

WHY HEALTH IS NOT SPECIAL: ERRORS IN EVOLVED BIOETHICS INTUITIONS*

BY ROBIN HANSON

I. INTRODUCTION

There is a widespread feeling that health is special; the rules that are usually used in other policy areas are not applied in health policy. Health economists, for example, tend to be reluctant to offer economists' usual prescription of competition and consumer choice, even though they have largely failed to justify this reluctance by showing that health economics involves special features such as public goods, externalities, adverse selection, poor consumer information, or unusually severe consequences.

Similarly, while some philosophers argue for bioethical conclusions based on very general ethical intuitions,¹ many others rely on moral intuitions that are specific to health and medicine to draw conclusions that are meant to apply mainly in health and medicine. For example, many authors appear to start from the strong moral intuition that it typically seems wrong to deny poor people access to health care, and then seek moral principles that can both account for such intuitions and justify the claim that people have some sort of right to health care.²

In metaethics, opinions on moral intuitions range from an extreme intuitionism, which accepts all case-specific moral intuitions at face value as reliable moral guides, to an extreme foundationalism, which rejects such intuitions as evidence regarding correct general moral principles. Between these extremes, opinions vary on how severe the errors in our moral intuitions are. The practice of bioethics seems to favor the extreme intuitionist end of this spectrum, and thus implicitly expects mild errors.³ In contrast, this essay will suggest that common practice in bioethics has seriously underestimated the errors in our moral intuitions.

In this essay, I consider the evolutionary origin of our moral intuitions, but avoid the extreme positions of moral skepticism and "whatever evolved must be good," both of which are commonly associated with evolution-

* For their comments on earlier drafts of this essay, I thank Allen Buchanan, Baruch Brody, Bryan Caplan, Tyler Cowen, Frances Kamm, Loren Lomasky, Arti Rai, and Eliezer Yudkowsky.

¹ See, for example, Allan Gibbard, "Health Care and the Prospective Pareto Principle," *Ethics* 94, no. 2 (1984): 261-82.

² See, for example, Norman Daniels, "Health Care Needs and Distributive Justice," *Philosophy and Public Affairs* 10, no. 2 (1981): 146-79; Allen E. Buchanan, "The Right to a Decent Minimum of Health Care," *Philosophy and Public Affairs* 13, no. 1 (1984): 55-78; and Robert M. Veatch, *A Theory of Medical Ethics* (New York: Basic Books, 1981).

³ See Samuel Gorovitz, "Baiting Bioethics," *Ethics* 96, no. 2 (1986): 356-74.

ary ethics. I instead explore the position that our errors are neither negligible nor overwhelming, and that we can identify signs that suggest that particular moral intuitions may be in error. I will argue that humans evolved particular health behaviors and moral intuitions because of specific contingent features of our ancestors' environments, features that are largely irrelevant today. More importantly, evolution has left us largely unaware of how self-serving and in-group-oriented (i.e., favoring our family, tribe, or ethnicity) the functions performed by those behaviors and intuitions were. In general, we distrust especially contingent and self- or in-group-serving moral intuitions. We should therefore consider our health-care-specific moral intuitions to be less reliable than other sorts of intuitions.

I will begin by reviewing the status of intuition-errors in an ethical reflective equilibrium, and then consider some standard signs of errors in our moral intuitions. I will then summarize a particular theory of the evolution of our moral intuitions, and show how this theory suggests that our health-care intuitions display many standard signs of error. I conclude with a discussion of how we might respond to learning that our health-care intuitions are in substantial error.

II. REFLECTIVE EQUILIBRIUM AND INTUITION-ERRORS

Before discussing signs of errors in our moral intuitions, let us review the role that intuition-errors play in reflective equilibrium. Metaethical theorists are often ambivalent about the reliability of moral intuitions.⁴ Most fields of applied ethics, however, rely heavily on moral intuitions, and attempt to find something like a "reflective equilibrium" between case-specific moral intuitions and the general principles that are supposed to account for them.⁵ A person who has no doubt regarding his or her case-specific moral intuitions,⁶ or regarding apparently self-evident general principles, believes that he has no need to engage in further moral analysis. However, one may instead agree with John Rawls that

[a]n allowance must be made for the likelihood that considered judgments are no doubt subject to certain irregularities and distortions despite the fact they are rendered under favorable circumstances. . . .

⁴ See Michael Smith, ed., *Meta-Ethics* (Aldershot, UK: Dartmouth Publishing Company, 1995); Robert L. Frazier, "Intuitionism in Ethics," in Edward Craig, ed., *The Routledge Encyclopedia of Philosophy* (London: Routledge, 1998), 4:853-56; and Steven D. Hales, "The Problem of Intuition," *American Philosophical Quarterly* 37, no. 2 (2000): 135-47.

⁵ On reflective equilibrium, see Richard B. Brandt, "The Science of Man and Wide Reflective Equilibrium," *Ethics* 100, no. 2 (1990): 259-78; Norman Daniels, "Wide Reflective Equilibrium and Theory Acceptance in Ethics," *Journal of Philosophy* 76, no. 5 (1979): 256-82; and John Rawls, *A Theory of Justice* (Cambridge, MA: Harvard University Press, 1971).

⁶ For an example of a theorist who believes we ought not doubt such intuitions, see Jonathan Dancy, "Ethical Particularism and Morally Relevant Properties," *Mind* 92, no. 368 (1983): 530-47.

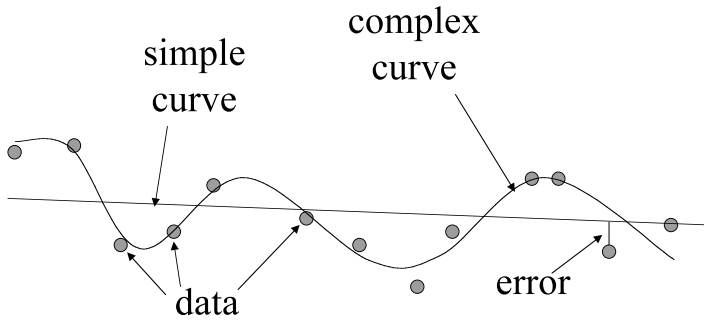


FIGURE 1. Fitting curves to data.

[If so, a person] may well revise his judgments to conform to . . . principles [of justice] . . . especially . . . if he can find an explanation for the deviations which undermine his confidence in his original judgments.⁷

This effort to allow for possible irregularities and distortions in our moral intuitions can be usefully compared to Bayesian curve-fitting.⁸ (There are other statistical approaches to curve-fitting, but these give one results similar to those obtained by the Bayesian method; the differences between these various methods are not important for my purposes here.) As depicted in Figure 1, in curve-fitting one starts with some set of data points, and then looks for a curve that best “fits” that data. That is, one tries to explain the data as being the result of a combination of two things: first, an underlying simple curve, and second, an error process that makes the data deviate from this curve.

In general, there is a trade-off in curve-fitting between simple curves and small errors. One can always fit data with a very simple curve, as is done in Figure 1, if one is willing to tolerate large enough errors. One can also make the errors in one’s curve very small, if one is willing to postulate a complex enough curve. Unless the data is especially cooperative, however, one cannot fit it with a very simple curve and have very small errors.

In a two-dimensional model, such as in Figure 1, a curve might be a line, a parabola, or a more complex polynomial. Alternatively, a curve might be a “spline” that interpolates a smooth curve between “prototypical” points. In models with more dimensions there are many more possibilities.

⁷ Rawls, *A Theory of Justice*, 48.

⁸ On Bayesian curve-fitting, see Andrew Gelman et al., *Bayesian Data Analysis* (London: Chapman and Hall, 1995).

Curve-fitting can involve many possible error processes. In the “normal” error process, for example, all deviations from a curve come independently from the same bell-curve-shaped distribution. In more complex error processes, the errors for different data points can come from distributions with different shapes and sizes, and the errors at nearby points may be correlated with each other.

A “curve-model” is a description of a particular curve, together with a specific error process characterizing deviations from that curve. The mathematics describing such a curve-model say how likely one would be, if the curve-model were the true model of the plotted phenomenon, to see a new data point in any given place. This mathematics also provides a measure of how likely the entire actual set of data points is under the given curve-model—this measure is the curve-model’s “likelihood.” The product of a curve-model’s likelihood and the curve-model’s Bayesian “prior,” which expresses how likely one would have thought the curve-model was before seeing any data, provides an indication of how well the model fits the data. One can then use this product as a relative figure of merit when searching among all possible models for the best-fitting one.

Priors are supposed to embody all the information one has relevant to selecting a curve-model, except for the data at hand. Priors also usually give priority to “simpler” models.⁹ In Figure 1, the strength of the priority given to simpler curves would determine which of the two curves shown would be preferred. Priors also embody expectations about the typical magnitude of errors. The larger one expects errors to be, the more one tends to prefer the simpler of two curves. This is because larger errors will tend to produce larger local fluctuations in data points, and these make it harder to discern local changes in the underlying curve.

In ethical “curve-fitting,” one’s “data” is a set of moral intuitions about what the right actions are in various particular circumstances. Regarding ethical choices made by a group, this data might consist of intuitions from all group members, while for choices made by an individual, the data might be limited to that person’s intuitions. One’s “curves” are sets of ethical “principles,” generally conceived. These can be very general principles, so-called “mid-level” principles, or perhaps the set of ethical choices made in certain prototypical cases (together with the relative salience of considerations used to interpolate between these cases).¹⁰ Together, a set of ethical principles should suggest right actions, and how they vary across some relevant range of circumstances.

To perform ethical curve-fitting, one needs not only data and curves, but error-process models and a prior as well. Models of ethical error processes describe the many ways in which one’s moral intuitions might

⁹ See Richard Swinburne, *Simplicity as Evidence of Truth* (Milwaukee, WI: Marquette University Press, 1997).

¹⁰ For an example of this third approach, see John Arras, “A Case Approach,” in Helga Kuhse and Peter Singer, eds., *A Companion to Bioethics* (Oxford: Blackwell, 1998), 106–14.

deviate from moral truth. And an ethical prior embodies all reasons, except those derived from case-specific moral intuitions, that one has for favoring some ethical principles over others. Putting all this together, we can consider a reflective equilibrium to be the end-result of a curve-fitting-like process. In this process, one searches for a best-fit set of ethical principles, principles one will then treat as one's best estimate of moral truth.

This ethical curve-fitting framework can help us to understand how people with similar case-specific moral intuitions can nonetheless end up with rather different reflective equilibria. For example, those whose ethical priors strongly favor certain moral principles or moral skepticism, and who expect large and highly correlated intuition-errors, will naturally favor those same principles or skepticism in reflective equilibrium. This is because for a given curve-model, a large enough prior can overcome even a very small likelihood. On the other hand, those whose priors do not favor simplicity, and who expect only small or rare errors in intuitions, might naturally come to a reflective equilibrium close to extreme intuitionism. This is because when the priors one assigns to various curve-models do not differ by much, the best models are the ones with the largest likelihoods.

This framework can help us categorize the kinds of arguments one can offer for or against an ethical claim. The most common sorts of these arguments offered are a priori arguments about general principles and arguments pointing out specific moral intuitions that do or do not fit well with certain principles. One can also argue, however, about the degree and types of simplicity we should expect to see in our moral principles, and about the degree and types of errors we should expect to see in our moral intuitions.

This essay focuses on this last, error-based kind of argument. In response to a claim that certain case-specific moral intuitions fit well with certain general principles, a critic can respond that the intuitions in question have characteristics that suggest that they are especially likely to be error-prone, relative to other intuitions one might draw on. This essay will argue that an examination of the evolutionary origins of our common moral intuitions about health care suggests that these intuitions are especially likely to be in error, because they have at least two specific features that are widely considered indicative of errors.

III. CHARACTERISTICS OF INTUITION-ERRORS

What characteristics of moral intuitions indicate that they are especially likely to be error-prone? Consider the opinions of the following six philosophers. Rawls writes:

We can discard those judgments made with hesitation, or in which we have little confidence. Similarly, those given when we are upset

or frightened, or when we stand to gain or lose one way or the other can be left aside. All these judgments are likely to be erroneous or to be influenced by an excessive attention to our own interests.¹¹

Peter Singer writes:

All the particular moral judgments we intuitively make are likely to derive from discarded religious systems, from warped views of sex and bodily functions, or from customs necessary for the survival of the group in social and economic circumstances that now lie in the distant past.¹²

Norman Daniels writes:

[Moral] opinions are often the result of self-interest, self-deception, historical and cultural accident, hidden class bias, and so on.¹³

John Arras writes:

It may well be that some of our most strongly felt convictions, far from being obviously right, are actually the fruit of profoundly unjust social practices and institutions. If we could just step back and gain some critical distance, the injustice might become visible.¹⁴

Jan Crosthwaite writes:

Feminist analysis is concerned less with conscious motivations than with discerning underlying assumptions, and patterns of thinking and practices, of which people may be quite unaware. . . . It also prompts reflection on the possibility of gender biases in the theoretical frameworks of bioethics.¹⁵

Finally, Richard Brandt offers a more detailed discussion:

One important source of ethical errors is doubtless the complexity of ethical problems. . . . Another important cause to which differences of ethical opinion may be ascribed is the degree of maturity of the person judging. . . . A third source of "error" or disagreement [is] . . .

¹¹ Rawls, *A Theory of Justice*, 47.

¹² Peter Singer, "Sidgwick and Reflective Equilibrium," *The Monist* 58, no. 3 (1974): 490-517.

¹³ Daniels, "Wide Reflective Equilibrium and Theory Acceptance in Ethics," 265.

¹⁴ Arras, "A Case Approach," 111.

¹⁵ Jan Crosthwaite, "Gender and Bioethics," in Kuhse and Singer, eds., *A Companion to Bioethics*, 35.

the distortion or warping of the organ of insight by passion, disinclination to make any sacrifice of pleasure or the lower values, and so on. . . . [I]nterest may prevent us from seriously examining a problem although we may talk as if we had done so; or it may make us rest with an incompletely analyzed conception; or it may make us ignore certain factors in the situation of which we are vaguely aware and which we would consider if we were honestly interested in an impartial and objective view.¹⁶

Philosophers are clearly concerned about a great many different types and sources of error. But when a mature individual has a relatively clear and strong moral intuition, two of the most commonly accepted signs that this intuition might nevertheless be in error seem to be that the intuition's origin is excessively historically contingent and that the intuition reflects a hidden bias toward one's self or one's in-group.

First, consider excessive historical contingency of origin. Presumably, every moral intuition has some causal origin, be it cultural, genetic, or otherwise. So the mere fact that an intuition has a causal origin cannot be a reason to suspect that it is in error. The more arbitrary or contingent that causal origin, however, the more we seem to suspect error, all else being equal. A moral intuition that we might expect would evolve in most intelligent social creatures in the universe, perhaps like "Keep your promises," seems more trustworthy than a moral intuition that arises only in very unusual circumstances. Imagine, for example, that you learned that your personal moral intuitions were produced in an experiment that randomly generated hundreds of odd moral intuitions just to see what people would do with them. It seems that learning this fact would reduce your confidence that your moral intuitions are reliable guides to moral truth.

Of course, in a deterministic world there is a sense in which perhaps no event is arbitrary; if you knew enough context, you might know that everything has to be exactly as it is. So to tell which events are more contingent or less so, we must make a principled distinction between local "random" fluctuations and more systematic robust processes. If an event occurs due to robust processes, then that event would have tended to occur even if many of the circumstances surrounding the action had differed. If, in contrast, an event occurs due to random fluctuations, then changes in circumstances could well have caused something else to happen. It is this latter sort of event—and "event," of course, includes occurrences like the development of moral intuitions—that is more contingent.

Second, consider hidden biases toward one's self or one's in-group. If a moral intuition has been selected for by some biological or cultural

¹⁶ Richard B. Brandt, "The Significance of Differences of Ethical Opinion for Ethical Rationalism," *Philosophy and Phenomenological Research* 4, no. 4 (1944): 484–86, 491.

competition, it is probably “self”-interested at some level or another. Presumably, only a rare accident would produce a moral intuition that did not help some relevant unit of selection to reproduce. Thus, the fact that a moral intuition reflects a preference for one’s self or one’s in-group cannot be, by itself, a strong reason to suspect that the intuition is in error.

Self-deception about such a preference, however, does seem to be widely taken as a reason for suspecting error. Consider, for example, the moral intuition that slaveholders should protect slaves from freedom because slaves are incapable of managing such freedom well. Or consider the related intuition that the upper class in a given society should rule due to its superior education and intelligence. Such intuitions are widely suspected of being mere fronts for self- and group-interest, even when they seem to be quite sincerely felt.

These suspicions of self-deception about self-interest receive empirical support from the social sciences. For example, people seem to overestimate their own generosity, even though they can accurately predict the generosity of others.¹⁷ Furthermore, people are more generous toward richer relatives, rather than those who most need help.¹⁸ And when choosing whether to reward those who help them based on effort or on output, people tend to choose in each situation the metric that costs them the least.¹⁹

A moral intuition about the worthiness of some action often comes packaged with a rationale stating that the intuition’s purpose is to benefit some third party. Nevertheless, a careful study of the intuition’s details and origin might well suggest that it functions primarily to benefit the actor or his close associates. Moral intuitions of this sort are commonly considered especially likely to be in error, all else being equal, as a result of some complex process of self-deception. For example, an intuition that one should badger bystanders with lifestyle advice seems suspect if we believe that the real subconscious motivation behind the intuition is a desire to make oneself feel superior or to hear oneself talk.

In Section II, I noted that one can criticize a proposed moral principle by claiming that the case-specific moral intuitions it is based on are especially likely to be in error. In this section, I showed that a quick survey of the opinions of six philosophers suggests that two of the most commonly accepted signs of error in such intuitions are excessive historical

¹⁷ See Nicholas Epley and David Dunning, “Feeling ‘Holier Than Thou’: Are Self-Serving Assessments Produced by Errors in Self- or Social Prediction?” *Journal of Personality and Social Psychology* 79, no. 6 (2000): 861–75.

¹⁸ See Susan M. Essock, Michael T. McGuire, and Barbara Hooper, “Self-Deception in Social-Support Networks,” in Joan S. Lockard and Delroy L. Paulhus, eds., *Self-Deception: An Adaptive Mechanism?* (Englewood Cliffs, NJ: Prentice Hall, 1988), 200–211.

¹⁹ See Norman Frohlich, Joe Oppenheimer, and Anja Kurki, “How Prevalent Is Self-Interest and How Can We Tell?” (working paper, last revised March 2001), available on-line at <http://www.bsos.umd.edu/gvpt/oppenheimer/research/dictprod2.pdf> [cited July 12, 2001].

contingency of origin and an intuition's reflecting a hidden bias in favor of one's self or one's in-group. These commonly expressed standards are certainly open to question, but this essay will not further analyze or examine them. I will instead try to apply these standards to common moral intuitions about health care.

To apply these standards to those aspects of health-care intuitions that we think are determined largely by culture, we would want to examine how contingent the relevant features of a given culture are. We would also need to consider to what extent the details and context of those intuitions suggest that they serve purposes at odds with the rationales commonly given for them in that culture. In this essay, however, I focus on common "evolved" elements of our health-care intuitions. Such elements should be shared by most human cultures, and were inherited from our distant ancestors. Therefore, I want to ask here how contingent the contextual features that determined those evolved moral intuitions were, and to what extent the details and context of those intuitions suggest that they serve purposes at odds with the rationales commonly given for them.

Before we can address these topics, however, we will need to consider in more detail the likely evolutionary origins of our health-care intuitions. And unfortunately, there is so far no widely accepted overall account of the evolutionary origin of human health-care behavior, though there are many partial accounts that agree in many respects. This essay will thus focus on a specific, somewhat comprehensive account that I have elaborated upon in more detail in another paper.²⁰ I will therefore now take a detour to consider some of the empirical data on which this account is based, and then present the account's main theoretical elements. After this, I will return to the topic of errors in our health-care intuitions.

IV. HEALTH-POLICY PUZZLES

Modern-day health policy displays a great variety of phenomena, many of which are puzzling from standard theoretical perspectives. These phenomena are important clues about the origins of our attitudes toward health and health care.

Perhaps the most striking puzzle is that in the aggregate, medicine seems to have almost no marginal effect on health. That is, when we observe external factors inducing people to consume more medicine in some situations than in others, the extra medicine consumed seems, on average, to have no measurable effect on health. Thousands of published clinical trials appear to demonstrate the benefits of specific treatments, at

²⁰ Robin Hanson, "Showing That You Care: The Evolution of Health Altruism" (working paper, last revised November 2000), available on-line at <http://hanson.gmu.edu/showcare.pdf> [cited July 12, 2001].

least in cases where the best possible treatment practice is applied to the patients deemed most likely to benefit. However, regional comparisons, such as across states or across nations, usually reveal no significant health differences that are attributable to variations in aggregate health-care spending.²¹

Even more telling is the RAND health-insurance experiment, an expensive randomized trial of the aggregate benefits of medicine; the study used five thousand people in the United States for three to five years in the 1970s.²² Some of these people were randomly assigned to have basically free medicine, while others were assigned to pay almost full price. Compared with those who had to pay full price, those given free health care obtained about one-third more health care, as measured by dollars of care received. As a result, they had more pairs of eyeglasses, more teeth filled, and more “restricted activity days,” that is, days when they could not do other things they wanted to do because they were sick or dealing with the medical system. But otherwise there was *no* significant difference between the groups’ levels of health, and the statistically insignificant difference that was estimated suggested that free care gave people an average of only about seven weeks of added life. This compares poorly with rough estimates of the average benefits associated with other health-related factors: one gets three years of added life by not smoking, six years by living in rural areas rather than cities, and roughly fifteen years each for exercising more, not being poor, and being a woman.²³ Yet today the United States spends about 14 percent of its gross domestic product (GDP) on medicine. If in practice U.S. citizens basically pay very little for their health care, thus resembling the group in the RAND study that got health care almost for free, the RAND study should be informative about the last one-quarter (that is, $(\frac{1}{3})/(1 + \frac{1}{3})$) of U.S. medical spending, which is about 3.5 percent of GDP. Do people in the United States really spend 3.5 percent of GDP to increase their lifespan by less than 0.2 percent? Spending that money in other ways (for example, spending it to encourage more exercise) would seem to offer much larger benefits.

Another puzzle of contemporary health policy is that it is difficult to find evidence that medicine affects mortality rates over time. It seems that age-specific mortality in developed nations fell at a steady exponential rate for the entire last century, independent of famous medical advances

²¹ See Jonathan Skinner and John E. Wennburg, “How Much Is Enough? Efficiency and Medicare Spending in the Last Six Months of Life,” in David Cutler, ed., *The Changing Hospital Industry: Comparing Not-for-Profit and For-Profit Institutions* (Chicago: University of Chicago Press, 2000), 169–93.

²² This paragraph’s information on the RAND experiment is from Joseph P. Newhouse, *Free for All? Lessons from the RAND Health Insurance Experiment* (Cambridge, MA: Harvard University Press, 1993).

²³ See Paula Lantz et al. “Socioeconomic Factors, Health Behaviors, and Mortality: Results from a Nationally Representative Study of US Adults,” *Journal of the American Medical Association* 279, no. 21 (1998): 1703–46.

or vast changes in health-care spending.²⁴ It is even difficult to find evidence that the introduction of famous medical treatments affected the mortality rates of those who the treatments were designed to help.²⁵

In addition to the aforementioned drop in age-specific mortality rates in the last century, there has been a striking increase in life-spans over the past two centuries. It is not clear, however, what is responsible for this increase. Sanitation improvements appeared late in the process, household variations in water sources and sanitation do not seem to explain household mortality variations, and those who ate nearly as well as we do now a century or two ago died much more often than we do now.²⁶

Also mysterious is the dramatic impact of social status on mortality, present even if one controls for the effects of medicine and health-risking behaviors. One study found that even after controlling for social support, income, education, access to medical care, unhealthy behaviors, and standard demographic characteristics, impoverished areas had mortality rates 50 percent higher than those found in areas that were not impoverished.²⁷

A perhaps related puzzle is the "placebo effect," a phenomenon in which the mere appearance of giving a patient medicine often seems to make him feel better. The entire apparent benefit of common antidepressant medications, for example, may be due to a placebo effect triggered by a patient's ability to discern, because of its larger side effects, that he has been given a "real" drug.²⁸

A final health riddle is the apparent ineffectiveness of preventive measures. Randomized trials have found that special programs to counsel teens regarding smoking have no significant effect on the health of those counseled; programs to counsel adults at high risk for heart attacks or mothers with low-weight babies have similar results.²⁹

²⁴ See Ronald D. Lee and Lawrence R. Carter, "Modeling and Forecasting U.S. Mortality," *Journal of the American Statistical Association* 87, no. 419 (1992): 659-71.

²⁵ See John McKinlay and Sonja McKinlay, "The Questionable Contribution of Medical Measures to the Decline of Mortality in the United States in the Twentieth Century," *Milbank Quarterly* 55, no. 3 (1977): 405-28.

²⁶ See Leng-fei Lee, Mark Rosenzweig, and Mark Pitt, "The Effects of Improved Nutrition, Sanitation, and Water Quality on Child Health in High-Mortality Populations," *Journal of Econometrics* 77, no. 1 (1997): 209-35.

²⁷ See Mary Haan, George A. Kaplan, and Terry Camacho, "Poverty and Health: Prospective Evidence from the Alameda County Study," *American Journal of Epidemiology* 125, no. 6 (1987): 989-98.

²⁸ See Irving Kirsch and Guy Sapirstein, "Listening to Prozac but Hearing Placebo: A Meta-Analysis of Antidepressant Medication," *Prevention and Treatment* 1, no. 2a (1998). *Prevention and Treatment* is an on-line journal available at <http://journals.apa.org/prevention>.

²⁹ See Arthur V. Peterson et al., "Hutchinson Smoking Prevention Project: Long-Term Randomized Trial in School-Based Tobacco Use Prevention—Results on Smoking," *Journal of the National Cancer Institute* 92, no. 24 (2000): 1979-91; Multiple Risk Factor Intervention Trial Research Group, "Mortality after 16 Years for Participants Randomized to the Multiple Risk Factor Intervention Trial," *Circulation* 94, no. 5 (1996): 946-51; and Cecelia M. McCarton et al., "Results at Age 8 Years of Early Intervention for Low-Birth-Weight Premature Infants: The Infant Health and Development Program," *Journal of the American Medical Association* 277, no. 2 (1997): 126-32.

Aside from all these numerous curious aspects of modern health, there are many puzzles associated more directly with health-care spending. For example, richer people within a given nation spend a smaller fraction of their income on health care than poorer people do, but richer nations spend a larger fraction of their income on health care than poorer nations do.³⁰ Thus, how much people spend on health depends not only on their income, but also on the income or health-related spending of those around them. This suggests that some sort of “keeping up with the Joneses” effect is relevant to health-care spending. Another puzzle is that women, those with children, and married people, even though they tend to be healthier than other people are, seem to go to the doctor more often, even after controlling for other relevant characteristics.³¹

Furthermore, note that consumers seem rather uninterested in private information about medical quality. For example, Pennsylvania once published risk-adjusted heart-surgery mortality rates for individual doctors and hospitals. Patients about to undergo heart surgery were surveyed; heart-surgery mortality rates at local hospitals they could choose from varied from 1 percent to 5 percent. Yet only 4 out of 474 patients stated that their choices of doctor or hospital were influenced by knowledge of the free public statistics. In addition, when offered the statistics, only 8 percent said they were willing to pay \$50 to see them.³² In another study, the publication of risk-adjusted information on hospital deaths in the United States had minimal effects on consumer behavior; it was estimated that a hospital’s having twice the risk-adjusted mortality rate of another hospital led to only 0.8 percent fewer patients. A press report of a single suspicious fatality at a hospital, however, resulted in 9 percent fewer patients.³³

Another important oddity in health-care spending is the strong political support for something like national, but not international, health insurance. Germany introduced national health insurance to Europe in 1883, and Japan introduced universal insurance to Asia in 1922; in both cases the national governments were apparently acting in part to gain allegiance from workers unhappy with industrialization. Economists have tried and largely failed to find market failures that might justify such government intervention; furthermore, beliefs about the factual nature of health-care markets do not seem to predict normative positions on na-

³⁰ See Thomas Getzen, “Health Care Is an Individualized Necessity and a National Luxury,” *Journal of Health Economics* 18, no. 2 (2000): 259–70.

³¹ See Newhouse, *Free for All?*

³² The Pennsylvania study is discussed at length in Eric Schneider and Arnold Epstein, “Use of Public Performance Reports: A Survey of Patients Undergoing Cardiac Surgery,” *Journal of the American Medical Association* 279, no. 20 (1998): 1638–42.

³³ See Stephen T. Mennemeyer, Michael A. Morrisey, and Leslie Z. Howard, “Death and Reputation: How Consumers Acted upon HCFA Mortality Information,” *Inquiry* 34, no. 2 (1997): 117–28.

tional health insurance among physicians, economic theorists, or health economists.³⁴

In addition to the aforementioned puzzles posed by contemporary medical trends and by health-care spending, we should note several indications that people view health in a special way. For example, while there are many charities devoted to helping those with health crises, few are devoted to helping people with similarly severe crises such as divorce, falling out of love, unemployment, failing in one's career, losing a friend, and so on. Health's special role is also illustrated by the fact that groups of people given fruit to divide among themselves tend to divide it more equally when told that the fruit is a health aid rather than merely something that tastes good.³⁵

One way current health policy reflects the special role we assign to health is by exhibiting an unusually high level of apparent paternalism—here, an unwillingness to defer to individual judgments regarding trade-offs between health and other considerations. This unwillingness is exemplified by, among other things, the provision of in-kind benefits (e.g., to the poor and elderly); the professional licensing of physicians; regulations covering foods, drugs, and medical devices; and safety rules pertaining to transportation, consumer appliances, and the workplace.

Paternalism seems particularly strong toward low-status individuals. Much more concern is expressed about the risks babies are exposed to during teen pregnancies than is expressed about the risks babies face during the pregnancies of women over the age of forty, even though the latter risks are far more clearly documented.³⁶ Furthermore, liquor stores in poor neighborhoods are the subject of much more concern than liquor stores in rich ones, and the penalties for using the kinds of cocaine favored by the poor are much larger than those for using the kinds favored by the rich.

Related to paternalism is the fact that public health researchers focus solely on bettering health outcomes, and largely disregard other things that people might sacrifice their health to obtain, such as money, fun, better appearance, and so on. Officials often encourage people to eat right, exercise, and get sufficient sleep, but rarely tell them to live a little and take more risks. That health seems to have this sort of priority over other things of value bolsters the idea that we treat health as being, in some way, distinct or special.

In summary, standard theoretical perspectives have trouble explaining many aspects of modern health care. These include numerous curious

³⁴ See Victor R. Fuchs, "Economics, Values, and Health Care Reform," *American Economic Review* 86, no. 1 (1996): 1–24.

³⁵ See Menahem E. Yaari and Maya Bar-Hillel, "On Dividing Justly," *Social Choice and Welfare* 1, no. 1 (1984): 1–24.

³⁶ See Kristin Luker, *Dubious Conceptions: The Controversy over Teen Pregnancy* (Cambridge, MA: Harvard University Press, 1996).

features of modern medicine: the marginal value of medicine with respect to health seems to be extremely low, and we find ourselves unable to adequately explain various phenomena such as the rapid improvement in health over the last century, the health benefits of status, and placebo effects. Further puzzles arise when we look at patterns in health-care spending—for instance, levels of support for national health insurance seem hard to explain as a response to standard market failures. Finally, the standard perspectives cannot account for numerous features of health policy that indicate that we give health a distinct priority over other valued goods. These features include the disproportionate attention we give health crises as opposed to crises that do not directly involve one's health, as well as the fact that modern health care strongly reflects paternalistic attitudes, particularly toward the poor.

V. THE EVOLUTION OF HEALTH-CARE BEHAVIOR

One approach to explaining the puzzling phenomena described in the previous section is to invoke evolutionary psychology. Under this approach, human minds are described as being not generic intelligences whose specific features are determined by cultural environments, but as culturally modifiable bundles of more specific features, features that enabled humans' distant ancestors to survive and reproduce. As a result, instead of seeing modern health phenomena as a response by generic creatures to modern circumstances, we might instead try to see such phenomena as a translation into modern environments of the health-related behaviors and intuitions that we believe would have benefited our distant ancestors.

Explanations of behavior that are based on evolutionary psychology tend to focus on distant explanations, not proximate ones. That is, they focus on the social and personal functions that behaviors once fulfilled, rather than the particular cognitive mechanisms that produce such behaviors. For any given behavioral context, such explanations also tend to focus on normal human behaviors rather than extreme sorts of conduct, and on behaviors common to most human cultures. Once a picture of "regular" human behavior is established, one can turn to specific factors (including cultural and institutional factors) associated with particular groups to explain why their behavior differs from the standard picture.

As with other methods of historical inquiry, many have expressed a concern that evolutionary psychology lends itself to explaining any possible phenomena by appealing to "just so" stories. To retain some discipline when using evolutionary psychology, then, one should try to explain a wide variety of otherwise puzzling phenomena in terms of a small number of assumptions that are plausible a priori. (This claim is basically the core element of standard Bayesian accounts of the relation between theory and evidence.)

The natural place to begin in constructing an evolutionary-psychology explanation of modern health behavior is to examine the health behaviors common among modern primates and human hunter-gatherers. Mammals, such as whales and elephants, are often kind to sick or injured associates; primates are no exception. Primates born without hands or feet have even survived, thanks to food and protection given to them by companions, to raise their own healthy children. Taking advantage of the trend toward aiding those who have been harmed, chimps have been seen faking injuries in order to induce other chimps to avoid hurting them. Primates can also be cruel to sick associates, however. For example, groups of chimps sometimes expel those with contagious diseases, most likely condemning them to death.³⁷

Modern human hunter-gatherers, such as the Ache of Paraguay, get sick or injured about once a month. The typical episode lasts about four days, but sometimes it can go on for months. Sick or injured people are given food by family members and associates, and those who give more tend to get more when they themselves are sick. Giving such aid is treated like sharing meat from a hunt, helping a work-party to build a hut, or joining in an attempt to avenge the killing of a friend. If one fails to do these things, one is seen as being less loyal to one's associates.³⁸

Primates are social animals, and humans are especially social. Therefore, the most important element of our ancestors' environments was probably other ancestors. Complex social groups have intricate, shifting structures of alliances, and such alliances were probably the most important determinants of a person's social status and future success. That is, our ancestors' most important capital assets were their allies, including their children, spouses, and friends. As a result, in good times our ancestors should have invested in finding and maintaining allies, allies they hoped would be there for them in bad times.

How does one find and maintain allies? One does this by giving potential or current allies credible signals that one is both able and loyal, and that one believes them to be able and loyal as well. Most primates do this in part by grooming others and giving them gifts of food. Gifts can signal that one is relatively able: those who are more able are those who can give things away and still survive. This aspect of gift-giving is not limited to primates. In some bird species, some birds literally fight to shove food down the throats of other birds; having one's gift accepted indicates one's higher status.³⁹ And by giving certain associates more gifts than one gives to others, one can signal that one is more loyal to the favored recipients.

³⁷ See Frans de Waal, *Good Natured: The Origins of Right and Wrong in Humans and Other Animals* (Cambridge, MA: Harvard University Press, 1996).

³⁸ This information concerning the Ache is from Michael Gurven et al., "It's a Wonderful Life: Signaling Generosity among the Ache of Paraguay," *Evolution and Human Behavior* 21, no. 4 (2000): 263–82.

³⁹ See Amotz Zahavi and Avishag Zahavi, *The Handicap Principle: A Missing Piece of Darwin's Puzzle* (Oxford: Oxford University Press, 1997).

Larger social groups have advantages in fending off predators, but also have more difficulty maintaining internal peace. Primates that live in larger social groups tend to have larger forebrains, and also tend to spend a larger fraction of their time grooming. Among primates, humans have the largest brains and live in the largest groups. This suggests that human brains may be large in part to support the complex social reasoning required to deal with Machiavellian local politics; language may even have evolved in part to allow people to verbally “groom” several people at a time through conversation.⁴⁰

Our ancestors had many ways to signal loyalty and ability, including grooming, gossiping, sharing food, hosting visitors and visiting others, building homes, throwing feasts (at weddings or funerals, for example), caring for sick or injured allies, helping with wars and revenge killings, and adopting local customs of speech, dress, or music. These were not all equally effective signals, however. Grooming and gossiping involve frequent small costs, and can thus only signal short-term allegiance. For example, if you are planning on betraying someone and groom them once a day, you may well groom them up until the day before you betray them. In contrast, activities involving infrequent large costs can signal long-term allegiance. For example, if you throw an expensive feast once a year to signal your loyalty to others, then just before the next feast you might reasonably ask yourself how likely it is that you will remain allies with your guests for most of the next year.

Among signals of loyalty, the most salient should probably be the infrequent high-cost things done for others when they most need them, especially when times are hard. Thus, caring for the sick or injured and helping with wars and revenge killings should be stronger signals of loyalty than throwing feasts or building homes. Less salient, but still important, would be more frequent low-cost signals like grooming and gossiping.

If people had evolved a very general high-level capacity for inferring and signaling loyalty, then the sorts of actions that signal loyalty might have quickly adapted to modern circumstances. It seems difficult to substitute new signals for the ancient ones, however. For example, when people want to signal their health, wealth, and intelligence to potential mates, they still primarily use the ancient signals of physique, participation in sports, fancy clothes, expensive possessions, and witty conversation. They do not use medical tests, bank statements, educational degrees, or IQ scores, even though it is quite plausible that using these signals would be more reliable and cheaper than using the ancient ones. This suggests that the ways we infer and signal health, wealth, and intelligence are relatively specific and hardwired.

⁴⁰ See Robin Dunbar, *Grooming, Gossip, and the Evolution of Language* (Cambridge, MA: Harvard University Press, 1996).

Analogously, modern humans probably rely heavily on specific ancient unconscious processes when determining who is loyal (and how loyal they are), making choices based on estimates of others' loyalty, and deciding how to signal loyalty to others. People today may consciously know that their world is very different from that of their distant ancestors, but the unconscious processes they use may not reflect this. Our unconscious processes may instead try the best they can to fit the world around us into the categories that were relevant to our distant ancestors.

Perhaps when we see a doctor that someone else paid for caring for us, we feel good toward that benefactor in the same way that our ancestors felt good about people who brought them food when they were sick. That is, when we are sick or injured, if we see ourselves receiving care and infer that certain benefactors helped to induce that care, then ancient parts of our minds may produce in us some standard package of mental reactions. Among other things, these reactions probably make us feel more comfortable with relying on and associating with such benefactors, and probably lead us to conclude that these benefactors are relatively "moral," since helping others seems to have strong associations with morality. Also, when we try to decide whether to be a benefactor for someone else when he or she is sick, similar ancient parts of our minds may generate a standard package of influences over that decision process. These influences probably weigh in favor of our helping the other person, at least to the extent that we want him or her to feel more comfortable with us, and they also probably make us think of ourselves as more moral if we do help.

Of course, according to the theory I outlined above, similar ancient parts of our minds should respond in similar ways when we are concerned with engaging in wars or revenge killings. In our modern world, however, the state has taken over most of the management and funding of these activities (capital punishment being one contemporary equivalent of revenge killings). This leaves health care as one of our few remaining ancient strong signals of long-term loyalty to associates, a signal that has likely retained its importance even as the world has changed. Thus, to the extent that modern health-care behavior reflects ancient instincts, we should expect people to buy health care for their associates just to show that they care, even if on the margin such care does little to improve others' health. In addition, we might expect to see more health care given to those who need more assurance of loyalty, such as women, those with kids, and married people. We should also expect people to be insensitive to private signals about the quality of health care. As I noted in the previous section, we do in fact see all these things.⁴¹

Consider, as an analogy to the provision of health care, giving a gift of a box of chocolates. It is nice if the chocolates are tasty, but your main goal

⁴¹ Of course, we may also expect similar if perhaps weaker sorts of effects to emerge with respect to other once-infrequent signals, such as throwing feasts or building houses. Such possibilities are beyond the scope of this essay, however.

in giving them to another is to show that you care. Therefore, you do not choose the size of the box based on your estimate of how hungry the recipient is; instead, the idea is to distinguish yourself from someone who does not care as much as you do. Things you know about the quality of the chocolates that you do not expect the recipient to know are irrelevant with respect to getting credit for generosity. Similarly, things that the recipient knows about the chocolates' quality that he or she does not expect you to know will not affect his or her estimate of your generosity. Only common opinions about quality matter.

Notice that it is not the case that you do not honestly care about the person to whom you give chocolates or health care. You do care, but you may care even more that they believe that you care. You care about this because almost nobody cares equally about everyone else. People care more about their allies than about others; a promiscuous concern for everyone equally is not evolutionarily stable.

It turns out that this selective concern can directly produce paternalistic attitudes regarding health. Here is how this works: You care about your allies, but you care about them conditionally on their remaining your allies. They care about themselves unconditionally, however. If you are uncertain about whether your allies will remain your allies, you want them to make choices as if they were sure to remain allies with you. Furthermore, if they would be more likely to remain allies with you if they had numerous other allies, you would want your allies to act as if they were confident they would have many allies. And if we think of having high status as being equivalent to having many good allies, then you want your allies to act as if they are confident of being of high status. Those allies, however, will make their choices with reference to the actual status they estimate that they have, not the status that you want them to have.

What choices make more sense for someone with many allies than for someone with few allies? For most primates, being of high status tends to protect one from crisis events that discourage investments in health. This is because mammals have a common "stress response" that suddenly heightens awareness and turns off the body's systems of growth, digestion, and immunity. This response can help a mammal escape from a predator, though at the expense of the mammal's long-term health. Social primates also invoke the stress response when their social status is low or threatened, since having a low social status is typically correlated with suffering crisis events such as beatings or worse. High-status primates, in contrast, invoke the stress response less, and therefore invest more resources in improving their health.⁴² Hence, if you want your allies to act as if they are of high status, you want them to invest a lot in improving

⁴² The information in the preceding sentences is from Robert M. Sapolsky, "Endocrinology Alfresco: Psychoendocrine Studies of Wild Baboons," *Recent Progress in Hormone Research* 48 (1993): 437-68.

their health. In fact, you want them to invest more in their health than they would choose to do themselves, given their own best estimates of their individual future social statuses.

In like fashion, your allies should care more about your health than they do about your happiness. This makes sense of the sorts of paternalism I discussed in the previous section. Furthermore, the status/stress relationship seems to be one of the main culprits behind another puzzle from Section IV—specifically, the puzzle of why the rich live longer than the poor.

The status/stress relationship also gives people another reason to signal loyalty: it calms their allies down. Caring for allies can reassure those allies that they are not going to be abandoned. In response, those cared for will generally have lower stress levels and will avoid trying desperate measures. Today, people visit friends in the hospital for these sorts of reasons; this suggests that similar reasons probably motivated our ancestors to care for each other.

Modern nations and corporations seem to have found many ways to tap into ancient loyalty signals so as to induce people to think of those organizations as their “tribes.” Wars, for example, have been very important in creating national loyalty. A similar functional rationale probably underlies the oddly strong support for national health insurance noted in the previous section: programs of national health insurance and employer-based health insurance likely function in part to strengthen citizen and employee loyalty. While many factors no doubt contributed to the development of national health insurance in Germany and Japan, its early introduction in those nations does seem to have been in part consciously designed to promote national loyalty.⁴³

My discussion so far has been evocative, but not especially precise. In addition, I have explained why the evolutionary-psychology theory I have presented explains some of the puzzles I mentioned in Section IV, but there are numerous other puzzles from that section that I have yet to address. Let me therefore now try to more precisely relate my theoretical assumptions to the phenomena those assumptions might explain. (I have described formal mathematical models detailing these relationships in other work.⁴⁴) This presentation should make it clearer whether the evolutionary-psychology theory I have presented meets the evidential standard I proposed earlier in this section—that is, whether it explains many otherwise puzzling phenomena with a few assumptions that are plausible a priori.

My first theoretical assumptions are that our ancestors were directly and honestly concerned about their allies, that being of higher status was essentially equivalent to having more and better allies, and that people

⁴³ See Charles Phelps, *Health Economics*, 2d ed. (New York: Addison-Wesley, 1997).

⁴⁴ See Hanson, “Showing That You Care.”

had private information about who would remain allies with whom. These assumptions imply that our ancestors were altruistic toward their allies, both for real and for show, and that our ancestors could have quite rationally neglected information about the quality of health care. The assumptions further suggest that health care had a low marginal health value, especially for those of high status (such as, perhaps, the elderly), and that more care was given to those who needed more assurance of loyalty, such as women, those with children, and married people.

I also assume that among our ancestors, allies protected one another from crisis events that discouraged investments in health. Together with the first set of assumptions, this assumption helps explain certain items from Section IV that I discussed earlier in this section—that there is a strong correlation between one’s health and one’s status, and that our ancestors cared more about their allies’ health than they did about their allies’ happiness. This new assumption also helps explain other phenomena noted in Section IV—the existence of the placebo effect, for example, is predicted because the appearance of costly care should induce calm in anxious sick people. The new assumption also helps explain the fact that paternalism seems especially strong toward those low in status, and sheds light on the common perception that high-status people should not get much more care than average-status people do.

Another assumption I make is that for our ancestors, social allies were more fundamentally useful in bad times, when they might prevent one’s starvation, than they were in good times. This encouraged people to spend more to preserve allies during good times, and hence made it the case that activities devoted to helping one find and maintain allies became societal “luxury goods”—that is, goods on which societies spend a higher proportion of their income as their income rises. That this trait emerged amidst our distant ancestors implies that health care and “leisure” socializing should be societal luxury goods among contemporary humans generally, even if within societies the rich spend a smaller percentage of their money on health care and leisure than the poor do. This may also explain why in the worst of times, such as famines, we often see examples of extreme cruelty to associates, such as young people laughing while stealing food from the mouths of their elders.⁴⁵ During such times, the value gained from previously created allies is very high, but the value gained from new investments to obtain future allies is relatively low.

In summary, one can use evolutionary psychology to explain a wide range of otherwise puzzling health behaviors in terms of a small number of assumptions that are plausible a priori. As I will show in the next section, this fact has important implications for the reliability of our health-related moral intuitions.

⁴⁵ This sort of activity is described in de Waal, *Good Natured*.

VI. WHICH EVOLVED HEALTH-CARE INTUITIONS ARE IN ERROR?

The topics touched on by the theory I presented above include charitable giving, national health insurance, health regulation, and health-care behavior among friends and family. These topics are widely thought to have an important moral component. If we provisionally accept the above explanation of common health-care behaviors, what does this tell us about related moral intuitions?

If moral intuitions substantially influence real behavior, and if the theory outlined above truly does explain many health-related behaviors, then that theory must explain a lot about the origins of those moral intuitions. The same event cannot have two independent determinate causes. It thus seems reasonable to suppose that many of the common elements in human moral intuitions about health-care behavior arose to legitimize and police ancient patterns of such behavior. That is, perhaps our ancestors were induced to signal loyalty by caring for sick and injured allies in the way the theory predicts in part because not doing so tended to seem immoral. (This tendency is consistent, of course, with there being many other influences producing local, cultural, and other variations in moral assessments of actions.) Certainly, many of our most important moral intuitions concerning health care are about who is obligated to give what sorts of help to which sorts of sick associates.

Many scholars have considered the ethical implications of the evolutionary origins of our moral intuitions.⁴⁶ The practice of evolutionary ethics, however, has tended to lead scholars toward one of two extreme positions: moral skepticism or forms of ethical naturalism that conclude that “whatever evolved is good.” That is, some have concluded that the evolutionary origins of our moral intuitions make them entirely untrustworthy as indications of moral truth,⁴⁷ while others have been at least accused (sometimes incorrectly, in my view) of presuming that all evolved moral intuitions are correct, at least for the species in which they evolved.⁴⁸ In this essay, I take the intermediate position that moral intuitions are error-prone indications of moral truth, and that we have ways of characterizing which intuitions are more likely to be in error. Specifically, we can examine our evolved health-care intuitions with respect to the two common indicators of intuition-error that I discussed in Section III: excessive historical contingency of origin and hidden bias toward one’s self or one’s in-group.

⁴⁶ See the various essays in Matthew H. Nitecki and Doris V. Nitecki, eds., *Evolutionary Ethics* (Albany: State University of New York Press, 1993); and in Paul Thompson, ed., *Issues in Evolutionary Ethics* (Albany: State University of New York Press, 1995).

⁴⁷ See, for example, Richard Joyce, “Darwinian Ethics and Error,” *Biology and Philosophy* 15, no. 5 (2000): 713–32.

⁴⁸ William F. Harms has been accused of holding this view; see William F. Harms, “Adaptation and Moral Realism,” *Biology and Philosophy* 15, no. 5 (2000): 699–712.

So how historically contingent are the origins of our health-care intuitions? The many cultural variations on such intuitions likely depend on cultural events of varying degrees of contingency. But what about the common, evolved element of such intuitions?

The need to signal loyalty to allies is likely a robust feature of species that acquired high intelligence through expanding sociability and that managed that sociability in part through morality. However, such signaling would probably not have focused as heavily on health care were it not for the happenstance that in most ancient environments health care was an infrequent large expense. The fact that it can be otherwise is demonstrated by the modern world, where health insurance has converted health-care costs into the frequent smaller expense of a health-insurance premium. (Also, if wars and law enforcement were not now financed and managed on such large scales, families and employers could engage in those activities to signal loyalty, and health care probably would not be relied on so strongly for that purpose.)

Does the common, evolved component of our health-care intuitions contain the second indicator of intuition-error, hidden bias toward one's self or one's in-group? As it was described, the above model of evolved health-care behavior allowed, but did not require, people to be fully aware of the nature and function of their behaviors. If we accept this model, however, then we must conclude that in practice people seem to be quite misinformed. People seem to think, for example, that they buy health care for family members and other associates in order to improve the health of those individuals. People also seem to believe that they are very concerned about the quality of this health care. In fact, as was discussed earlier, such care has little effect on health and people do very little to really evaluate health-care quality. When confronted with these facts, it has been my experience that even people who accept them continue on as before, making various weak excuses to explain their behavior.

People seem to be generally aware that they have apparently paternalistic attitudes concerning the health-risking behaviors of others, especially the poor. They most commonly explain these attitudes as being due to those other people being misinformed about risks. This induces endless attempts to "educate" people about those risks. Yet, as was discussed in Section IV, these attempts generally fail. This failure is blamed on people's irrationality, which is then cited as a reason for product bans and limits on behavior.⁴⁹ The theory described above, however, suggests that paternalism is instead due to the fact that lower-class people tend to have fewer allies, which suggests that any one low-class ally of yours is less likely to remain your ally.

⁴⁹ I discuss this at length in Robin Hanson, "Four Puzzles in Information and Politics: Product Bans, Informed Voters, Social Insurance, and Persistent Disagreement" (Ph.D. diss., California Institute of Technology, 1997). See also Robin Hanson, "Warning Labels as Cheap Talk: Why Regulators Ban Drugs," forthcoming (2002) in *Journal of Public Economics*.

Another example of people being misinformed about the functions of their behaviors is the fact that people seem to think that they now spend more of their income on medicine due to improved medical technologies, even though the theory described above says people spend more because they are now richer and because health care is a societal luxury good. Furthermore, many people seem to think that they want to nationalize health care due to some imagined market failure, but the theory we have explored suggests that the important factors are that voters want to show their national loyalty and that national leaders want to induce more national solidarity.

Why did humans evolve to be so ignorant about their motivations? As biologist Richard Alexander says, "It is difficult to understand why individual control over the emotions and conscious understanding of motivations are so obviously incomplete or imperfect, or even deceptive, to the actor him/herself."⁵⁰ To explain this, Alexander has attempted to sketch a theory based on the idea that people who are ignorant of their motivations are more trustworthy. He suggests that when we lack knowledge of the deeper motivations behind our behavior, "[w]e demonstrate that we are prepared to be beneficent without thinking about it, without a cold calculation of costs and benefits."⁵¹

Alexander's ideas are interesting, but for our purposes, it is enough to notice that one's health-care intuitions contain hidden biases toward one's self and one's in-group. According to the theory explored in Section V, we may think that we buy health care to directly benefit others, but a large part of the real reason for our actions involves persuading others that we are able and loyal. While we believe that our apparent paternalism is a response to the ignorance of those we are supposedly helping, it seems to actually be the direct result of an ancient fear that such people will not remain in our group of allies. Also, as I noted above, while many believe that they want national health insurance in order to deal with some failure in the health-care market, the actual function of such insurance appears to be to promote national solidarity.

All of these cases of hidden group bias should lead us to seriously question the reliability of our moral intuitions regarding health care—at least, they should if we follow most philosophers in treating such bias as a strong sign of errors in moral intuitions. If we consider excessive historical contingency of origin to be another sign of error, then we should also be concerned by the arbitrary factors that led health care to become a primary signal of loyalty today.

Two further reasons for suspicion about our health-care intuitions should be mentioned: the standard inefficiencies of both signaling and loyalty. In

⁵⁰ Richard Alexander, "Biological Considerations in the Analysis of Morality," in Nitecki and Nitecki, eds., *Evolutionary Ethics*, 193.

⁵¹ *Ibid.*

the standard economic analysis of signaling, “good” individuals take some costly action to distinguish themselves in the eyes of observers from “bad” individuals. The signaling action needs to be cheaper somehow for good individuals, and its cost must be high enough to deter bad individuals from trying to imitate good ones. In evaluating whether signaling helps or hurts the agents involved, the signaling action’s cost must typically be weighed against any applicable “sorting gain”—that is, social benefits that may come from observers being able to distinguish between good and bad individuals. If the action’s cost is greater than the sorting gain, there is a net signaling loss.⁵²

Let us consider one example of these concepts. In the standard market-failure story of adverse selection in health insurance, people who know that they have a low risk of illness buy less insurance in order to convince insurers that they (the buyers) in fact have that low risk. High-risk people do not attempt this, because they would suffer too much from being underinsured. There is a social loss associated with low-risk people being underinsured, a loss that is not counterbalanced by a sorting gain. After all, every reduction in insurance premiums for low-risk individuals leads to the raising of premiums for high-risk individuals. There are thus only signaling losses, and so in many cases all parties can be made better off by requiring everyone to have the same amount of health insurance. (It should be noted that these problems emerge in a theoretical account of the health-insurance market. In actuality, however, there does not seem to be an adverse-selection problem with respect to individual purchases of insurance; low-risk people actually buy more insurance than high-risk people, not less.⁵³)

For loyalty signaling, the relevant question becomes: Do the sorting gains from loyalty signaling outweigh the signaling losses? It seems unlikely that there are direct sorting gains, because this would require that the gains experienced by those who are found to be more loyal than expected must be larger than the losses experienced by those who are found to be less loyal than expected. In fact, the opposite appears to be the case. However, there might be indirect sorting gains; the threat of being found out as a result of signaling may induce people to be more loyal to their groups.

The intrinsic inefficiency of loyalty, however, makes it unlikely that such indirect sorting gains outweigh the signaling losses of loyalty signaling. Loyalty is intrinsically inefficient because the benefits that groups get from stronger internal loyalty typically come in part at the expense of those outside the group. For example, from inside a family, nepotism may

⁵² This analysis of signaling is derived from A. Michael Spence, *Market Signaling: Informational Transfer in Hiring and Related Screening Processes* (Cambridge, MA: Harvard University Press, 1974).

⁵³ See John Cawley and Tomas Philipson, “An Empirical Examination of Information Barriers to Trade in Insurance,” *American Economic Review* 89, no. 4 (1999): 827–46.

seem like a good thing, and may even seem morally required. Outside that family, however, it can seem like simple corruption. Another example is that while many see the primary advantage of national solidarity as being that it promotes internal peace, such solidarity has in fact been used to help nations in their conflicts with other nations. Let us not forget that Germany and Japan were the first nations in their regions to introduce universal national health insurance in order to promote national solidarity, and that this strong national solidarity is part of what emboldened those nations in the great wars of the twentieth century.

The main problem with this analysis of loyalty signaling is that it applies much more clearly to the happiness of our genes than to the happiness of the people involved. It is possible, and quite probable, that people prefer, and are made happy by, things other than what most helps their genes reproduce. This is because genes' limited ability to express preferences probably led them to make do with some "close enough" approximations. Thus, to do normative analysis in terms of the happiness of people, we need to know more than we seem to about how "genetic happiness" translates into human happiness. For example, on the one hand, evolution might have made people prefer to take specific loyalty-signaling actions. On the other hand, evolution might have simply made people prefer to be thought of as being more loyal than their neighbors. In both cases, evolution has induced people to signal loyalty, but in the former case loyalty signaling would directly satisfy one's preferences, while in the latter case it might not.

Even if a more careful analysis finds positive functional benefits from inducing more loyalty, however, those benefits must be discounted because in the modern world, health insurance means that health care can no longer credibly signal long-term loyalty. Now, it may be that people's evolved habits have fooled them into only paying health-insurance premiums for their allies when they expect to remain allied with them for a long time. But it may also be that people send this signal even when they expect shorter-term alliances, and that it is those who get health care that are fooled into taking that as a signal of long-term loyalty. Or it may be that neither side is fooled.

VII. CONCLUSION

Common moral sensibilities—that is, common reflective equilibria—have changed over the centuries. It seems that long ago, moral sensibilities focused heavily on fealty to one's local lord, group, and spouse, with faithfulness and courage in battle and a thirst for revenge being signs of such loyalty. While there do seem to be strong moral intuitions supporting this sensibility, the modern world is less enthusiastic, suspecting these intuitions of containing substantial errors. In contrast, modern moral sen-

sibilities are much less suspicious of strong common moral intuitions about health care.

This essay has presented a plausible evolutionary-psychology account intended to explain an important component of the origin of those intuitions concerning health care. This account suggests that like those intuitions that reflect older moral sensibilities, common moral intuitions regarding health care also evolved from our ancestors' preferential concern for their allies and from efforts by allies to signal their loyalty. Are our health-care intuitions therefore also in error? Several factors suggest that they are. First, the evolutionary origins of our intuitions were disturbingly contingent. Furthermore, the fact that people are largely unaware of the evolutionary functions of our moral intuitions regarding health care indicates that these intuitions reflect a bias toward one's self and one's in-group, and such bias is widely taken to be a sign of error in moral intuitions. Finally, even if these intuitions had pristine origins and reflected no bias whatsoever, there are reasons to suspect that nevertheless they will provide us with poor moral guidance. As shown in the previous section, economic analysis identifies some ways in which loyalty signaling might fail to make people happy; moreover, in the modern world the provision of health care can no longer credibly signal long-term loyalty.

If we come to distrust most of our health-care-specific moral intuitions, then what do we do then? One response is to adopt full moral skepticism—that is, to mistrust all moral intuitions. This is not, however, the only possible response. If some but not all of our moral intuitions come under a cloud of suspicion, we can simply rely more heavily on our other intuitions. In other words, if moral intuitions taken from contexts outside this cloud of suspicion are presumed to have smaller errors, then we can

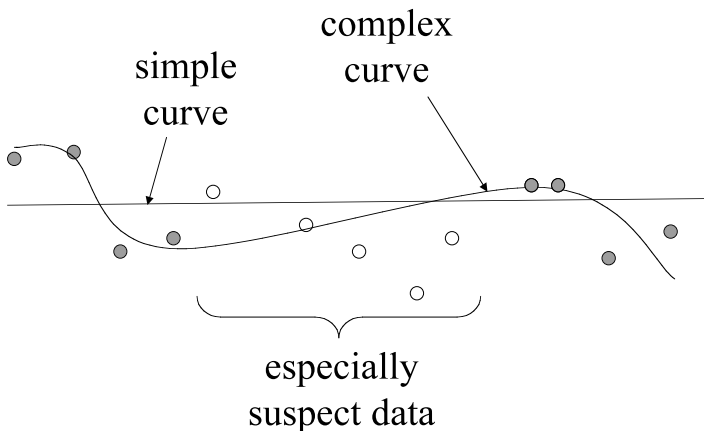


FIGURE 2. Interpolating curves across suspect data.

seek moral principles that primarily fit our data in those contexts and apply those principles to health care.

Figure 2 illustrates this in the curve-fitting context. The curves in Figure 2 reflect a curve-fitting process applied to some, but not all, of the data points that were present in Figure 1. When large regions of one's data are suspect and for that reason given less credence, even complex curves will tend to look simpler as they are interpolated across such suspect regions. In general, the more error one expects in one's intuitions (one's data, in the curve-fitting context), the more one prefers simpler moral principles (one's curves) that are less context-dependent. This might, but need not, tip the balance of reflective equilibrium so much that we adopt very simple and general moral principles, such as utilitarianism. This might not be appealing, but if we really distrust some broad set of our moral intuitions, this may be the best that we can do.

Economics, George Mason University