. Then in each subproblem, M is as good as an arbitrary solution. Let σ be a finite output sequence of M in a given subproblem. So every solution to the subproblem produces an output sequence as bad as σ, so the scientist has no winning strategy in the σ-avoidance game. So by forcing-determinacy, the demon has a winning strategy, so σ is demonic. So an efficient method is as good as the demonic sequences in each subproblem. Conversely, suppose that M is as good as the set of demonic sequences in a given subproblem. By definition, the set of demonic sequences is as good as an arbitrary solution in the subproblem. So M is as good as an arbitrary solution in the subproblem and, hence, is efficient in the subproblem. q.e.d.

So not only is an efficient solution as good as any solution, it is as good *because* it is as good as the demonic sequences, which are as good as any solution.[20]

13 Efficient Solutions = Ockham Solutions

Here is the main result. Ockham is indeed necessary *and* sufficient for efficiency in an extremely broad range of problems. The hypothesis of forcing-determinacy makes the proof surprisingly easy.

Proposition 5 (Ockham Equivalence Theorem). Let the problem be forcing-determinate. Then the efficient solutions are exactly the Ockham solutions.

Proof. Let a forcing-determinate problem be given. For the necessity argument, suppose that solution M violates Ockham's razor upon entering some subproblem by producing non-Ockham answer A. Let D be the set of demonic sequences for the subproblem. Since A is not Ockham and M is a solution, there exists (by Proposition 1) a demonic sequence σ in the subproblem such that $A * \sigma$ does not map into any demonic sequence. Hence, $M \not\leq D$. So by Proposition 4, M is not efficient.

For sufficiency, it suffices to argue that *every finite sequence of Ockham answers encountered in subproblems successively reached as experience increases maps into some demonic sequence in the first subproblem*. For then an Ockham solution, which produces only sequences of Ockham answers interspersed with question marks, is as good as the demonic sequences in an arbitrary subproblem and, hence, is efficient, since the demonic sequences are, by definition, as good as an arbitrary solution. In the base case, each Ockham answer A in a subproblem has an information sequence for it, so the singleton sequence (A) can be forced by the demon in the subproblem and, hence, is demonic in the subproblem. Now consider a finite, downward-nested sequence of non-empty open sets $(E_0, \ldots, E_{n+1})$ determining respective sub-problems with respective Ockham answers $(A_0, \ldots, A_{n+1})$. By the induction hypothesis, $(A_1, \ldots, A_{n+1})$ is demonic in P_1. Furthermore, since E_1 is a non-empty subset of E_0, whatever the demon can force in P_1 he can force in P_0, so $(A_1, \ldots, A_{n+1})$ is demonic in P_0. So since A_0 is Ockham in P_0, $(A_0, A_1, \ldots, A_{n+1})$ is demonic in P_0, which proves the italicized claim and, hence, the theorem. q.e.d.

More can be shown for the particle-counting problem and for others of its attractive kind. For such problems have the special feature that in each

[20] Such a result is called a *universal factorization* [Mac71, p. 1–2].

subproblem, if A is an Ockham violation upon entering the subproblem, then there exists an Ockham answer U upon entering the subproblem such that the binary sequence $A * U$ is not as good as any demonic sequence for the subproblem. Say that such problems are *stacked*.[21] Examples of non-stacked problems illustrate intuitive ideas about empirical symmetry and will be considered in the next section. The result is:

Proposition 6 (Strong Ockham Necessity for Stacked Problems). In a stacked, forcing-determinate problem, if a solution violates Ockham's razor upon entering a sub-problem, the solution is stongly worse than each efficient solution in the same sub-problem.

Proof. Consider an efficient solution to a stacked, forcing-determinate problem and suppose that M solves a given subproblem but violates Ockham's razor upon entering it by producing A. Let U be the Ockham answer promised by the stacking property. Then since M already says A and U is compatible with the current subproblem, the demon can force M to produce U after producing A. (That is the initial U-turn resulting from the Ockham violation). Consider an arbitrary, finite output sequence τ of the efficient solution. Then for some demonic δ, $\tau \leq \delta$ (Proposition 4) and, hence, $\tau \leq A * U * \delta$. Since U is Ockham and δ is demonic, $U * \delta$ is demonic (by Proposition 1). So since M already says A and $U * \delta$ is demonic, $A * U * \delta$ is as good as one of the output sequences of M in the current subproblem. Of course, $\tau \leq \delta \leq A * U * \delta$, so M is as bad as the efficient solution. Furthermore, $A * U$ maps into no demonic sequence, so neither does $A * U * \delta$. Since all the optimal method's output sequences map into demonic sequences, it follows that $A * U * \delta$ is as good as none of the optimal method's output sequences. Hence, M is *strongly* worse than the efficient solution in the current subproblem. q.e.d.

This fits closely with the spirit of the freeway example and with what is going on in particle counting and curve fitting. The property of being "stacked" can be viewed as the topological essence underlying the very strong Ockham intuitions attending such problems.

[21] To see that the particle-counting problem is stacked, suppose that A is not Ockham upon seeing, say, four particles. Let U be the Ockham answer "four". Then the binary sequence $A * U$ maps into no demonic sequence in the subproblem. For if A posits fewer than four particles, A maps into no demonic sequence since the demon can't force an arbitrary solution into a refuted answer. If A posits more particles, then (A, U) maps into no demonic sequence since all such sequences are ascending.

14 Testing as an Instance of Ockham's Razor

Suppose that you want to decide whether some theory is true or not. That poses a binary question: the theory vs. the theory's denial. Typically, the theory is refutable and, hence, closed. Everyone chooses the refutable (*e.g.*, point) hypothesis as the null hypothesis and its denial as the alternative. On the proposed account of simplicity, the decision to accept the refutable hypothesis until it is refuted is an instance of Ockham's razor and is underwritten by the U-turn argument, so that the proposed account of efficient theory choice subsumes this aspect of testing practice as a special case.

First, observe that the demon can force you to conclude the refutable hypothesis H (by showing you a boundary point in the hypothesis, since closed sets contain all of their boundary points). Then he can show you data refuting the theory. So only $(H, \neg H)$ and its subsequences are demonic. Hence, only H is Ockham (Proposition 1), so (by Proposition 5) every efficient solution says H (or suspends) until H is refuted, which reflects practice. Finally, that practice is efficient (since its output sequences are all demonic), so Ockham's razor bites and you should heed his advice.

The trouble with standard conceptions of hypothesis testing is that they ignore the possibility of extra mind-changes. Yes, it is refutable to say that the bivariate mean of a normal distribution is precisely $(0, 0)$, since $\{(0, 0)\}$ is closed (and hence refutable) in the underlying parameter space. But what if you want to test the non-refutable and non-verifiable hypothesis that exactly one component of the mean is zero? Solving this binary question requires multiple mind-changes, as in particle-counting and other model selection problems. For the demon can make it appear that both components are zero until you probably say "no" (as sample size increases) and can then reveal deviation of one component from zero until you probably say "yes" and then can reveal deviation of the other component from zero until you probably say "no" again, for a total of two mind-changes (in probability). Essentially, you are just counting deviations of mean components from zero as you were counting particles before. So the demonic sequences are all the sequences that map into "yes, no, yes", so the obvious method of counting nonzero mean components is efficient and the unique Ockham hypothesis at each stage is the one that agrees with the current nonzero mean count. So you should heed Ockham's advice, (as you naturally would in this example).

Since testing theory usually isn't applied until all the parameters in a model are fixed by point estimates, it appears as though a testing theory for refutable (closed) hypotheses is good enough. Hence, the essential involvement of Ockham's razor in testing theory is missed and so the strong analogy between model selection and testing with multiple mind-changes is missed as well. The proposed account of Ockham's razor, therefore, suggests a new, more unified foundation for classical statistics, whose development lies beyond the scope of this explorative essay.

15 Ockham and Respect for Symmetry

When there are two maximally simple answers compatible with the data, Ockham can't help you decide among them and the strong intuition is to wait for nature to "break the symmetry" prior to choosing. For example, modify the particle-counting problem so that particles come in two colors, green and blue and you have to specify the total number of each that will ever be emitted. Assume also that you can hear particles rattle down the faucet before they emerge from the machine (Figure 37). Having seen no

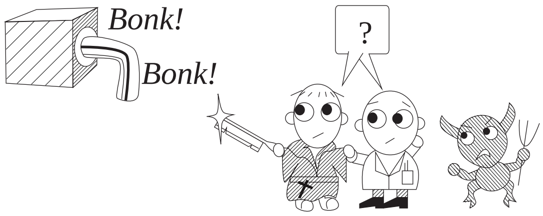

Figure 37. Breaking Symmetry.

particles, you hear the first one coming. What do you conclude? It is hard to say, for both colors of marbles will stop appearing, eventually, so there is no general "pattern" to detect in the data, and there is no obvious primacy of one color over the other so far as the underlying problem is concerned. This is not mere skepticism, since after the next marble is observed, you will eventually have to leap to the bold Ockham hypothesis that no more particles are coming. Instead, it is respect for symmetry, one of the strongest intuitions in science since Greek times. That leads to an intriguing idea. Perhaps the U-turn argument also explains our strong hesitance to break symmetries in experience. Then respect for symmetry would simply be Ockham's razor conforming itself to the structure of symmetrical problems. That is correct.

Consider how Ockham's razor applies to the case at hand. When you hear the rattling that announces the first particle, you have entered a new subproblem. There are two equally simple answers at that point, "one green, zero blue" or symmetrically "zero green, one blue". But *neither* of these answers is Ockham. For each answer constitutes a unit demonic sequence, but neither binary sequence consisting of the two symmetrical competitors is demonic in the subproblem, since the demon can't take back the first particle after its color is observed. So Ockham demands silence, and we already know from Proposition 5 that every efficient solution to the problem must heed this advice. Is there an efficient solution? Indeed, just heed Ockham's advice by counting the total number of particles whose colors are seen and by suspending judgment when the next rattle is heard. Respect for symmetry follows from Ockham's razor.

The symmetrical problem under discussion is a nice example of a non-stacked problem. For consider the answer "zero green, one blue". There is no Ockham answer one can concatenate to this answer in the subproblem entered with the first rattle because there is no Ockham answer at all. And the violator is not strongly worse than the Ockham method just described in that subproblem, because the demon can force even an optimal method to say the same answer the violator chose in advance and the violator produces no output sequence worse than that.

The same argument works even after a run of a thousand exclusively green particles, in which case it might be objected that past experience does break symmetry between blue and green. But the subproblem so entered is topologically equivalent to the original problem prior to seeing any marbles. Hence, no non-circular, efficiency-based account of Ockham's razor could possibly explain why it is more efficient to say "green" rather than "blue" upon entering the subproblem.

In the preceding problem, the counting question slices the problem into topologically invariant simplicity degrees corresponding to particle counts in spite of occasional symmetries (*e.g.*, when the particle rattles and has not yet been seen). In other problems, symmetry is so pervasive that Ockham's razor doesn't bite at all (Figure 38). For example, suppose you have

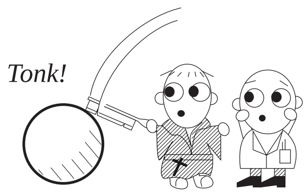

Figure 38. Overly Symmetrical Problems.

to report not only how many particles will appear, but when each one will appear (forgetting about color). It might seem, at first, that the simplest answer is to say that you have seen all the particles already and that they appear exactly when they were observed to, since if you were asked only how many particles there are, you would only be permitted to say the number seen so far. That is so, if you choose to conceive of the sequence identification problem as a refinement of the particle counting problem. The trouble is that the sequence identification problem also refines alternative problems that lead to incompatible simplicity judgments. For example, *oneicles* are non-particles up to stage one and particles thereafter. There are finitely many particles if and only if there are finitely many oneicles, so the underlying space Ω is unaltered by the translation. But the answers to the two counting problems are different and the U-turn argument leads to correspondingly different recommendations (*i.e.*, to count particles or to count onecles, respectively). Since the sequence identification problem refines the problems of counting particles, oneticles, twoticles, threeticles, *etc.*, it can't consistently favor one kind of counting over another without making a global, symmetry-breaking choice in favor of one of its possible coarsenings. The only sensible resolution of this Babel of alternative coarsenings is for Ockham to hold his tongue.

And that's just what the proposed theory says. First of all, no answer is Ockham in this problem, since every demonic sequence is of unit length. For consider a single answer. The answer is true in just one world, which the demon can present until you take the bait. So each unit sequence of answers can be forced. For each alternative answer (satisfied by an alternative world), there is a least stage by which the two cease agreeing and diverge. But some solution refuses to be convinced of the first answer (on pain of converging to the wrong answer) until the divergence point is already passed (Figure 39). So the demon can force no binary sequence of answers

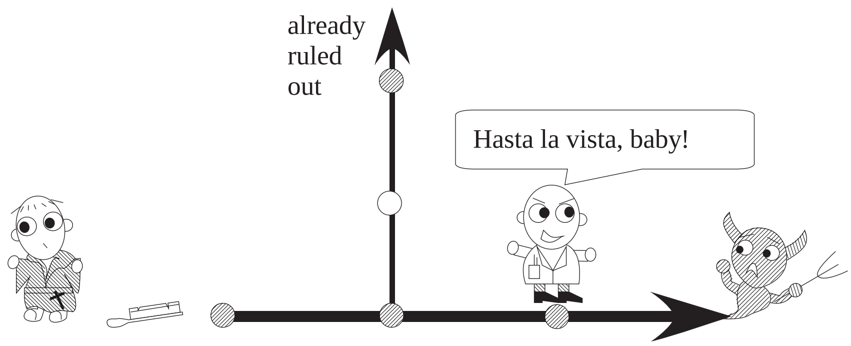

Figure 39. The Trouble With Singleton Answers.

from an arbitrary solution. Hence (Proposition 4), there can be no efficient solution, since no solution to this problem succeeds without mind-changes.

So there are lots of solutions to this problem, but no efficient ones. Hence, even if there were an Ockham answer, there would be no efficient method to put normative teeth into the U-turn argument! Ockham is both mute and toothless in this problem.

Again, that is the correct answer. The sequence-identification problem is completely symmetrical in the sense that any homeomorphism of the space into itself results in the very same problem (since each permuted world still ends up in a singleton answer over the same topological space). So there is no objective, structural sense in which one answer is simpler than another any more than there is any objective physical sense about where zero degrees longitude is. Coordinate systems are not physically real because they aren't preserved under physical symmetries; philosophical notions of simplicity (*e.g.*, brevity, sequential uniformity, entrenchment) are not real because they aren't preserved under problem symmetries. To seek objective, truth-finding efficiency in distinctions that really aren't in the problem is like trying to generate electricity by formally spinning coordinate axes. The situation is different in the counting problem. There exist homeomorphisms of the underlying topological space that materially alter the original problem (*e.g.*, the unique Ockham hypothesis "no particles" would become "no oneicles", which means "one particle at stage 1"). It is precisely this lack of symmetry in the particle-counting problem that allows Ockham to slice it into objective simplicity degrees.

The usual attempts to use coding, "entrenchment", or prior probability to force a foliation of the sequence identification problem into simplicity degrees must involve the imposition of extraneous considerations lacking in the problem's intrinsic structure as presented. Therefore, such considerations couldn't possibly have anything objective to do with solving the problem (as stated) efficiently. So the proposed account yields precisely the right judgment in this example when its true nature is properly understood.

One can also arrive at overly-symmetrical problems by coarsening the particle-counting problem. For example, consider the question whether there is an even or an odd number of particles. Since this coarsens the particle-counting problem, one again expects "even" to be the Ockham answer when an even number of particles have been observed and "odd" to be the right answer otherwise (Figure 40). But the proposed theory of Ockham's razor doesn't agree. Ockham is once again silenced, but this time the difficulty is exactly reversed: every solution is efficient and every answer is Ockham in every subproblem so every method satisfies Ockham's razor and the U-turn argument can't even begin (Figure 41).[22]

[22]In the theory presented in [Kel07], there is no Ockham solution to this problem. Either way, Ockham refuses to choose among potential solutions to the problem.

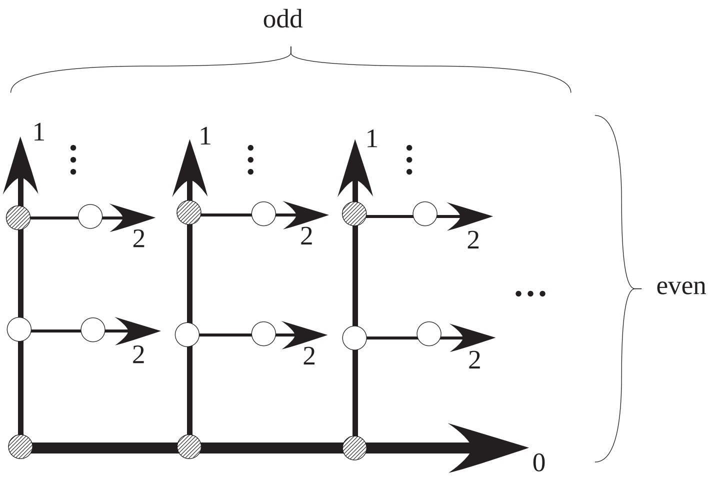

Figure 40. Even/Odd as Particle Counting.

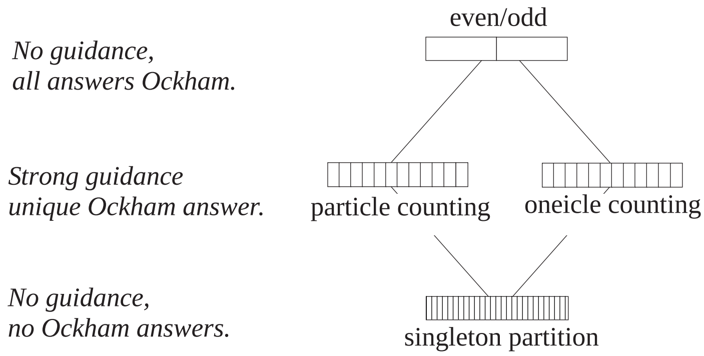

Figure 41. Ockham Under Refinement.

The theory is right. Yes, if one thinks of the problem as a coarsening of particle counting, “even” must come first. But one could also think of it as a coarsening of counting oneicles instead of particles. Then the zero oneicle world is an “odd” world. The one oneicle worlds include the zero particle world as well as all the two particle worlds in which the first appears right away. These are all “even particle” worlds. Continuing in this way one obtains a oneicle-counting simplicity foliation (Figure 42) in

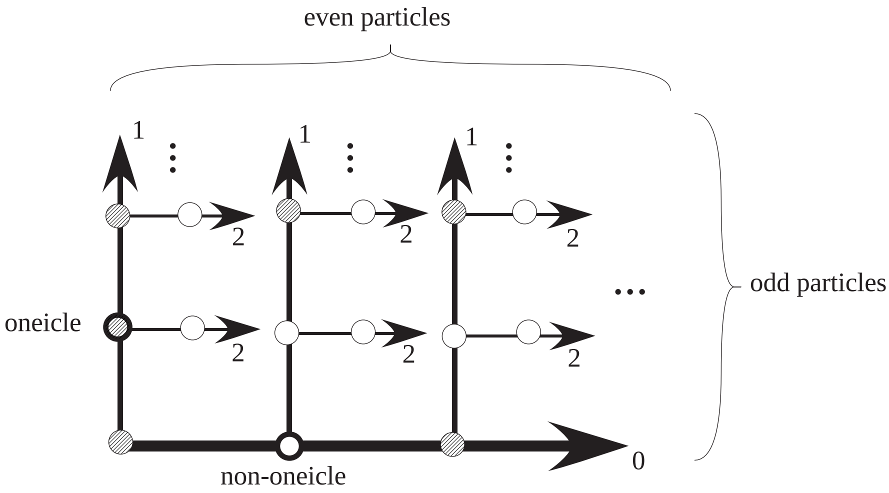

Figure 42. Even/Odd as Oneicle Counting.

which the obvious "first" conjecture is "odd". But the oneicle translation is a homeomorphism of the space that reflects each answer onto the other, so a prior preference for "even" couldn't have anything to do with the objective efficiency of solutions to the even/odd problem as stated.

The urge to extend Ockham's advice to symmetrical problems is understandable — guidance is most precious when there is none. And even in light of the proposed account, nothing prevents us from saying that we are really interested in counting marbles rather than merely saying whether they are even or odd, in which case the problem is no longer symmetrical. But it is quite another matter to smuggle extra structure into a symmetrical problem without acknowledging that one has done so, for such practice is not warranted by truth-finding efficiency in the problem addressed (Figure 43).

16 Conclusion: Ockham's Family Secret

Ockham is beloved as an inexhaustible source of free information that somehow parlays the scientist's limited viewpoint into sweeping generalizations about unseen realities (Figure 44). But his very appeal is his undoing, for it is impossible to explain how his fixed advice could be true without assuming exactly what we rely upon him to tell us.

This paper presents an alternative view, according to which Ockham

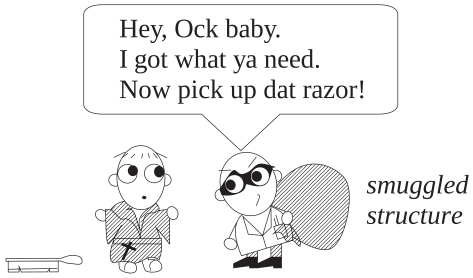

Figure 43. Theft Over Honest Toil.

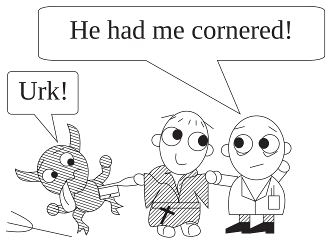

Figure 44. Ockham's Day Job.

helps us find the truth, but in an unexpected way. He doesn't provide any guarantee that the theory he selects is true or probably true. He doesn't point at the truth. He can't even bound the number of future surprises or U-turns you will have to make in the future on your way to the truth. All he does is save you the trouble of *needless* surprises beyond those arbitrarily many surprises nature is objectively in a position to exact from you. But in that respect, his advice is still uniquely the best.

The proposed explanation is unnerving because it singles out simplicity as the right bias to have, but falls so far short of our craving for certainty, verification, and guarantees against future surprises. That is far harder to dismiss than the usual, academic sort of skepticism, which finds no connection between simplicity and truth and urges rejection of simplicity-based conclusions altogether.

It is also ironic that Ockham is viewed as a comforting figure when, in fact, he is built out of the inductive demon's opportunities to successively force science to reverse course. Indeed, Ockham and the demon work together as a coordinated team, since Ockham changes his recommendations each time the demon uses up one of his opportunities to fool the scientist (Figure 45).

Figure 45. But by Night. . .

The key to understanding Ockham's razor is to set aside our instinctive appetite for certainty and to focus squarely on the objective complexity properties of empirical problems that underly unavoidable reversals of scientific opinion through time. A similar focus on problem complexity has long been the norm in the mathematical theories of computability, computational complexity, and descriptive set theory. In these established, scientific subjects, nobody would dream of "victory" over complexity. It is late in the day for the philosophy of science and induction to be dreaming still.

References.

[$\text{Car}_0\text{50}$] Rudolf **Carnap**, Logical Foundations of Probability, University of Chicago Press 1950

[Ear83] John **Earman** (*ed.*), Testing Scientific Theories, Universitiy of Minnesota Press 1983

[$\text{For}_1\text{Sob94}$] Malcolm R. **Forster** and and Elliott **Sober**, How to Tell When Simpler, More Unified, or Less Ad Hoc Theories will Provide More Accurate Predictions, **The British Journal for the Philosophy of Science** 45 (1994), p. 1–35

[$\text{Fri}_1\text{83}$] Michael **Friedman**, Foundations of Space-Time Theories, Princeton University Press 1983

[$\text{Fri}_2\text{GoeHar}_0\text{07}$] Michèle **Friend**, Norma B. **Goethe**, and Valentina S. **Harizanov** (*eds.*), Introduction to the Philosophy and Mathematics of Algorithmic Learning Theory, Springer 2007 [Handbook of Philosophical Logic 9]

[Ger75] Carl J. **Gerhardt** Die Philosophischen Schriften von G. W. Leibniz, Volume IV, Weidmann 1875

[Gly80] Clark N. **Glymour**, Theory and Evidence, Princeton University Press 1980

[GlyKel04] Clark N. **Glymour** and Kevin T. **Kelly**, Why Probability Does Not Capture the Logic of Scientific Justification, *in:* [Hit04, p. 94–114]

[GlySch$_1$Spi00] Clark N. **Glymour**, Richard **Scheines**, and Peter **Spirtes**, Causation, Prediction, and Search, 2nd edition, MIT Press 2000

[Goo83] Nelson **Goodman**, Fact, Fiction, and Forecast, 4th edition, Harvard University Press 1983

[Hal74] Paul R. **Halmos**, Measure Theory, Springer 1974

[Hit04] Christopher **Hitchcock** (*ed.*), Contemporary Debates in the Philosophy of Science, Blackwell 2004

[Jai+99] Sanjay **Jain**, Daniel **Osherson**, James S. **Royer**, and Arun **Sharma**, Systems That Learn: An Introduction to Learning Theory, MIT Press 1999

[Kec95] Alexander S. **Kechris**, Classical Descriptive Set Theory, Springer 1995 [Graduate Texts in Mathematics 156]

[Kel96] Kevin T. **Kelly**, The Logic of Reliable Inquiry, Oxford 1996

[Kel02] Kevin T. **Kelly**, Efficient Convergence Implies Ockham's Razor, Article for the 2002 International Workshop on Computational Models of Scientific Reasoning and Applications, Las Vegas, USA, June 24-27, 2002

[Kel04] Kevin T. **Kelly**, Justification as Truth-finding Efficiency: How Ockham's Razor Works, **Minds and Machines** 14 (2004), p. 485–505

[Kel07] Kevin T. **Kelly**, How Ockham's Razor Helps You Find the Truth — Without Pointing at It, *in:* [Fri$_2$GoeHar$_0$07, p. 111–143]

[Kuh$_1$62] Thomas **Kuhn**, The Structure of Scientific Revolutions, University of Chicago Press 1962

[Lau81] Larry **Laudan**, A Confutation of Convergent Realism, **Philosophy of Science** 48 (1981), p. 19–48

[Lei14] Gottfried W. **Leibniz**, Monadologie, 1714, *in:* [Ger75, p. 607–23]

[LiVit97] Ming **Li** and Paul **Vitanyi**, An Introduction to Kolmogorov Complexity and Its Applications, Springer 1997

[LiVit00] Ming **Li** and Paul **Vitanyi**, Minimum Description Length Induction, Bayeisanism, and Kolmogorov Complexity, **IEEE Transactions on Information Theory** 46 (2000), p. 446–464

[LohZhe95] Wei-Yin **Loh** and Xiaodong **Zheng**, Consistent Variable Selection in Linear Models, **Journal of the American Statistical Association** 90 (1995), p. 151–156

[Mac71] Saunders **MacLane**, Categories for the Working Mathematician, Springer 1971

[Mar$_1$75] Donald A. **Martin**, Borel determinacy, **Annals of Mathematics** 102 (1975), p. 363–371

[Mit97] Tom **Mitchell**, Machine Learning, McGraw-Hill 1997

[NatWin08] Paul **Natorp**, Wilhelm **Windelband** (*eds.*), Kant's gesammelte Schriften, Die Akademie-Ausgabe, Abteilung I: Werke, Band V: Kritik der praktischen Vernunft, 1788, Kritik der Urtheilskraft, 1793, Berlin-Brandenburgische Akademie der Wissenschaften 1908

[NieSch$_6$82] Mogens **Nielsen** and Erik M. **Schmidt** (*eds.*), Proceedings of the 9th Colloquium on Automata, Languages and Programming, Aarhus, Denmark, July 12-16, 1982, Springer 1982 [Lecture Notes in Computer Science 140]

[Pop68] Karl **Popper**, The Logic of Scientific Discovery, Harper 1968

[Put65] Harry **Putnam**, Trial and Error Predicates and a Solution to a Problem of Mostowski, **Journal of Symbolic Logic** 30 (1965), p. 49–57

[Qui69] Willard Van Orman **Quine**, Natural Kinds, *in:* [Res$_0$69, p. 5–23]

[Res$_0$69] Nicholas **Rescher** (*ed.*), Essays in Honor of Carl D. Hempel, A Tribute on the Occasion of His Sixty-Fifth Birthday, Reidel 1969

[Ros83] Roger **Rosenkrantz**, Why Glymour is a Bayesian, *in:* [Ear83, p. 69–98]

[Sch$_9$99a] Oliver **Schulte**, The Logic of Reliable and Efficient Inquiry, **The Journal of Philosophical Logic** 28 (1999), p. 399–438

[Sch$_9$99b] Oliver **Schulte**, Means-Ends Epistemology, **The British Journal for the Philosophy of Science** 50 (1999), p. 1–31

[Sch$_9$01] Oliver **Schulte**, Inferring Conservation Laws in Particle Physics: A Case Study in the Problem of Induction, **The British Journal for the Philosophy of Science** 51 (2001), p. 771–806

[Sco$_1$82] Dana S. **Scott**, Domains for Denotational Semantics *in:* [NieSch$_6$82, p. 577–613]

[Was04] Larry **Wasserman**, All of Statistics: A Concise Course in Statistical Inference, Springer 2004

Received: April 12th, 2005;
In revised version: December 2nd, 2005;
Accepted by the editors: June 6th, 2007.