

## TRUE ENOUGH

Catherine Z. Elgin

*Abstract:* Truth is standardly considered a requirement on epistemic acceptability. But science and philosophy deploy models, idealizations and thought experiments that prescind from truth to achieve other cognitive ends. I argue that such felicitous falsehoods function as cognitively useful fictions. They are cognitively useful because they exemplify and afford epistemic access to features they share with the relevant facts. They are falsehoods in that they diverge from the facts. Nonetheless, they are true enough to serve their epistemic purposes. Theories that contain them have testable consequences, hence are factually defeasible.

Epistemology valorizes truth. Sometimes practical, or prudential, or political reasons convince us to accept a known falsehood, but most epistemologists deny that there can be cognitively good reasons to do so. Our overriding cognitive objective, they maintain, is the truth, preferably the whole truth, and definitely nothing but the truth.<sup>1</sup> If they are right, then at least insofar as our ends are cognitive, we should accept only what we consider true, take pains to insure that the claims we accept are in fact true, and promptly repudiate any previously accepted claims upon learning that they are false. I suggest, however, that the relation between truth and epistemic acceptability is both more tenuous and less direct than it is standardly taken to be. Sometimes, I contend, it is epistemically responsible to prescind from truth to achieve global cognitive ends.

At first blush, this looks mad. To retain a falsehood merely because it has other epistemologically attractive features seems the height of cognitive irresponsibility. Allegations of intellectual dishonesty, wishful thinking, false consciousness, or worse immediately leap to mind. But science routinely transgresses the boundary between truth and falsehood. It smoothes curves and ignores outliers. It develops and deploys simplified models that diverge, sometimes considerably, from the phenomena they purport to represent. Even the best scientific theories are not true. Not only are they plagued with anomalies and outstanding problems, but even where they are successful, they rely on laws, models, idealizations and approximations that diverge from the truth. Truth-centered epistemology, or *veritism* as

---

<sup>1</sup>Goldman, Lehrer, Dretske,

Alvin Goldman calls it,<sup>2</sup> easily accommodates anomalies and outstanding problems, since they are readily construed as defects. The problem comes with the laws, models, idealizations, and approximations which are acknowledged to be untrue, but which are nonetheless critical to, indeed constitutive of, the understanding that the science delivers. Far from being defects, they figure ineliminably in the success of science. If truth is mandatory, much of our best science turns out to be epistemologically unacceptable and probably intellectually dishonest as well. Our predicament is this: We can retain the truth requirement and construe science either as cognitively defective or as non-cognitive, or we can reject, revise, or relax the truth requirement and remain cognitivists about, and fans of science. I take it that science provides an understanding of the natural order. By this I do not mean merely that an ideal science *would* provide such an understanding or that in the end of inquiry science *will* provide one, but that some actual science has done so and continues to do so. I take then it that much mature science is cognitively reputable. So an adequate epistemology should explain what makes good science cognitively good. Too strict a commitment to truth stands in the way. Nor is science the only casualty. In other disciplines such as philosophy, and in everyday discourse, we often convey information and advance understanding by means of sentences that are not literally true. An adequate epistemology should account for this as well. Nevertheless, I do not think that we should jettison concern for truth completely. The question is what role a truth commitment should play in a holism that recognizes a multiplicity of sometimes conflicting epistemological desiderata.

It is widely agreed that epistemic acceptability requires something like justified and/or reliable, true belief. The justification, reliability, and belief requirements involve thresholds. '[O]ne may need to be confident enough and well enough justified and one's belief must perhaps derive from a reliable enough source, and be little enough liable to be false.'<sup>3</sup> But unlike the other requirements, truth is supposed to be an absolute matter. Either the belief is true or it is not. I suggest, however, that the truth

---

<sup>2</sup>Alvin Goldman, *Knowledge in a Social World*, Oxford: Oxford University Press, 1999.

<sup>3</sup>Ernest Sosa, 'Skepticism and Contextualism,' *Skepticism*, ed. Ernest Sosa and Enrique Villanueva, Boston: Blackwell, 2000, p. 2.

requirement involves a threshold too. I am not saying that truth itself is a threshold concept. Perhaps such a construal of truth would facilitate treatments of vagueness, but that is not my concern. My point is rather that epistemic acceptability turns not on whether a proposition is true, but on whether it is true enough -- that is, on whether it is close enough to the truth. 'True enough' obviously has a threshold.

I should begin by attempting to block some misunderstandings. I do not deny that (unqualified) truth is an intelligible concept or a realizable ideal. We readily understand instances of the (T) schema:

'Snow is white' is true  $\Leftrightarrow$  snow is white

'Power corrupts' is true  $\Leftrightarrow$  power corrupts

'Neutrinos have mass' is true  $\Leftrightarrow$  neutrinos have mass

and so on. A disquotational theory of truth suffices to show that the criterion expressed in Convention (T) can be satisfied. One might of course want more from a theory of truth than the satisfaction of Convention (T), but to make the case that the concept of truth is unobjectionable, such a minimalist theory suffices. Moreover, not only does it make sense to call a sentence true, we can often tell whether it is true. We are well aware not only that 'Snow is white' is true  $\Leftrightarrow$  snow is white, but also that 'Newly fallen snow is white' is true. The intelligibility and realizability of truth, of course, show nothing about which sentences are true, or which truths we can discover. Nevertheless, as far as I can see, nothing about the concept of truth discredits veritism. Since truth is an intelligible concept, epistemology can insist that only truths are epistemically acceptable. Since truth is a realizable objective, such a stance does not lead inexorably to skepticism. I do not then deny that veritism is an available epistemic stance. But I think it is an unduly limiting one. It prevents epistemology from accounting for the full range of our cognitive achievements.

To be sure, if epistemic acceptance is construed as belief, and epistemic acceptability as knowledge, the truth requirement seems entirely reasonable. For cognizers like ourselves, there does not seem to be an epistemically significant gap between believing that  $p$ , and believing that  $p$  is true. Ordinarily, upon learning that our belief that  $p$  is false, we cease to believe that  $p$ . Moreover, we consider

it cognitively obligatory to do so. One ought to believe only what is true. Perhaps a creature without a conception of truth can harbor beliefs. A cat, for example, might believe that there is a mouse in the wainscoting, even though she lacks the resources for semantic ascent. We might therefore conclude that she believes that there is a mouse in the wainscoting without believing that 'There is a mouse in the wainscoting' is true.<sup>4</sup> In that case, the connection between believing that  $p$  and believing that  $p$  is true is not exceptionless. But whatever we should say about cats, it does not seem feasible for any creature that has a conception of truth to believe that  $p$  without believing that  $p$  is true. If epistemic acceptance is a matter of belief, acceptance is closely linked to truth. So is assertion. Although asserting that  $p$  is not the same as asserting that  $p$  is true, it seems plain that one ought not to assert that  $p$  if one is prepared to deny that  $p$  is true or to suspend judgment about whether  $p$  is true; nor ought one assert that  $p$  is true if one is prepared to deny that  $p$  or suspend judgment about whether  $p$ . Assertion and belief then seem committed to truth. So does knowledge. Whether or not we take knowledge to be equivalent to justified or reliably generated true belief, once we discover the falsity of something we took ourselves to know, we withdraw the claim to knowledge. We say, 'I thought I knew it, but I was wrong', not 'I knew it, but I was wrong'.

Being skeptical about analyticity, I do not contend that a truth commitment is part of the meanings of 'belief', 'assertion' and 'knowledge'. But whatever the explanation, such a commitment so tightly intertwines with our views about belief, assertion, and knowledge that we do best to retain it and revise epistemology by making compensatory adjustments elsewhere in the system. Once those adjustments are made, knowledge and belief turn out to be less central to epistemology than we tend to think. I do not then claim that it is epistemically acceptable to *believe* what is false or that it is linguistically acceptable to *assert* what is false. Rather, I suggest that epistemic acceptance is not restricted to belief.<sup>5</sup> Analogously, uttering or inscribing seriously and sincerely for cognitive purposes -- call it 'professing' -- is not limited to asserting. Understanding is often couched in and conveyed by

---

<sup>4</sup>I do not have strong intuitions about this case, but I do not think it is clearly wrong to say that the cat has such a belief.

<sup>5</sup>Cf., L. Jonathan Cohen, *An Essay on Belief and Acceptance*, (Oxford: Clarendon Press, 1992).

symbols that are not, and do not purport to be, true. Where such symbols are sentential, I call them felicitous falsehoods. I contend that we cannot understand the cognitive contributions of science, philosophy or even our ordinary take on things if we fail to account for such symbols.

Let's look at some cases:

*Curve smoothing:* Ordinarily, each data point is supposed to represent an independently ascertained truth. (The temperature at  $t_1$ , the temperature at  $t_2$ , . . .). By interpolating between and extrapolating beyond these truths, we expect to discern the pattern they instantiate. If the curve we draw connects the data points, this is reasonable. But data rarely fall precisely on the curve adduced to account for them. The curve then reveals a pattern that the data do not instantiate. Veritism would seem to require accepting the data only if we are convinced they are true, and connecting these truths to adduce more general truths. Unwavering commitment to truth would seem then to require connecting all the data points no matter how convoluted the resulting curve turned out to be. This is not done. To accommodate every point would be abandon hope of finding order in most data sets, for jagged lines and complicated curves mask underlying regularities. Nevertheless, it seems cognitively disreputable to let hope triumph over experience. Surely we need a better reason to skirt the data and ignore the outliers than the fact that otherwise we won't get the kind of theory we want. Nobody, after all, promised that the phenomena would accommodate themselves to the kind of theory we want.

There are often quite good reasons for thinking that the data ought not or at least need not be taken as entirely accurate. Sometimes we recognize that our measurements are relatively crude compared with the level of precision we are looking for. Then any curve that is within some  $\delta$  of the evidence counts as accommodating the evidence. Sometimes we suspect that some sort of interference throws our measurements off. Then, in plotting the curve, we compensate for the alleged interference. Sometimes the measurements are in fact accurate, but the phenomena measured are complexes only some of whose aspects concern us. Then in curve smoothing we, as it were, factor out the irrelevant aspects. Sometimes we have no explanation for the data's divergence from the smooth curve. But we may be rightly

convinced what matters is the smooth curve the data indicate, not the jagged curve they actually instantiate. Whatever the explanation, we dismiss the data's deviation from the smooth curve as negligible.

*Explanation sketches* 'present ... the general outlines of what might well be developed by gradual elaboration and supplementation, into a more closely reasoned explanatory argument.'<sup>6</sup> Because of their sketchiness, Hempel believes, they are not full-fledged explanations. For the same reason, they are often not true. The currently accepted account of the extinction of the dinosaurs is roughly this:

(E)The Earth was hit by an asteroid. The impact caused an enormous cloud of dust which blocked the sunlight, impeding photosynthesis. This disrupted the food chain. The mass starvation that followed caused dinosaurs to become extinct.

This account is, let us assume, in the right ballpark. But it is only a sketch. Not only is (E) a sketch, so is the counterpart that paleontologists accept. Extinction requires universal reproductive failure, so a true explanation of the extinction of the dinosaurs has to explain why every member of the most recent generation died without reproducing. If, as seems likely, some members of that generation died without reproducing for reasons that had nothing to do with the asteroid, an account that adduces only the effects of the asteroid does not explain why dinosaurs are extinct. Since a species can rebound from near extinction, the fact that most members of a generation die without reproducing does not guarantee extinction. (E) overstates its case. The final sentence is not true. Still, an account like (E), or at least, the counterpart that paleontologists accept, strikes us acceptable.<sup>7</sup> (E) largely explains why the dinosaurs are extinct. It is not quite true as it stands, but it affords a pretty good understanding of the phenomenon.

*Ceteris paribus laws:* Many scientific laws obtain only *ceteris paribus*. The familiar law of gravity,

---

<sup>6</sup>Carl G. Hempel, 'Aspects of Scientific Explanation', *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science*, Macmillan, 1965, 424.

<sup>7</sup> If a single factor were responsible for most of the remaining reproductive failures, scientists might be interested in it. But they are unlikely to have any interest in large numbers of individual reproductive failures that do not generalize.

$$F=Gm_1m_2/r^2$$

is not universally true, for other forces may be in play. The force between charged bodies, for example, is a resultant of electrical and gravitational forces. Nevertheless, we are not inclined to jettison the law of gravity. The complication that charge introduces just shows that the law obtains only *ceteris paribus*, and when bodies are charged, *ceteris* are not *paribus*. This is no news. ‘*Ceteris paribus*’ is Latin for ‘other things being equal’, but it is not obvious what makes for equality in a case like this. Sklar glosses it as ‘other things being normal’<sup>8</sup>, where ‘normal’ seems to cash out as ‘typical’ or ‘usual’. Then a ‘*ceteris paribus*’ law states what usually happens. In that case, to construe the law of gravity as a *ceteris paribus* law is to contend that although there are exceptions, bodies usually attract each other in direct proportion to the product of their masses and in inverse proportion to the distance between them. The law usually obtains.

This construal may work for the law of gravity, but it does not always work. For there are some laws which do not even usually hold. Snell’s law,

$$n_1\sin i = n_2\sin r$$

expresses the relation between the angle of incidence and the angle of refraction of a light ray passing from one medium to another.<sup>9</sup> As standardly stated, the law is perfectly general, ranging over every case of refraction. But is not true of every case; it obtains only where both media are optically isotropic. The law then is a *ceteris paribus* law. But it is not even usually true, since most media are optically anisotropic.<sup>10</sup> One might wonder why physicists don’t simply restrict the scope of the law: ‘For any two optically isotropic media,  $n_1\sin i = n_2\sin r$ ’. The problem is that such a restricted law would afford no insight into cases where the restriction does not obtain. Snell’s law is more helpful. Even though the law is usually false, it is often not far from the truth. Most media are anisotropic, but lots of them -- and lots

<sup>8</sup>Lawrence Sklar, ‘Philosophy of Science,’ *Cambridge Dictionary of Philosophy*, ed. Robert Audi, Cambridge: Cambridge University Press, 1999, p. 702.

<sup>9</sup> $i$  and  $r$  are the angles made by the incident beam to the normal and  $n_1$  and  $n_2$  are the refractive indices of the two media.

<sup>10</sup>Nancy Cartwright, *How the Laws of Physics Lie*, Oxford University Press, 1983, pp. 46-47.

of the ones physicists are interested in -- are nearly isotropic. The law supplies good approximations for nearly isotropic cases. So although explanations and calculations that rely on Snell's do not yield truths, they are often not off by much.

The law may be valuable for another reason as well. Sometimes it is useful to first represent a light ray as conforming to Snell's law, and later introduce 'corrections' to accommodate anisotropic media. If we were only interested in what happens to a particular light ray, such a circuitous approach would be unattractive. But if, for example, we are interested in optical refraction in general, it might make sense to start with a prototypical case, and then show how anisotropy perturbs. By portraying anisotropic cases as perturbations, we may point up affinities that direct comparisons would not reveal. The issue then is what sort of understanding we want. Showing how a variety of cases diverge from the prototypical case may contribute valuable insights into the phenomenon we are interested in. And what makes the case prototypical is not that it usually obtains, but that it cleanly exemplifies the features we deem important. *Idealizations*: Some laws never obtain. They characterize ideal cases that do not, perhaps cannot, occur in nature. The ideal gas law represents gas molecules as perfectly elastic spheres that occupy negligible space and exhibit no mutual attraction. There are no such molecules. Explanations that adduce the ideal gas law would be epistemically unacceptable if abject fidelity to truth were required. Since helium molecules are not dimensionless, mutually indifferent elastic spheres, an account that represents them as such is false. If veritism is correct, it is epistemically unacceptable. But, at least if the explanation concerns the behavior of helium in circumstances where the divergence from the ideal gas is negligible (roughly, where temperature is high and pressure is low), scientists are apt to find it unexceptionable.

*Stylized facts* are close kin of ceteris paribus laws. They are 'broad generalizations true in essence, though perhaps not in detail'.<sup>11</sup> They play a major role in economics, constituting explananda that economic models are required to explain. Models of economic growth, for example, are supposed to

---

<sup>11</sup>Graham Bannock, R. E. Baxter, Evan Davis, *Penguin Dictionary of Economics*, 1998.



explain the (stylized) fact that the profit rate is constant.<sup>12</sup> The unvarnished fact of course is that profit rates are not constant. All sorts of non-economic factors -- e.g., war, pestilence, drought, political chicanery -- interfere. Manifestly, stylized facts are not (what philosophers would call) facts, for the simple reason that they do not actually obtain. It might seem then that economics takes itself to be required to explain why known falsehoods are true. (Voodoo economics, indeed!) This can't be correct. Rather, economics is committed to the view that the claims it recognizes as stylized facts are in the right neighborhood, and that their being in the right neighborhood is something economic models should account for. The models may show them to be good approximations in all cases, or where deviations from the economically ideal are small, or where economic factors dominate non-economic ones. Or they might afford some other account of their often being nearly right. The models may diverge as to what is actually true, or as to where, to what degree, and why the stylized facts are as good as they are. But to fail to acknowledge the stylized facts would be to lose valuable economic information (for example, the fact that if we control for the effects of such non-economic interference as war, disease, and the president for life absconding with the national treasury, the profit rate is constant.) Stylized facts figure in other social sciences as well. I suspect that under a less alarming description, they occur in the natural sciences too. The standard characterization of the pendulum, for example, strikes me as a stylized fact of physics. The motion of the pendulum which physics is supposed to explain is a motion that no actual pendulum exhibits. What such cases point to is this: The fact that a strictly false description is in the right neighborhood sometimes advances understanding of a domain.

*A fortiori arguments:* Some accounts focus on a single, carefully chosen case and argue that what holds in that case holds in general. If so, it does no harm to represent the phenomena as having the features that characterize the exemplary case. *A Theory of Justice* represents people as mutually disinterested. Rawls is under no illusion that this representation is accurate. He recognizes that people are bound to one another in ties of affection of varying degrees of strength, length, and resiliency. But, he

---

<sup>12</sup>'The profit rate is the level of profits in the economy relative to the value of capital stock.'  
Ibid.

believes, if political agents have reason to co-operate even under conditions of mutual disinterest, they will have all the more reason to co-operate when ties of affection are present. I don't want to discuss whether Rawls is right. I just want to highlight the form of his argument. If what holds for the one case holds for the others, then it does no harm to represent people as mutually disinterested. That people are mutually disinterested is far from the truth. It is possible that no one on Earth is wholly indifferent to the fates of everyone else. But if Rawls is right, the characterization's being far from the truth does not impede its function in his argument.

The foregoing examples show that in some cognitive endeavors we accept claims that we do not consider true. But we do not indiscriminately endorse falsehoods either. The question then is, what makes a claim acceptable? Evidently, to accept a claim is not to take it to be true, but to take it that the claim's divergence from truth, if any, is negligible. The divergence need not be small, but whatever its magnitude, it can be permissibly neglected. We accept a claim, I suggest, when we consider it true enough. The success of our cognitive endeavors indicates that we are often right to do so. If so, a claim is acceptable when its divergence from truth is negligible. In that case it is true enough.

In practical, political or prudential contexts, both the acceptance and the acceptability of falsehoods are widely recognized. One can accept, and be right to accept the dean's latest pronouncement, if what matters is that it is the dean's opinion, not that it is true. But epistemic contexts are supposed to be different. Many epistemologists contend that when our concerns are cognitive we should accept only what we consider true. I disagree. I suggest that *to accept* that *p* is to take it that *p*'s divergence from truth, if any, does not matter. *To cognitively accept* that *p* is to take it that *p*'s divergence from truth, if any, does not matter cognitively. The falsehood is 'as close as one needs for the purposes at hand'.<sup>13</sup> In the remainder of this paper, I take 'acceptance' to mean 'cognitive acceptance'.

This raises a host of issues. Perhaps the most pressing is to say something about what I mean by 'cognitive'. A familiar line is that for a consideration to be cognitive is for it to aim at truth or be truth

---

<sup>13</sup>Robert Stalnaker, *Inquiry*, (Cambridge, Massachusetts: MIT Press, 1987), p. 93.

conducive.<sup>14</sup> Plainly, I can say no such thing. I suggest rather that a consideration is cognitive to the extent that it figures in an understanding of how things are. This is admittedly vague, but I am not sure that it is any worse than untethered remarks about truth-conduciveness and the like.

It might seem that my characterization just postpones the evil day (and not for long!), since ‘understanding how things are’ must itself be explicated in terms of truth or truth-conduciveness. To see the problem, compare three concepts -- belief, thought, and understanding. Belief aims at truth. Roughly, a belief fulfills its goal in life only if it is true.<sup>15</sup> Thought, however, can be aimless. Musings, fantasies, and imaginings can be fully in order whether or not they are true. Understanding, the argument goes, is more like belief than like thought. Since there is such a thing as misunderstanding, understanding is subject to a standard of rightness. It has an aim. Misunderstanding evidently involves representing things as they are not. This suggests that the aim of understanding is truth. If so, it may seem, divergences from truth, even if unavoidable, are always cognitive defects.

The argument goes too fast. That misunderstanding involves representing things as they are not does not entail that whenever we represent things as they are not we misunderstand them. At most it indicates that understanding is not indifferent to truth. But it does not follow that every sentence -- or indeed any sentence -- that figures in an understanding of how things are has its own truth as an objective. Understanding involves a network of cognitive commitments. It is not obvious that an aim of the network must be an aim of every, or indeed any, sentential node in the network. A goal of the whole need not be a goal of each of its parts. I don’t expect these sketchy remarks to persuade anyone that I am right to loosen the tie between understanding and truth. My hope is that they are enough to persuade you that the jig is not yet up, that a willing suspension of disbelief is still in order.

Let us turn then to acceptance. To accept that  $p$ , I said, is to take it that  $p$ ’s divergence from truth, if any, is negligible. In that case,  $p$  is true enough. Whether this is so is manifestly a contextual matter.

---

<sup>14</sup>Mike Williams, Ralph Wedgwood, *Bonjour?*

<sup>15</sup>Ralph Wedgwood, ‘The Aim of Belief’, *Philosophical Perspectives*, forthcoming; Jonathan Adler, *Belief’s Ethics*, Cambridge: MIT Press, 2002.

A sentence can be true enough in some contexts but not in others. A variety of factors contribute constraints. Background assumptions play a role. 'A freely falling body falls at a rate of 32 ft./sec.<sup>2</sup>' is often true enough, assuming that the body is within the Earth's gravitational field, that nothing except the Earth exerts a significant gravitational force on the body, that the effects of non-gravitational forces are insignificant, and so on. But even when these assumptions are satisfied, the formula is not always true enough, since gravity varies slightly with longitude. Sometimes it matters where in the gravitational field the freely falling body is. Whether 'G = 32 ft./sec.<sup>2</sup>' is true enough depends on what we want the formula for, what level of precision is needed for the calculation or explanation or account it figures in. There is no saying whether a given contention is true enough independently of answering, or presupposing an answer to the question 'True enough for what?' So purposes contribute constraints as well. Whether a given sentence is true enough depends on what ends its acceptance is supposed to serve.

Function is critical too. If to accept that  $p$  is simply to take it that  $p$ 's divergence from truth does not matter, it might seem that we accept all irrelevant propositions. None of my projects, cognitive or otherwise, is affected by the truth or falsity of the claim that Ethelred the Unready was a wise leader. So its divergence from truth, if any, does not matter to me. Since acceptance can be tacit, the fact that I have never considered the matter is not decisive. Nevertheless it seems wrong to say that my indifference makes the claim true enough. The reason is that the contention is idle. It performs no function in my cognitive economy. Owing to my indifference, there is no answer to the question 'True enough for what?'

Context provides the framework. Purposes fix the ends. Function is a matter of means. The sentences that concern us tend not to have purposes in isolation. Rather, they belong to larger bodies of discourse such as arguments, explanations, or theories that have purposes. The function of such a sentence consists in its contribution to such a body of discourse. In accepting a sentence then we treat it in a given context as performing a function in a body of discourse which seeks to achieve some end. Whether 'G = 32 ft./sec.<sup>2</sup>' is acceptable depends on whether the body of discourse it figures in serves its cognitive purpose -- whether, that is, it yields the understanding of the domain that we seek.

A statement's divergence from truth is negligible only if that divergence does not hinder its performing its cognitive function. Hence whether a contention is true enough depends not just on its having a function, but on what its function is -- on what role it plays in the account it belongs to. To determine whether a statement is true enough, we thus need to identify its function. It might seem that for cognitive purposes only one function matters. If the criterion for felicity is being true enough, one might think, the function of all felicitous falsehoods is to approximate. There is, as it were, a tacit 'more or less' in front of all such claims. This will not do.

One reason is that the proposal is not sufficiently sensitive. Not all approximations perform the same function. Some are accepted simply because they are the best we can currently do. They are temporary expedients which we hope and expect eventually to replace with truths. We improve upon them by bringing them closer to the truth. Such approximations are, in Sellars's terms, promissory notes that remain to be discharged. The closer we get to the truth, the more of the debt is paid. They are, and are known to be, unsatisfactory. But not all approximations have this character. Some are preferable to the truths they approximate. For example, it is possible to derive a second order partial differential equation that exactly describes fluid flow in a boundary layer. The equation, being non-linear, does not admit of an analytical solution. We can state the equation but not solve it. This is highly inconvenient. To incorporate the truth into the theory would bring a line of inquiry to a halt, saying in effect: 'Here's the equation; it is impossible to solve.' Fluid dynamicists prefer a first-order partial differential equation, which approximates the truth, but admits of an analytical solution.<sup>16</sup> The solvable equation advances understanding by providing a close enough approximation that yields numerical values which can serve as evidence for or constraints on future theorizing. The approximation then is more fruitful than the truth. There is no hope that future inquiry will remedy the situation, for it is demonstrable that the second order equation cannot be solved numerically, while first order equation can.<sup>17</sup> That reality forces such a choice

---

<sup>16</sup>Margaret Morrison, 'Models as Autonomous Agents,' *Models as Mediators*, ed. Mary S. Morgan and Margaret Morrison, Cambridge: Cambridge University Press, 1999, pp. 56-60.

<sup>17</sup>Conceivably, of course, the equations in question will be superseded by some other understanding of the subject, but the fact that the equation we consider true does not have an analytical

upon us may be disappointing, but under the circumstances it does not seem intellectually disreputable to accept and prefer the tractable first order equation. One might say that acceptance of the first-order approximation is only practical. It is preferable merely because it is more useful. This may be so, but the practice is a cognitive one. Its goal is to understand fluid flow in boundary layers. In cases like this, the practical and the theoretical inextricably intertwine. The practical value of the approximation is that it advances understanding of a domain. A felicitous falsehood thus is not always accepted only in default of truth. Nor is its acceptance always 'second best'. It may make cognitive contributions that the unvarnished truth cannot match.

Moreover, not all felicitous falsehoods are approximations. Idealizations may be far from the truth, but not epistemically inadequate on that account. Political agents are not mutually disinterested. They are not nearly mutually disinterested. Nor is it the case that most political agents mutually disinterested. There is, no way I can see to construe Rawls's model as approximately true. Nonetheless, for Rawls's purposes, the characterization of political agents as mutually disinterested is felicitous, if the features it highlights are constitutive of fair terms of cooperation underlying the basic structure of a democratic regime. There is no reason to think that in general the closer it is to the truth, the more felicitous a falsehood.

I suggest that felicitous falsehoods figure in cognitive discourse not as mistaken or inaccurate statements of fact, but as fictions. Before I can make my case, though, I need to supply another piece of the puzzle -- exemplification. To exemplify is to highlight, exhibit, display or otherwise make manifest.<sup>18</sup> A commercial paint sample consists of a patch of color on a card. The patch is not merely an instance of the color, but a telling instance -- an instance that exemplifies the color. By so doing, the sample equips us to recognize the color, to differentiate it from other similar shades. The sample then affords epistemic access to the color. Although the patch on the sample card has a host of other features -- size, shape, solution provides no reason to think so. Nor does it provide reason to think that the considerations that supersede it will be mathematically more tractable, much less that in the long run science will be free of all such equations.

<sup>18</sup>See Goodman, *Languages of Art*.

location, and so on -- it standardly does not exemplify them. Exemplification is selective. It brings out some features of an exemplar by overshadowing, downplaying, or marginalizing others. Nothing in the nature of things dictates that the patch's color is worthy of selection, but its shape is not. What, if anything, an item exemplifies depends on its function. The very same item might perform any of a variety of functions. The patch on the sample card could be used to teach children what a rectangle is. In that case it would exemplify its shape, not its color. The sample card could be used as a fan. Then the patch would not exemplify at all.

Exemplification is not restricted to commercial and pedagogical contexts. Whatever an item exhibits, highlights or displays, it exemplifies. A poem might exemplify its rhyme scheme, its imagery, its style. A water sample might exemplify its mineral content, its flavor or its impurities. Exemplification, I have argued, is ubiquitous in art and science.<sup>19</sup>

Treating paint samples as paradigmatic exemplars may encourage the idea that exemplified features are all like expanses of color -- homogeneous qualities spread out before us, lacking in depth and complexity, hence able to be taken in at a glance. Many are not like that. Pick up a rock containing iron ore. It might serve as a sample of iron, or of hematite, or of something dug out of the Black Hills, or of something that bears a striking resemblance to your high school algebra teacher. It can exemplify such features only where certain background assumptions are in place. Not just anyone looking at the rock could tell that it exemplified these features.

Moreover, although in principle any item can serve as an exemplar and any feature can be exemplified, sometimes a good deal of effort is required to bring about the exemplification of a recondite feature. Some of that effort is mental. Just as we ignore the shape of the paint sample and focus on the color, we can ignore the fact that the rock looks like your algebra teacher and focus on its hardness. This is a start. But some irrelevant features so intricately intertwine with relevant ones that more drastic measures are called for. If we seek to exemplify some recondite feature of iron, mental agility alone may

---

<sup>19</sup>See *Considered Judgment*, Chapter 6.

not be enough bracket the effects of other minerals in the rock. So we refine the ore and filter out the impurities. The result of our efforts is pure iron. It is the product of a good deal of processing<sup>20</sup> which eliminates complicating factors and brings to the fore characteristics that are hard to detect and difficult to measure in nature. To facilitate the exemplification of the feature of interest, we do not just mentally sideline features we consider irrelevant, we physically remove some of them.

Even then, we do not just contemplate the bit of iron as we might a paint sample. We subject it to a variety of tests. We seek to produce circumstances where the features of interest stand out. We not only investigate the iron's behavior in standard conditions, we study what happens in extreme conditions -- very high or low pressure or temperature, in a vacuum, under intense radiation, and so forth. Although we recognize that the test conditions do not ordinarily (or perhaps ever) obtain in nature, we take it that the behavior of the refined metal in the test conditions discloses something about the natural order. If so, by understanding what happens in the lab, we can understand something of what happens in the world. The connection is, of course, indirect. It involves a complicated extrapolation from situations and materials that are highly artificial and carefully contrived. One might argue that the lab itself is a fictional setting and the conclusions we draw about nature on the basis of our findings in the lab are projections from fiction to fact. I don't quite want to say that (although I suspect that Nancy Cartwright does.) But I do want to point out that experimentation involves a lot of stage setting.

There is a tendency to think of experiments as processes that generate information, hence as ways to find things out. This of course is true. But it is worth noting that an experiment is not like an oracle,<sup>21</sup> or an anchorman or a fortune cookie. It does not just issue a report that states its results. It displays them. It shows what happens to the magnetic properties of iron in conditions near the melting point. The experiment exemplifies its results.

---

No matter how carefully we set the stage, irrelevancies remain. We do not and ought not then

<sup>20</sup>This is why Nancy Cartwright thinks the laws of physics lie. The laws are developed on the basis of, and are strictly true only of the processed samples, not their naturally occurring counterparts.

<sup>21</sup>Actually, of course, I don't know how oracles are supposed to operate. I have always assumed that they simply emit true sentences like 'Socrates is the wisest of men'.



read every aspect of the experimental result back onto the world. Not only are there irrelevant features, there are issues about the appropriate vocabulary and level of precision for characterizing what occurs. The fact that the experiment occurred in Cleveland is unimportant. The fact that the iron has a certain mass or lattice structure may or may not be significant. The fact that the temperature is within 100 °C of the melting point of iron may matter, while the fact that it was within 5 °C of 1500 °C does not. Some features of the iron in the experimental situation are telling features. Others are not. The telling features are the ones that the experiment discloses, or makes manifest. By exemplifying certain features, the experiment brings them to light and affords epistemic access to them. That is its cognitive contribution. Other features, equally real, are not exemplified. The experiment embodies an understanding of the phenomenon in question through its exemplification of telling features. By making these features manifest it affords an understanding of the phenomenon.

Now if the cognitive contribution of an exemplar consists in the features it exemplifies, then anything that exemplifies exactly those features can, in a suitable context, make the same contribution. Return for a moment to the paint sample. I have been speaking of it as though it is a sample *of paint*, a telling instance of the stuff you might use to paint the porch. This is not true. The sample does not consist of paint, but of an ink or dye of the same color as the paint whose color it exemplifies. If the sample were supposed to exemplify the paint itself, or the chemical features of the paint, the fact that it is not paint would be objectionable. But since it only exemplifies the color, all that is needed is something that is the same color as the paint. The exemplar need not itself be paint. Similarly, in scientific cases. Consider a DNA molecule that exemplifies its molecular structure. Anything that exemplifies the same structure has the capacity to perform the same function in our understanding of DNA. No more than the paint sample need consist of paint, does the exemplar of DNA's molecular structure need to consist of DNA. A schematic model that exemplifies the same features but has a different material (or even immaterial) substrate could do the job. Here is where felicitous falsehoods enter the picture. Something other than paint can serve as a paint sample, affording epistemic access to a color also instantiated by

paint. Something other than a molecule can exemplify molecular structure, affording epistemic access to a structure also instantiated by the molecule. A felicitous falsehood then is a fiction that exemplifies a feature in a context where the exemplification of that feature advances understanding. The utility of such a falsehood is plain. It is sometimes inconvenient, difficult, or even impossible to bring it about that all and only the features that interest us are exemplified by the phenomena of interest. (DNA molecules are very small, charged pions are short lived.) If we introduce a falsehood that exemplifies those features -- a bigger, longer-lasting model, for example -- we can highlight them, and display their significance for the understanding of the phenomenon in question. The camel's nose is now officially inside the tent.

There is more than one role that such fictions can play. Some serve as points of reference. We understand things in terms of them. In the simplest cases, like the model displaying the helical structure the DNA molecule, they are simply schemata that exemplify factors they share with the phenomena they concern. They qualify as fictions because they diverge from the phenomena in unexemplified properties. (DNA molecules are not made of pipe cleaners.) In other cases the connection to the facts is less direct. No real gas has the properties of the ideal gas. The model is illuminating though because we understand the properties of real gases in terms of their deviation from the ideal. In such cases, understanding involves a pattern of schema and correction. We represent the phenomena with a schematic model, and introduce corrections as needed to bring it into closer accord with the facts. The fictional ideal then serves as a sort of least common denominator that facilitates reasoning about and comparison of actual gases. We 'solve for' the simple case first, then introduce complications as needed.

Acknowledging the role of corrections might seem to suggest that the detour through fictions is just a circuitous route to the truth. Rather than a simple true description of the behavior of neon, we give a complicated truth that makes reference to deviations from some ideal. But the cognitive content of the exercise resides entirely in the truth. I don't think this is right for a couple of reasons. The first is that sometimes the corrections that would be needed to yield a truth are unnecessary or even counter-productive. A fortiori arguments succeed because the corrections don't matter. If a consideration holds

for one case, even if that case is a fiction, it holds for all. If Rawls' argument is sound, 'correcting' for ties of affection just muddies the waters. In other cases fidelity to the facts may prove a hinderance. It follows from the ideal gas law that pressure becomes infinite as volume goes to zero. This would not happen. Given a fixed number of molecules, pressure increases as volume decreases -- not to infinity, but only to the point where the container explodes. No one of course denies this. But to understand what would happen in the limit, we need to prescind from such material concerns and pretend that the walls of the container are infinitely strong. We need then to introduce not corrections that bring us back to the facts, but further idealizations. The second reason is that even where the corrections are needed, the fiction may be more than a *façon de parler*. It can structure our understanding in a way that makes available information we would not otherwise have access to. If, e.g., we draw a smooth curve that skirts the data and construe the data as a complex of relevant and irrelevant factors (noise), or construe a transaction in terms of an economic model overlaid with non-economic factors which skew the outcome, we impose an order on things, highlight certain aspects of the phenomena, reveal connections, patterns and discrepancies, and make possible insights that we could not otherwise obtain. We put ourselves in a position to see affinities between disparate occurrences by recognizing them as variations on a common theme.

Explanation sketches like the one I mentioned earlier are incomplete in a different way. They exemplify some aspects of things by leaving other factors out. The explanation sketch of the extinction of the dinosaurs is not true, for it fails to account for some reproductive failures among the last generation of dinosaurs. But it seems to be a good explanation anyway. The reason is that it highlights the changes in the environment that tipped the ecological balance. Although there were surely other causes of reproductive failure, there was presumably nothing unusual about them. Such things had been happening all along without endangering survival of the species. The sketch omits mention of ordinary reproductive failures precisely because they were ordinary. They were part of the background against which the asteroid had its effect. It might seem that we could evade the charge of falsehood if we simply embedded

the sketch in something like ‘Besides the usual reproductive failures . . .’ But it is worth noting that we don’t think we need to do any such thing. We unblushingly consign the ordinary to the unarticulated background in order to highlight what was new.

In dealing with the sample we discriminate between factors that matter and those that do not. The ones that matter are exemplified by the sample. By studying them, we advance our understanding of the chemical features of iron. But if all we care about is the exemplified features, we should be happy with anything that exemplifies those features, whether or not it was originally a bit of iron ore. This is where fictions come in. We could simply create an entity that exemplifies the important features of iron ore and see how that entity behaves under various tests. The creation would not be an instance of iron, but it would display the key features and afford epistemic access to those features. This is the way a model can work. Strictly it is a schema or a simulation, but if the features it exemplifies are features of the phenomenon, it reflects or embodies an understanding of the phenomenon. That it has other features that it does not share with the phenomenon is neither here nor there. For if we know how to interpret the model correctly, we do not project the unexemplified features back onto the phenomenon. The model then need not even *be* iron. It need only exemplify the same features as are important to the understanding of as it might be, the magnetic features of iron.

Moreover, there are Gricean reasons for thinking that framing our explanation that way would be counterproductive. Once we mention the usual reproductive failures, they acquire a salience, and perhaps seem in need of an explanation of their own. Grice’s second maxim on quantity is ‘Do not make your contribution more than is required’ for the current purposes of the exchange.<sup>22</sup> If what is required for the

---

<sup>22</sup>Paul Grice ‘Logic and Conversation’, *Studies in the Way of Words*, Cambridge: Harvard, 1989, p. 26.

current purposes of the exchange is information about what changed in the reproductive circumstances of the dinosaurs, omitting mention of circumstances that were unchanging is not only appropriate but required. But as is well known, utterances that satisfy the Gricean maxims are often false.

[short interlude about fictions]

A falsehood is acceptable only if its divergence from truth is negligible. Therefore a falsehood's acceptability is discredited if its divergence is non-negligible. To know whether this is so requires not only knowing where and how much it diverges from truth, but also, how much divergence would be negligible. This depends on its function or role in the theory. Felicitous falsehoods then have defeasibility conditions. Since an ordinary factual statement purports to be true, a single counterexample discredits it. So if in a given context, a single counterexample is sufficient to discredit a claim, that claim functions as a statement of fact. A generality purports to be true in most cases. One counterexample does not discredit it, but a sufficient number of counterexamples do. An approximation purports to be within some  $\delta$  of the truth. If (enough of) its instances fall outside that range, it is discredited. Models admit of a variety of defeasibility conditions. Some purport to be accurate in relevant respects, or others in important regions, or others for a significant range of cases. The models are discredited if they are not sufficiently accurate in the cases in question. By establishing what it would take to undermine a claim then, we establish what role the claim is playing. Such propositions do not, of course, possess defeasibility conditions in isolation. The issue is what would discredit them in the cognitive contexts in which they are found. Although there may be some vagueness about the matter, standardly the answer is reasonably determinate.

A major worry with this account is that my position makes the world safe for postmodernist claptrap. If truth is not required for epistemic acceptability, why isn't a flagrantly false account acceptable? Why not say that theories attesting to the healing powers of crystals are as acceptable as mainstream crystallography? My seemingly wimpy requirement that an acceptable account must yield an

understanding of how things are provides the answer. An account that yields such an understanding must accommodate the facts in a domain. The accommodation may be indirect. Strictly false idealizations may be deployed. Detours through stylized facts may be made. Etc. The justification for the falsehoods is that they figure in accounts that make sense of the facts. A cognitively acceptable account sheds light on its subject. Where felicitous falsehoods are involved the light may be oblique.

A theory can only claim to make sense of a range of facts if it is factually defeasible -- if that is, there is some reasonably determinate, epistemically accessible factual arrangement which, if it were found to obtain would discredit the theory. A felicitous falsehood is acceptable only if the theory or system of thought it belongs to accommodates the (epistemically accessible) facts. Exactly what this requires needs to be spelled out. The usual considerations about evidence, simplicity, scope, etc. come into play. Even though some of the sentences in a theory are not supposed to be true, the way the world is constrains the acceptability of the theory.

What I need to do:

A. discuss some of the ways felicitous falsehoods advance understanding:

1. Schema -- Gombrich
2. Exemplification
3. Provide a scaffolding -- Stalnaker
4. Gricean stuff on do not be prolix. We bracket confounding factors in order to say only what is necessary (Note 'say' is wrong here). If one were even to say 'That's the main cause of dinosaur extinction' rather than 'That's the cause ...' it would raise the issue of what the other causes are.

For the purposes at hand, we have reason to leave that issue aside.

B. You have to know how to read:

1. Be able to tell which sentences do not purport to be true
2. Be able to tell how they function.

C. The bodies of discourse, systems of thought, or theories have to be defeasible.

I want to understand and justify the patterns of understanding. The cognitive web (which I can't any longer call a web of belief) contains nodes that are not and do not purport to be true. We accept those nodes because they figure in networks that constitute understanding. The network as a whole has to be testable and defeasible. There must be some identifiable situation which would discredit the network. Such a situation, does not directly tell against the falsehoods in the set. It only tells us that something is amiss. The falsehoods -- the stylized facts or idealizations or simplifications -- are at least as ripe for revision as the sentences that purport to be true.

Constellations afford a vivid example. They are fictional orders imposed on the stars. Our ancestors connected the dots in the sky to form a hunter or a bear or a big dipper. This is clearly a fanciful projection. You could connect dots and get a hunter or you could connect dots and get Fenway Park. The problem is not just that there are more potential figures than anyone has bothered to draw. The conviction that the stars really do instantiate the patterns in the traditional constellations does not withstand much scrutiny. Orion's belt is supposed to consist of three collinear stars. They are not really collinear though, or even close to it. Because they are so far away, differences in depth are not apparent to the naked, Earthbound eye. From a different celestial perspective, or with greater magnification, they would not even seem to line up. So we may be inclined to consign the constellations to the superstitious childhood of the race, and dismiss them as cognitively valueless. There is a problem, though. Not only are the constellations used effectively in celestial navigation, they are also used in astronomy. 'Current astronomy divides the sky into regions associated with prominent constellations. These constellation based boundaries provide a standard scientific map for locating objects in the heavens.'<sup>23</sup> Astronomers use the fictions to orient themselves toward the facts.

0. Our ordinary philosophical view is that pragmatics is somehow on the side: There's the knowledge, and there's what you do with it. I am suggesting that because Gricean considerations infuse our thinking

---

<sup>23</sup>Robert Schwartz, 'Starting from Scratch: Making Worlds,' *Erkenntnis*, 52 (2000), p. 155.

about the domain, the ‘what you intend to do with it’ stuff is integral to the knowledge or understanding we manage to get.

0'. Felicitous falsehoods are nodes in the net that facilitate comparison, by identifying a common factor.

0' Make it clear that the understanding that science produces cannot simply factor out the falsehoods. The laws, etc. figure ineliminably in the understanding. You can't then do a sort of Lehrer move of saying that the scientifically justified stuff is justified in the ‘total science system’ and the verific system that results when all falsehoods are removed.

One is by providing a frame of reference or an orientation for locating facts. The system of latitude and longitude is a scheme of imaginary lines projected onto the Earth. They enable us to locate any position on the planet and equip us with resources for stating a host of truths about geography, astronomy, meteorology and history. They constitute a framework for stating and investigating matters of fact. Although they are purely conventional, that does not undermine their utility. I don't quite want to say that statements that make reference to latitude and longitude are false. I think this is a very delicate question involving the interpretation of ascriptions of mathematical features to material objects.<sup>24</sup> I bring up this case because it is a borderline case. Still, it is telling that latitude and longitude are characterized in reference books as *imaginary* lines. It is at least not implausible to think of them as fictions which facilitate our thinking about, describing, and investigating matters of geographical fact by providing a schema for fixing location.

1. We want a more finely textured understanding than the characterization of felicitous falsehoods as approximations provides. For once we give up the idea that a good theory is simply true, that it mirrors its

---

<sup>24</sup>The question is whether something material can possess purely mathematical, that is immaterial features. The answer might be ‘yes’, but it would take some difficult metaphysics to explain how. The answer might be ‘literally no’. In that case when we say that a material object has mathematical features we are speaking metaphorically.



domain, or whatever, we need to understand not only that it represents its domain and is approximately true of its domain, we need to know how it represents the domain, and this requires understanding how and where it diverges from the facts.

2. Re stylized facts in economics. Stylized facts facilitate comparisons. A straight explanation of the profit rate in Senegal might have to take account of the effects of the drought. An explanation of the rate in Chechenya would have to account for the effects of the war. An explanation of the rate in Botswana would have to factor in the effects of the AIDs epidemic and so on. The complications are significant, and the complicating factors different. By factoring out the complications that war, disease, and drought present, economists can discover common economic patterns. Obviously, they can't set non-economic factors permanently aside. But to structure an understanding of the domain in terms of a streamlined economic schema augmented with noneconomic corrections<sup>25</sup> facilitates the ends of the discipline. An economic model prescind from the effects of political, medical, and meteorological phenomena to concentrate on purely economic factors. The description it yields is not true, for the fact that, e.g., the president for life ran off with the national treasury is apt to have non-negligible effects on the state of the national economy. To account for the actual state of the economy, then would require introducing corrections to the explanation that the model yields. The schema provides a bare bones description; the corrections put meat on the bones. By accounting for the real by reference to the stereotype, it facilitates comparisons. It shows how different situations, resulting from different causal forces (war, pestilence, corruption, etc.) are variants on the same economic structure. One could, of course, account for the same situation using a political model instead. Then different commonalities and differences would be brought to the fore. A straight account that gives each significant causal factor its due might well explain how Canada came to be in the economic situation it is in. But because the significant non-economic factors are apt to vary considerably from one case to the next, straight stories about each of a number of economies might make it look as though they had nothing in common. If they can all be shown to be

---

<sup>25</sup>On schema and corrections, see Gombrich, *Art and Illusion*, pp.

variants on the same schema, the common factors are brought to light.

3. I have not said that one need entertain  $p$ . Nor have I said that one need be expressly aware of one's attitude toward  $p$ . But this does not jeopardize my account though for I do not *take it* that the proposition's divergence from truth, if any, does not matter. To take it that  $p$ , is to have a policy or disposition to use  $p$  as a premise.<sup>26</sup> Since we are concerned with cognitive acceptance, to take it that  $p$  is to have a policy of treating  $p$  as a premise in reasoning whose goal is to understand how things stand. 'Premise' has to be widely construed. Acceptance can be tacit, so the premise that  $p$  may be presupposed rather than expressly stated. It may, moreover, function as a constraint on the shape or scope of the argument rather than as a 'line' in it. Thus, for example, acceptance of stylized facts in economics requires that explanations proceed via them, rather than bypassing them.

4. Manifestly, 'true enough' does not have the logical properties of 'true'. The conjunction of two truths is true, but the conjunction of two true enough propositions is not always true enough. The result is a compartmentalization.

5. Let us say that to utter seriously and sincerely for cognitive purposes is to profess.

6. As the foregoing examples show, sometimes this is so. When a sentence's divergence from truth does not matter, it is, I want to say, true enough. First, acceptance is context sensitive.  $p$ 's divergence from truth might matter in some contexts but not in others. Second, it is relative to purposes or functions. The role a sentence or other symbol plays in a system of thought is critical to deciding whether its divergence from truth matters. For example, 'Amherst is 90 miles from Cambridge,' is not true. It is, however, close enough to be useful if we are considering how long it would take to drive there. It is not close enough if our HMO treats injuries outside of a 90 mile radius differently from injuries within that radius. A representation that represents extended bodies as point-masses may be acceptable for some purposes.

---

26L. Jonathan Cohen, *An Essay on Belief and Acceptance*, Oxford University Press

But for others, the actual size and shape of the body matter.

7. In a context where  $p$  is true enough, it can be used as a premise<sup>27</sup>, or a presupposition, or a constant. (E.g., sentence that is true enough can be used in explanation. It can even be treated as a statement of fact that needs to be explained (example: idealized facts). Moreover, it is also a reasonable basis for action (driving to Amherst is or is not reasonable based on the 90 mile reading.)

8. This by itself says nothing about how a true enough sentence might diverge from the truth or how far it might diverge from the truth. In fact both the character and the degree of divergence vary considerably. If we are just concerned with whether an inference form is truth preserving, we can simply pretend that any sentences we choose are true, and see what happens. The only constraints then would be at the sentences be self-consistent and logically independent of each other. If all we care about is the structure, or that structure at a given level of abstraction, any sentence that is isomorphic with the truths we care about might do. Sometimes vague or partially interpreted sentences may be true enough, if they are true or even true enough in the realm we are concerned about.

9. One consequence of the ideal gas law is that pressure becomes infinite as volume goes to zero. This would not happen. To be sure, given a fixed number of molecules, pressure increases as volume decreases -- not to infinity, but only to the point where the container explodes. No one of course denies this. But to understand the way the idealization functions, we need to prescind from such material concerns and see what would happen if the walls of the container were infinitely strong. We need, then to introduce a further idealization.

10. [Felicitous falsehoods are all not of a piece. A contention suited to one such role is not thereby suited to others. Rawls's principle functions well as an *a fortiori* principle, but would be a poor stylized fact,

being far from the truth.]

11. But, inconveniently, there seem to be multiple reasons for acceptance. Some falsehoods are favored because they are good approximations. Perhaps current science cannot do any better. Perhaps, and more interestingly, current science can produce a truth, but the cognitive cost is prohibitive. For example, we might come close to the truth with a first order differential equation that has an exact numerical solution, or express the actual truth with a higher order equation that does not. It does not seem like a sell-out to opt for an equation we can solve. Some falsehoods are accepted because they largely account for the phenomenon in interest, but does not entirely do so. It may account for most instances, or account for the dominant factors. Some falsehoods describe telling instances. They are not true in every case (indeed, it may not be true in any real case), but to one extent or another, actual cases can be understood in terms of their divergence from the telling instances.

12. Acceptability will have to be not that  $p$ 's divergence from truth does not matter, but that one rightly takes it that  $p$ 's divergence from truth does not matter.

13. shed light on, sometimes from an oblique angle.

14 We want a more finely textured understanding than the characterization of felicitous falsehoods as approximations provides. Once we give up the idea that a good theory is simply true, that it mirrors the phenomena, or whatever, we need to understand not only that it represents the phenomena and is true enough them, we need to know how it represents the phenomena, and this requires understanding how and where and to what effect it diverges from the truth

15 In the realm of fiction, falsity is unobjectionable. So if we can assimilate felicitous falsehoods to fictions, we can begin to make some headway. One might think that this is a non-starter, for the reason that falsity is unobjectionable in fiction is that fiction is frivolous. It really is, in Wittgenstein's terms,

'language on holiday'. If the sentences comprising a work of fiction are not put forth seriously, it hardly matters that they are false. The trouble is that much great fiction -- and a lot of lesser fiction -- is far from frivolous. Whatever accounts for the power of *Crime and Punishment* or *The Oresteia*, frivolity is unlikely to play a role. The sentences in such works are not literal, declarative statements of fact. But it does not follow that they are not put forth seriously. Rather it raises the question what serious roles besides literal fact stating sentences can play.

16. What needs to be done.

More on what exemplification is.

Argue that the experiment or other exemplar does not merely produce or cause understanding, it embodies and expresses understanding.

Argue that only the exemplified features figure in or contribute to the understanding the experiment or whatever embodies. The unexemplified features may be a necessary substrate, but they do not figure in the understanding. If this is right then anything that exemplifies the same features does equally well. That there is a different substratum or none at all, makes no difference.

Now, a felicitous falsehood can exemplify the same features. Hence it may be as useful as the experimental display. (Note this is true for the embodiment or manifestation of knowledge. A slightly or intricate story must be told about the context of discovery, although I would argue that manipulating a model can yield discoveries, just as manipulating material in the lab can.)

Consider now, not a substitution, but a model or other ff that plays a role that we cannot match in the lab. E.g., a model might manifest what happens to a transistor at absolute zero. We can approximate, but not reach this in the lab. A model reveals that as we compress a gas, pressure increases to infinity. Obviously in any real case where we kept increasing the pressure, the container would simply explode.

A plant cell might exemplify plasmolysis.

A critical feature of exemplification is that it is selective. An exemplar exemplifies only some of its features. The paint sample is rectangular, but does not exemplify its shape. The Michelson-Morley experiment was first performed in 1887, but does not exemplify its date of origin.

Every item belongs to a huge number of extensions. Let us call that which the members of an extension share a feature.<sup>28</sup> Then every item has a huge number of features. It is pointless to attempt to enumerate all the features an item has, all the relations it figures in, all the processes it undergoes or contributes to. Not only would the task be impossible to complete; even a first stab would yield such a vast and disorderly jumble of data that we could make no sense of it. In any case, most of the features a thing instantiates are of no interest whatsoever. Of the remainder, some are of interest for some purposes, others for others. Some objects do not merely instantiate features, they also exemplify them. Consider a commercial paint sample of the sort you find in a hardware store. It consists of a colored patch on a card. Its function is to display a particular color of paint. It does so, not merely by being an instance of that color, but by being a telling instance, an instance that brings the color to the fore and makes it manifest. The patch, of course, is an instance of a host of other features as well -- size, shape, location, and so on. But it makes nothing of them. The patch exemplifies its color, not its shape or size. By exemplifying the color, by displaying it or making it manifest, the sample affords epistemic access to it. It shows the color and equips us to differentiate it from other similar shades.

Exemplification is not restricted to commercial samples. It is the mechanism by which samples, examples, and other exemplars point up and afford epistemic access to the features that they do. An example in a textbook might exemplify a reasoning strategy or a poetic form. *The Rite of Spring* exemplifies atonality. The Michelson-Morley experiment exemplifies the finitude of the speed of light. Exemplification, I have argued, is ubiquitous in art and science.<sup>29</sup>

---

<sup>28</sup>This notion of a feature is akin to David Lewis's notion of a property. Cf., 'New Work for a Theory of Universals' [I think]. I use the term 'feature' rather than 'property', since I do not want to get embroiled in what are for my purposes needless disputes about the nature of properties. 'Feature' in my usage is a term of art that designates what the members of an extension have in common.

<sup>29</sup>See *Considered Judgment*, Chapter 6.

Goodman uses paint chips as a paradigm case of exemplification. This may encourage the idea that exemplified features are like expanses of color -- homogeneous qualities spread before us, lacking in depth and complexity, hence able to be taken in at a glance. Many features are not like that. The feature of that all bits of iron share is something enormously complex, which takes a good deal of effort to make manifest. .

[Through its display of these features the experiment embodies and expresses a partial understanding of the phenomenon