Are There Objective Values?: The Greatest Happiness Principle in Perspective

by Arnold Faden

This is the presentation of a whole new philosophical system—ethics, ontology, epistemology, politics, the works. Hence I cannot give too many details here, but just the overall vision, with some indications of the grounds for it and some consequences.

Alternate visions of the good compared

One major thesis is that values are in some sense "objective." But which value system is the correct one? I believe it is the Greatest Happiness Principle (GHP) (better known as utilitarianism, a word I dislike because of its misleading connotations). The GHP enjoins one to maximize the expected net happiness of the universe. Here "net" refers to the balance of happiness (joy, pleasure, bliss) over unhappiness (suffering, pain, agony), "expected" refers to expectation in the sense of probability theory, and "universe" means that the interests of all sentient beings are to be considered, no matter where, when, or in what manner they exist.

But how does one carry out such an injunction? Doesn't it require information that is not available? Doesn't it require an impossible amount of calculation? (That the GHP is "non-operational" is one of the standard objections to it.) The answer lies in the concept of **probability**, a key to understanding both epistemology and ontology.

Probability is often said to come in two forms: "subjective," referring to someone's degree of belief in a proposition, and "objective," referring to causation, scientific laws, the propensity of some properties to be associated with others. This is not correct: there is only one probability; it embraces both of these aspects. $P(S)$ is the credibility a judging agent assign to proposition S at a certain time. Propositions come in ranges, S_1, S_2, ... such that exactly one of the S_i is true. (Think of a question and the set of possible answers to it; in probability theory, these ranges are called random variables.) Random variables typically come in families, indexed say by time, place, or entities, e.g. the weather at times 1,2,..., or national income of countries, 1,2, ...

Now times, places, and entities have part-whole or **hierarchical** relations among themselves (short time stretches into long, persons into groups, counties into states, etc.). Indexing families of random variables to these induces a corresponding hierarchy among random variables. (The typical statistical model is a one-step hierarchy. For example, let X_n be the outcome of the n-th toss of a certain coin, the range being {head, tail}; the distribution of the X's is governed by θ, a variable indexed to the coin itself, say its propensity to fall heads; the range of θ is the interval [0, 1]. A more inclusive model might distinguish between θ_i for different coins, their joint distribution being governed by a variable attached to coins-in-general, and so on.)

Probability distributions over lower levels (the X's) are often called "objective," those over higher levels (the θ's) "subjective," but they are all the same kind of thing. Probabilities are not relative frequencies, though the latter when observed are evidence which should modify the probabilities. (Actually, any evidence whatever may be construed in terms of frequencies, or, more generally, measures. See Faden, Economics of Space and Time.)

The basic rule of inference is this: if P(S) is my probability assignment, and I observe E, then P(S) gets modified to P'(S) = P(S | E), the old probability conditioned on the evidence E. Usually, it is convenient to express this as follows: Given a propositional range $S_1, S_2, ...$, multiply the "prior" probability $P(S_n)$ by the "likelihood" $P(E | S_n)$; these products are then proportional to the new "posterior" probabilities $P'(S_n)$. In words, the alternatives are enhanced proportionately to how well they explain the evidence E.

How far over space, time, or entities does a given probability law extend? This is the problem of induction, and itself may be governed by random variables at a higher level (e.g. all known crows are black, hence—with high probability? -- all crows are black). The induction problem is solvable by what is called reliability theory, which I won't elaborate on here.

Note that certainty is a special case of probability (P = 0 or 1), just as rest is a special case of motion. Thus, strict causality (propensities of 0 or 1) is included here.

The probabilistic framework just outlined is, in principle, the correct way for dealing with uncertainty. But how does it compare with the way we actually do handle information? A huge literature indicates that we are subject to "heuristics, biases, and illusions," but no consensus exists on the extent of this irrationality. (Represent this by an "irrationality" random variable with range from 0 to 100% !)

Some general comments are in order.

(i) We don't run around probabilities in our heads (P(rain tomorrow) = 0.1763). But our language is full of semi-quantitative phrases (e.g. "barely possible," "likely," "beyond a reasonable doubt") indicating that we do make such judgments. (And how much exactness is worthwhile? After all, outside of the realm of pure mathematics, all our concepts have a fuzzy penumbra, so that we are not quite sure what our propositions mean. There is a theory of "negative significant figures" which covers this situation.)

(ii) The information processing capacity of the human mind is limited and thinking is costly. We economize in several ways. First, by having a limited range of concepts (few random variables). (Think of experts vs. amateurs; we are all amateurs in most areas.) Second, by truncating random variables, ignoring or overlooking possible answers to our questions. Thirdly, by excessive certainty: Probabilities constitute an extra layer of complexity, which we can avoid by fixing on one possibility and believing it. Fourthly, by focusing our attention on a small number of variables of interest, pushing the rest into the background.

(iii) The mind is largely geared to action, not information registry, and emotion plays a major role in swaying our thoughts.

Cognitive traps are a consequence of these observations. If P(S) = 0 or 1 for some sentence S, then P'(S) continues to be 0 or 1 from then on, by the rules of probability. If P(S) is close to one, it will not pay to investigate it, and contrary evidence will tend to be re-interpreted or overlooked. (S can refer to attitudes, values, or ideologies, as well as laws.) When combined with our general tendency to replace doubt by certainty, this can lead to a dogmatic position where S is accepted without possibility of challenge. This is fine if S is true, but what if S is false?

Back to the Greatest Happiness Principle:

There is now sufficient background to state my main thesis. The GHP is objective, in the sense that all are governed by it and strive to attain it, but each maximizes expected value within the world as he sees it (his personal probabilities, as outlined above).

Certain preliminary objections may be deflected.

The problem of **evildoers**: How can Hitler, Stalin, or Osama bin Laden be thought of as utilitarians? Well, given their respective versions of racial ideology, Marxism-Leninism, and radical Islamic theology, they were acting correctly. (In all three cases, developments appear to have falsified large portions of these outlooks.)

The problem of **selfishness**: The GHP gives equal weight to all pleasures and pains, no matter to whom, when, or where they are felt. So why do people seem concerned only with their own (or immediate family's) welfare? This is a more difficult question. Let me reformulate it. We may distinguish circles of increasing "social distance": self, family, friends, neighbors, members of one's social class, ethnic group, national, fellow humans, ..., with apparently decreasing concern as one moves outward. Now, something like this "localism" is inevitable. First, over a wide range, we are best positioned to provide self-services: Everyone must brush his own teeth. Second, the information needed to make good decisions is more richly available locally, and the causal structure is better known. Third, in response to these facts, rational utilitarians will devote less effort to investigating and acting the further the distances involved. Finally, a certain level of "egoistic" action tends to promote general welfare more than some apparently "altruistic" acts. (See the discussion of **competition** below.)

(Incidentally, the discussion of social distance carries over to **time**: an apparent decreasing concern with increasing distance in time—to our own future selves, future generations, ...)

The problem of **dissension**: If we are all basically governed by the GHP, why aren't we all conscious utilitarians? Here are some points relevant to this difficult question. A thing which abstractly seems good will not appear until the time is ripe. Intelligence is evidently advantageous for survival, yet it took billions of years for glimmerings of it to appear. We may assume it was not "cost-effective" before then. Probability is a concept of basic importance, yet people are not very good at using it, it appeared late in history (17th century), and the "probabilistic revolution" is still far from being completed. (Even the utilitarian thinkers Mill and Sidgwick made little use of it, though Bentham and Jevons did.) Happiness as a concept is in even worse shape. So, the time seems not quite ripe yet for general acceptance of the GHP. Note in this connection the theories in which the historical process itself is unbeknownst to the participants: Smith's invisible hand, Marx's vision of the capitalist process, Hegel's march of the world spirit, theological eschatologies, etc.)

How does this all work out in the social world? A very big question, to which I can here give a barely an outline under six headings (in alphabetical order): **capital theory, competition, externalities and institutions, justice and distribution, natural selection,** and **organizations**. The task, which is far from completed, is to put all these insights together into a unified theory.

Capital theory should be (but isn't yet) the central core of economic theory. (We also hear of social, cultural, and political capital, indicating that other social

sciences are picking up on these ideas.) Happiness extends over time, so the GHP is a problem in dynamic optimization. The mathematics of such problems imposes a system of "shadow" prices and values which guide action, and which mimic the structure of prices under free markets. In particular, the value of a resource reflects the value of the future stream of services emanating from it, and in part the value of the inputs going to produce it. The capital concept extends to people themselves, and households may be thought of as little factories producing human capital.

Consider virtue ethics. It stresses the acquisition of good habits of thought and action—the virtues—through practice. This is exactly the capital concept applied to people, provided the virtues are properly defined. (Aristotle, the main source of virtue ethics, omitted benevolence as a virtue.) Thus, virtue ethics is not a rival but actually a corollary to the GHP. Speaking of rival ethical theories, they should be thought of as providing partial insights which have to be properly interpreted and bounded, with the GHP providing more complete insight. This attitude follows from our all being unconscious utilitarians.

Competition. We may distinguish competition among ideas from competition for resources. To take the former, suppose S_1, S_2, ... is a propositional range, so that just one of these competing hypotheses is true, the others false. (They are competing for our credibility, since a total probability of one has to be allocated among them.) Now suppose S_1 is true, unbeknownst to us. Then, by the rules of probabilistic inference, $P(S_1)$ will **tend** to rise, since evidence favoring S_1 is more likely to appear. (All this can be expressed precisely within probability theory itself.) Thus, truth will out, at least in tendency. Much of our evidence is based on the testimony of others, so one must allow for biases, propaganda, etc., but the principle remains. The scientific method, insofar as it exists, is a set of procedures for reducing bias; the rules of evidence in law have a similar function.

One of the functions of the exchange of ideas—"free speech"—is to overcome cognitive traps, to bring forcefully to one's attention possibilities that have been dismissed, or not thought of at all. This applies to competition among ideologies and religions, as well as to science. (One does not usually think of competition as an essential ingredient of science, but it is. The mutual criticism of competing schools of thought keeps all participants on their toes.)

Competition for resources—in the form of wealth, power, prestige, and reputation—threatens to turn into the mutual destruction or the suppression of one side by the other. It works best when controlled: e.g., the market system under the law of property and contract, the political system under checks-and-balances and party competition.

Externalities and institutions. An externality arises by definition between agents A and B when the actions open to A impinge on the interests of B. The question is whether B can easily communicate with A so as to influence A's choice. Example: B offers a monetary reward to induce A to hand over a good desired by B—ordinary market exchange falls under this rubric. If communication is difficult or impossible, this may lead A to ignore or downplay this effect, which in general reduces overall welfare. (The classic example is smoky chimneys, but there thousands of these effects throughout social life.)

4

These uncompensated externalities tend to give rise to institutions that mitigate them. The market has been mentioned, and underpinning this the invention of money and credit. The entire legal system may be thought of in these terms. The law of contracts helps enforce obligations that arise in time-extended agreements between A and B. The law of property establishes clear rights, which facilitates agreements. The law of torts and the criminal law make A more aware of A's effects on other parties, etc. Other institutions include new customs, public opinion, one's own conscience, and the rise of new ideologies such as environmentalism.

Justice, desert, and distribution. Economic theory generally makes recommendations concerning allocative efficiency and avoids discussion of distribution. This is a mistake: income or wealth distribution is part of the allocative process, since it determines agents' relative share of power in the overall utilization of resources into the future. This makes a difference from the viewpoint of GHP maximization, since agents differ in their foresight, managerial abilities, and public spirit. But what principle should govern distribution? I will make a suggestion that has the following three merits. First, it accords with a widespread (but far from unanimous) sense of justice or desert. Second, there is at least a plausible argument in its favor. And third, it is what an ideal free market system would generate, so that any major departure from it would be very costly. The principle is that each agent should receive what he (and his property) produces.

The plausible argument in favor of this "marginal productivity" principle is that past good performance raises the probability of future good performance, so this principle puts more resources in the hands of people likely to use them well. This is not a knockdown argument, and should compete with others.

Conversely, deviations from marginal productivity provide arguments for redistribution. For example, a fair but unknown fraction of overall income appears due to luck. If this fraction correlates with income level, it provides an argument for progressive taxation and welfare programs.

Finally, a few words about inequality. There are deep reasons why the market process generates considerable inequality, arising first from uncertainty: few innovations succeed, and those that do bring great rewards. Further, people focus their attention on a few "stars" (other people, organizations, products, modes of conduct), amplifying their returns. Finally, there are snowballing (dynamic increasing returns-to-scale) processes which magnify inequalities still further.

The price system, which reflects these inequalities, provides valuable information. For example, the run-up in inequality in recent decades indicates that the returns to education have risen and explains the recent rise in interest in that institution.

Natural selection. Given a range of traits $T_1, T_2, \ldots$ that entities may possess (so that any entity has exactly one of these T's), suppose that, in an environment E, these traits influence the survival of the entities having them. Let $P(E \mid T_n)$ be the relative propensity of entities with trait T_n to survive in environment E. Then the prior population $P(T_n)$ of entities with trait T_n will be transformed into the later "posterior" population $P'(T_n)$, proportional to the product of the prior and propensity. (The alert reader will have noticed the isomorphism between this set-up and the movement from prior to posterior probabilities. Darwin and Wallace discover Bayes' Theorem!)

In the natural world, the T's refer to such things as anatomical features, speed, or intelligence. The concept of population needs some thought: should an elephant and a bacterium each be counted as one? A better index of the overall importance or influence of entities would be biomass.

Does biological evolution have anything to do with the GHP? We postulate that organisms seek pleasure and avoid pain. The resulting behavior patterns may or may not be conducive to survival. Culling by natural selection then establishes a positive correlation between traits that are good in the sense of pleasant and good in the sense of promoting survival. Might it then be inferred that natural selection tends to promote happiness aggregated over time?

Do these ideas carry over to the social world? Consider families with different propensities to save: the "grasshoppers" live day-to-day while the "ants" invest heavily. Other things being equal, the latter will gain in wealth relative to the former. (Wealth appears to be a better index of importance or influence here than population numbers per se; we must of course count in wealth in the form of human capital, which is by far the biggest component of wealth for most of us.) And "propensity to save" is just one of a panoply of trait classes relevant to the accumulation of wealth or power: energy, industriousness, intelligence, etc. Furthermore, these effects occur at all hierarchical levels: individuals, families, communities, nations.

These "social Darwinist" tendencies are important, but it is also important to put them in context. Sociobiology appears to exaggerate greatly the importance of genetics. It may be that "the genes hold culture on a leash" (E. O. Wilson), but the leash is long one. Further, some traits are good in one environment but bad in another—e.g. "treat all strangers as enemies" or "eat up all available food." Finally, the higher levels of organization have autonomous principles of their own, such as the institutions discussed above that internalize externalities. With all these considerations, it is not implausible that natural selection operates in the social world to promote the GHP.

Organizations. "Methodological individualism" is the injunction to explain social phenomena from the behavior of individuals. Is this principle sound? One should distinguish the epistemological from the ontological here. As for the former, there is no question that we understand things best in terms of the actions of individuals like ourselves: witness myths, legends, fables, the Gods, etc. But in terms of being, what we find is a hierarchy of levels of organization connected by part-whole relations, no level being any more real than another. (Compare "there is no society, only a bunch of individuals" to "there is no individual, only a bunch of organs.") So what is the extra that appears at each level? At a minimum, it is the spatial relations among the parts. In Newtonian mechanics, one particle moves in a straight line, two in conic sections, and three or more in paths not expressible in closed form. But in general, the parts get modified internally by their influence on each other, as when people become aware of each other.

Time is another aspect of hierarchical organization, with shorter intervals nested into longer. Relative to a given interval variables may be classified as slow, medium, and fast. A slow variable hardly changes at all over the interval. A fast variable is always fully adapted to the others and may be ignored. The medium variables are where the action is, causally interacting in a dynamic process. For example, inventory levels are slow variables when considering a simple transaction taking minutes, medium

variables in business cycles taking a year or so, and fast variables for development processes taking decades. Longer times go with larger organizations: In general, there is a characteristic time for the parts of an organization to interact with each other. Given a much shorter time, the organization supplies only slow variables, and the part itself is temporarily isolated. Given a much longer time, a super-organization at the next hierarchical level may kick in.

Thus, at the highest level we know of, the cosmological, the time scale would be of the order 10^{10} years, with corresponding spatial extent in light-years. At the human level, there are time scales ranging say from one second to many centuries, with corresponding extents from one's immediate vicinity to the entire earth. At the atomic level, the time scale might be 10^{-15} second, with corresponding minute extent.

I conceive all the processes discussed in this paper to be operating simultaneously at all levels, each with its characteristic time.

Some Speculations Concerning the Extent of Mind in Nature

According to what might called "standard scientific materialism," mind is an unimportant froth appearing accidentally on some minor planet: Like the toot of the whistle on a locomotive, according to T. H. Huxley. I will sketch an alternative view, more in keeping with the tenor of this paper (though not essential to it). I am thinking mainly of the affectual qualities of mind (pleasure and pain), though the other "qualia" such as colors and sounds also raise similar issues. There two lines of evidence for the thesis that mind has a far-reaching extent in nature, perhaps a universal extent.

The first is the obvious one: introspection. We note that we enjoy and suffer, and that these feelings are major motivators of our actions. Everything and everyone else we see from the outside, so to speak. So the one part of the universe that is open to us from the inside is found to be full of affect. Note that this is a genuine observation; it didn't have to be that way—we could have been zombies, robots, or machines, noting passionlessly our lack of passion. But once this observation is made it generalizes irresistibly to our fellow human beings, higher animals, and anything else that resembles us—in fact, we generalize too far (anthropomorphism: one suspects that all sources of "standard scientific materialism"—which itself goes well beyond the evidence—is to counteract anthropomorphism.)

The second line of evidence is from the "variational" structure of physical laws. They have the same form as arises from solving dynamic optimization problems, a realization that took a century from the work of Maupertuis, Euler, Lagrange, Jacobi, and Hamilton. Furthermore, this formulation has survived the relativity and quantum revolutions intact. What is one to make of this? Note again that it didn't have to be this way: the formulation was unknown to Newton and Leibniz, though both were inclined to look for optimality in nature. (There are some technical objections to this line of argument which I won't go into here. Still, I believe it has force.)

But then why is there so little direct evidence? A clue lies in the discussion above under organization. Different organisms live their lives at different rates: consider the hummingbird and the shrew, the sloth, and the oyster. Plants, usually considered to lack consciousness (why?) might also be thought of as having very slow reaction times, like the ents in Lord of the Rings. Now suppose there were mind-like entities but living

at a very different pace from us—either much faster or much slower. Would we not overlook them?

Here is another problem. Suppose some organization has a mind, has affects, like ourselves or more specifically our brains. This organization has parts, and these have parts, etc. Does pleasure or pain emerge only at the top, the way conic sections emerge in the two-body problem? That seems odd. I conjecture that pleasure and pain go all the way down. but each smaller part lives at a faster pace, and cannot be sensed in itself at the higher level.

Some Comments on Religion

Here is a new theory of religion or rather of why religion persists. It is because the viewpoint of "standard scientific materialism" is so incongruous with our intuitions. The world is basically mindless, values are subjective, a reflection of our emotions, which in turn are a reflection of our genetic endowment, the genes themselves being purposeless survival machines (Dawkins). Under these conditions, it is not too surprising that people cling to the most popular alternative, which is traditional theism in some form.

Again, take John Mackie's "queerness" argument against objective values, which is that value judgments are unlike any of the other judgments we make. Since a good half of our thoughts might be evaluations, this argument itself is queer. It's like saying "women are such strange creatures: there are no other human beings like them." The argument of this paper, and in particular the speculative "pan-psychic" vision outlined above, offers a third alternative, in which values arise from the bottom up, so to speak, rather than being prescribed from on high. And it is totally compatible with scientific knowledge.

Design in Nature?

The principle of natural selection tends to blur the distinction between design and non-design, since it tends to produce entities who are smart enough to begin designing for themselves. But is there evidence for design in nature outside the human race?

Two suggestions have been made. One is that the evolutionary process itself shows evidence of intelligent design (Dembski, Behe, Johnson). I don't find these arguments persuasive, and won't discuss them further here. The second is rather more interesting: that the physical parameters that govern our known universe have been selected to promote intelligent life (the cosmological anthropic principle), since the minutest change in them would make it impossible.

I find this plausible (probability 1/10, definitely toward the subjective pole), because I can outline a scenario under which it would happen: (i) Our known universe is an infinitesimal speck of what actually exists, which extends way back in time. (The Big Bang is not the beginning of everything.) (ii) Natural selection occurs at all levels of organization and tends to produce intelligent life. (iii) Intelligent life might be interested in propagating other life the way we plant trees. Thus, these cosmic Johnny Appleseeds set the parameters of our known universe to produce us.

Even if these Artificers exist, it would be a mistake to identify them with any traditional Deities. For at such a scale, they would lead very slow lives (maybe taking a

8

billion years to form a single thought). Thus, any kind of direct communication or prayer would be impossible.