

John R. Welch

Coping with Ethical Uncertainty

Diametros nr 53, 150-166

2017

Artykuł został opracowany do udostępnienia w internecie przez Muzeum Historii Polski w ramach prac podejmowanych na rzecz zapewnienia otwartego, powszechnego i trwałego dostępu do polskiego dorobku naukowego i kulturalnego. Artykuł jest umieszczony w kolekcji cyfrowej bazhum.muzhp.pl, gromadzącej zawartość polskich czasopism humanistycznych i społecznych.

Tekst jest udostępniony do wykorzystania w ramach dozwolonego użytku.

COPING WITH ETHICAL UNCERTAINTY

– John R. Welch –

Abstract. Most ethical decisions are conditioned by formidable uncertainty. Decision makers may lack reliable information about relevant facts, the consequences of actions, and the reactions of other people. Resources for dealing with uncertainty are available from standard forms of decision theory, but successful application to decisions under risk requires a great deal of quantitative information: point-valued probabilities of states and point-valued utilities of outcomes. When this information is not available, this paper recommends the use of a form of decision theory that operates on a bare minimum of information inputs: comparative plausibilities of states and comparative utilities of outcomes. In addition, it proposes a comparative strategy for dealing with second-order uncertainty. The paper illustrates its proposal with reference to a well-known ethical dilemma: Kant’s life-saving lie.

Keywords: uncertainty, decision theory, probability, probability of frequency, plausibility, plausibilistic expectation, utility, expected utility, Kant.

1. Beyond Negative Capability¹

In an 1817 letter to his brothers, John Keats praised the quality of “Negative Capability.” Possession of this quality, he explained, enables one to be in “uncertainties, Mysteries, doubts, without any irritable reaching after fact & reason.”² Shakespeare possessed the quality “enormously,” in Keats’s judgment; Coleridge, by contrast, did not.

Negative capability may be an admirable quality in a writer, but it is a dubious basis for action. Aristotle liked to say that every action is aimed at a good.³ Our aim at goods we mean to obtain is unlikely to improve without some of Keats’ “irritable reaching after fact and reason.” The irritation is caused by doubt, as Peirce pointed out.⁴ Because doubt—or uncertainty—is so pervasive in human life, we are forced to *reach* after fact and reason almost without pause. This is particu-

¹ The research was supported by a grant no. 2015/17/B/HS1/02279 funded by the National Science Centre, Poland.

² Keats (2009): 60.

³ Aristotle (1984): *Nic. Eth.* 1094a1-3; *Pol.* 1252a1-4.

⁴ Peirce (1986): 247.

larly true of our ethical decision making. How might we cope with ethical uncertainty in a reasonable way? This paper proposes one answer.

I should make clear from the outset that this answer is not meant for all terrains. It is meant to aid ethical decisions that must be made on the sparse terrain of information poverty. We may be painfully uncertain about the appropriateness of our goals, the relevance and accuracy of the available information, the consequences of a contemplated course of action, the possible reactions of other people, or the meaning of vague ethical terms.⁵ The paper's proposal concerns these information-poor situations. When we are better informed, other approaches adapted to richer stores of information are clearly preferable.

The paper aims to present the broad outlines of its proposal without entering into a great deal of detail. Because this proposal leans heavily on the concept of plausibility, Section 2 offers a few preliminary remarks on the concept. Drawing on a specific form of plausibility, Section 3 then sketches a comparative approach to ethical uncertainty, and Section 4 expands this approach to accommodate second-order uncertainty. Several objections are anticipated in Section 5. Finally, a few remarks on information poverty and the versatility of plausibility occupy Section 6.

2. On Plausibility

This paper builds on a specific form of plausibility. Because the concept of plausibility is not nearly as familiar to many readers as the cognate concept of probability, this section of the paper offers a few preliminary remarks on this less customary concept. The remarks are not meant to be comprehensive.

In one form or another, the idea of plausibility has been batted around for a long time. Walton notes that the sophists Corax and Tisias, Plato, Aristotle, Carneades, Locke, and Bentham all managed the concept,⁶ and Rescher provides a substantial bibliography that attests to its antiquity.⁷ Even a cursory review of this literature suggests strongly that there is no single concept of plausibility. Rather, there are a number of plausibility concepts that can be said to bear a certain family resemblance to each other. Although the literature on plausibility continues to grow, a convenient way of organizing it is to distinguish among numeric, non-numeric, and mixed uses of the term. A few examples of each follow.

⁵ Welch (2007).

⁶ Walton (2001): 149-155.

⁷ Rescher (1976): 120-122.

A well-known numeric treatment is Rescher's theory of plausible reasoning.⁸ The theory is meant to reduce cognitive dissonance that may be caused by imperfect or conflicting data. In these non-ideal circumstances, Rescher proposes that the plausibility of a proposition be determined by the reliability of its source. He uses numbers to express the reliability of sources, which are broadly understood to include persons, oral tradition, common knowledge, sense perception, memory, conjecture, inference, and even principles such as simplicity and uniformity. If three sources with reliability indices of .2, .5, and .7 support a given proposition, then the proposition is normally assigned a plausibility of .7: the reliability index of its highest-ranked source. A very different numeric approach is characteristic of Dempster-Shafer theory, which defines plausibility as one of two numbers that together represent uncertainty concerning a proposition. These numbers are supplied by a belief function (Bel) and a plausibility function (Pl) that are duals in the sense that, for a proposition A , $Pl(A) = 1 - Bel(-A)$.⁹ This two-number approach contrasts starkly with uses of a single real number, typically in the interval $[0, 1]$, to represent the plausibility of a proposition. Unlike Rescher's plausibilities, this single number is not determined exclusively by the reliability of sources.¹⁰

A classic non-numeric treatment of plausibility is Pólya's defense of plausible reasoning in—note the irony—mathematics.¹¹ Pólya contrasts demonstrative reasoning, which is safe, uncontroversial, and final, with plausible reasoning, which is hazardous, controversial, and provisional. He treats the operations of generalization, specialization, analogy, and induction as forms of plausible reasoning. A kindred concept of plausibility appears in the law, in the relative plausibility theory of evidence.¹² This theory arose from dissatisfaction with probabilistic models of evidence, whereby the overall probability of one hypothesis—the accused is guilty, for example—on the evidence is compared with the overall probability of another hypothesis—the accused is innocent—on the same evidence. Relative plausibility theorists argue that legal findings are not reached in this way. Instead, they are reached by comparing the relative plausibility of one hypothesis with that of another, where relative plausibility is understood as a form of inference to the best explanation. A more general approach that is not specific to math-

⁸ Rescher (1976).

⁹ Shafer (1987): 62. See also Paris (1994): 38; Klir (2006): 166–167.

¹⁰ Van Horn (2003): 6–7.

¹¹ Pólya (1954).

¹² E.g., Allen (1994); Pardo (2000); Allen, Pardo (2007).

ematics or the law is due to Walton, who argues that abductive reasoning is a third type of inference in addition to deductive and inductive reasoning.¹³ Abductive reasoning is a special kind of plausibilistic reasoning, he claims, and a plausible inference is “one that can be drawn from the given apparent facts in a case suggesting a particular conclusion that seems to be true. Both a proposition and its negation can be plausible, as the ancient legal case of the stronger and the weaker man showed.”¹⁴ Walton’s ancient legal case is discussed by Plato, who attributes it to Tisias (*Phaedrus* 273B–C), and by Aristotle, who credits it to Corax (*Rhetoric* 1402a18–25).

Other treatments of plausibility are mixed in the sense that they admit both numeric and non-numeric plausibility values. A number of publications by Friedman, Halpern, and Chu exemplify this approach.¹⁵ Plausibility in this sense is so amply defined that members of any partially ordered set can count as plausibility values. Hence non-numeric values such as ‘low’, ‘medium’, and ‘high’ as well as numeric values such as probabilities can all be treated as plausibilities. One offshoot of this perspective on plausibility is Friedman and Halpern’s axiomatization of qualitative default reasoning.¹⁶ An alternative approach to qualitative default reasoning that relies on a different sense of plausibility has been proposed by Rott.¹⁷ He defines the plausibility of a proposition as the dual of its entrenchment: “That some proposition *A* is *less plausible* than another proposition *B* means that the expectations militating against *A* are *stronger*, or *better entrenched*, than the expectations militating against *B*.”¹⁸ Even though the Friedman-Halpern and Rott senses of plausibility are similar, a Friedman-Halpern theorist and a Rott theorist can disagree about which of two propositions is less plausible without either having made a mistake.¹⁹ The two senses of plausibility are distinct.

The mixed sense of plausibility employed by Friedman, Halpern, and Chu is a point of departure for this paper. Even though a plausibility measure of this sort is “typically taken to be a probability measure,”²⁰ it need not be; and qualita-

¹³ Walton (2004).

¹⁴ *Ibidem*: 35.

¹⁵ E.g., Friedman, Halpern (1995); Halpern (2003); Chu, Halpern (2004, 2008).

¹⁶ Friedman, Halpern (1995): 182–183, (2001).

¹⁷ Rott (2014). As Rott points out in (2003): 259, his sense of entrenchment is not to be confused with Goodman’s familiar sense.

¹⁸ Rott (2014): 1222.

¹⁹ *Ibidem*: 1236.

²⁰ Halpern (2003): 165.

tive utilities such as ‘terrific’ and ‘terrible’ are admissible.²¹ The paper recommends that we draw on these qualitative possibilities in order to cope with ethical uncertainty.

3. A Comparative Approach to Uncertainty

Ethical discourse can be sorted into three interconnected levels.²² The teleological level stresses ultimate ends like the greatest happiness of the greatest number or the Kantian good will. The instrumental level highlights actions that ought to be performed to achieve a moral end; it also concerns instruments like the hedonistic calculus that may help to attain such ends. Finally, the phenomenal level includes fact-like descriptions of concrete moral phenomena—a utility score for a policy, say, or a claim that a certain maxim is universalizable, or a description of an action as honest or cruel.

We are often plagued by uncertainty at one (or more) of these levels. We might be torn by doubt over rival ultimate ends, or conflicting ought statements, or inconsistent moral properties ascribed to an action. One way to manage this uncertainty is to resort to decision theory. A decision-theoretic analysis would treat the choice of one of these options as a mental act. It could analyze this act as a decision under risk, producing a certain outcome when a given state of the world obtains. The analysis would require a probability function μ that assigns probabilities to the states of the world s_1, s_2, \dots, s_n relevant to the acts under consideration. It would also require a utility function v that assigns utilities to the possible outcomes o_1, o_2, \dots, o_n of performing each act under consideration. Then the expected utility EU of each act a could be calculated using Equation (1).

$$EU(a) = \sum_{i=1}^n v(o_i)\mu(s_i). \quad (1)$$

The analysis would conclude by invoking the decision rule to choose the act that maximizes expected utility.

To manage ethical uncertainty with the help of decision theory is a venerable aspiration, dating back at least to Jeffrey’s “Ethics and the Logic of Decision.”²³ But there are obvious limitations. Recall Kant’s life-saving lie, where an agent is beset with doubt about “whether it would be a crime to lie to a murderer who asked us whether a friend of ours whom he is pursuing has taken refuge in our house.”²⁴ The agent’s options include telling the truth and lying. Relevant to these

²¹ Chu, Halpern (2008): 4–5. See also Halpern (2003): 165.

²² Welch (2014): 11–14.

²³ Jeffrey (1965).

²⁴ Kant (1996): Ak 8:425.

options are possible states of the murderer's beliefs: the murderer may believe the agent or not. Say that if the agent tells the truth and the murderer believes her, the friend will be murdered; but if he does not believe her, the friend will escape. On the other hand, if the agent lies and the murderer believes her, the friend will escape; but if he does not believe her, the friend will be murdered. In order to calculate the expected utilities of lying and telling the truth, the agent would need to know the probability that the murderer believes the agent and the utilities of the various outcomes. An agent able to assign precise and reliable probabilities to the possible states of the murderer's beliefs would be rare indeed. But if the agent is operating under the usual human limitations, she would be unable to provide the requisite numbers. Hence the agent could not calculate the expected utilities of lying and telling the truth.

Nevertheless, there is a decision-theoretic approach that has a fighting chance of being applied in conditions of information poverty. The approach is comparative. To apply it, the concept of probability must be generalized as plausibility and that of expected utility as plausibilistic expectation.²⁵

As noted in Section 2, the concept of plausibility has been variously employed. Here, however, we rely primarily on the work of Friedman, Halpern, and Chu.²⁶ In this usage, plausibility values can be attributed to the members of any partially ordered set. Where \top and \perp are nonnumeric limits representing maximum and minimum plausibilities respectively, a plausibility measure π can be defined for propositions q and r as follows:

P1. If q is contradictory, $\pi(q) = \perp$.

P2. If q is tautologous, $\pi(q) = \top$.

P3. If q implies r , $\pi(q) \leq \pi(r)$.²⁷

Plausibility in this sense is the most general of current modes of representing uncertainty. Probability measures, Dempster-Shafer belief functions, possibility measures, and ranking functions are all special kinds of plausibility measures.²⁸

Since expected utility is probabilistic expectation, the generalization of probability as plausibility forces a correlative generalization of expected utility as plausibilistic expectation. Analogous to the expression of expected utility as a summation of products of probability values and utility values, plausibilistic

²⁵ These generalizations are carried out in detail in Welch (2014): ch. 3.

²⁶ Friedman, Halpern (1995); Halpern (2003); Chu, Halpern (2004, 2008).

²⁷ These axioms are propositional variants of the set-theoretic axioms in Chu and Halpern (2004): 209–210.

²⁸ Halpern (2003): ch. 2.

expectation can be expressed as a quasi-summation of Cartesian products of plausibility values and utility values. Although plausibilistic expectation can be calculated by using an equation analogous to Equation (1) above, this paper is meant to be a non-technical exposition of its main ideas, and we need not introduce this equation here.²⁹ However, the next two paragraphs will give a rough account of plausibilistic expectation that will suffice for present purposes.

Suppose we are confronted with a choice between ethical alternatives a_1 and a_2 . Let a_1 , when combined with state s_1 , produce outcome o_1 ; and a_2 , when combined with state s_2 , produce outcome o_2 . Say that the plausibilities of s_1 and s_2 can be compared in terms of less than ($<$), equal to ($=$), and greater than ($>$). The utilities of o_1 and o_2 are comparable in the same terms. Then there are nine possible cases. These cases are summarized in Table 1, where ' $<$ ' in the plausibility column, for example, abbreviates ' $\pi(s_1) < \pi(s_2)$ ', which says that the plausibility of state s_1 , which would produce a relatively favorable outcome for a_1 , is less than that of state s_2 , which would produce a relatively favorable outcome for a_2 . Similarly, ' $<$ ' in the utility column stands for ' $u(o_1) < u(o_2)$ ', which says that the utility of outcome o_1 from choosing a_1 is less than that of outcome o_2 from choosing a_2 .

Case	Plausibility	Utility
1	$<$	$<$
2	$<$	$>$
3	$<$	$=$
4	$>$	$<$
5	$>$	$>$
6	$>$	$=$
7	$=$	$<$
8	$=$	$>$
9	$=$	$=$

Table 1. The basic binary case

²⁹ The equation can be found in Welch (2014): 64.

Although plausibilistic expectations for this table can literally be calculated, it is not necessary to do so here.³⁰ The resolution of six of the nine cases is evident at a glance. In case 1, for example, plausibility considerations favor a_2 and utility considerations favor a_2 ; hence the obvious decision-theoretic advice would be to choose a_2 . Cases 2, 4, and 9 are a bit different, however. In cases 2 and 4, plausibility considerations favor one option and utility considerations favor the other; hence no purely comparative resolution is possible – unless, of course, we are willing to weight plausibility and utility unequally. In addition, case 9 does not provide a unique resolution like case 1; it provides a disjunctive resolution instead. Such a resolution need not paralyze action, however, for it is simply a tie, comparable to a disjunctive resolution of a moral dilemma.³¹ The results for all nine cases are summarized in Table 2.

Case	Plausibility	Utility	Resolution
1	<	<	a_2
2	<	>	no decision
3	<	=	a_2
4	>	<	no decision
5	>	>	a_1
6	>	=	a_1
7	=	<	a_2
8	=	>	a_1
9	=	=	a_1 or a_2

Table 2. The basic binary case with resolutions

The foregoing paragraphs present a bare outline of comparative decision theory. This form of decision theory is a branch of individual decision theory. As such, it simply assumes the beliefs and desires of individual decision makers as inputs. Although these inputs can range from purely personal tastes to carefully consid-

³⁰ Sample calculations can be found in *ibidem*: 68–71.

³¹ Greenspan (1983): 117–118; Gowans (1987): 19; Zimmerman (1996): 209, 220–221.

ered judgments, the decision-theoretic machinery will accept these inputs, regardless of their cognitive credentials, and produce a corresponding output. The outputs are conditionally rational – rational given the inputs – but they are not necessarily fully rational. They are fully rational only if the inputs on which they are based are rational. Possible objections to this approach are addressed in Section 5.

4. Second-order Uncertainty

In 1975 the U. S. Nuclear Regulatory Commission published its *Reactor Safety Study*, which estimated the risk of an early human fatality due to 100 commercial nuclear power plants in the United States to be 2×10^{-10} /year.³² This study, commonly referred to as WASH-1400, has been called “one of the best and most renowned risk assessments ever accomplished.”³³ Yet it was immediately controversial. The so-called Lewis Committee, which was commissioned to review the study’s conclusions, reported “We are unable to determine whether the absolute probabilities of accident sequences in WASH-1400 are high or low, but we believe that the error bounds on those estimates are, in general, greatly understated.”³⁴

The Lewis Committee’s focus on error bounds for the probabilities used in WASH-1400 sets the tone for this section of the paper. The committee repeatedly asked the question ‘How reliable is that probability?’ The indispensability of this question was pointed out by Aven in a different context:

The assigned probabilities are conditional on a specific background knowledge, and they could produce poor predictions. This leads to the conclusion that the main component of risk is uncertainty and not probability... Surprises relative to the assigned probabilities may occur, and by just addressing probabilities such surprises may be overlooked.³⁵

Questions about the reliability of probabilities were treated in a much-cited paper by Kaplan and Garrick, who proposed a two-level procedure that they termed “probability of frequency.”³⁶ The first level is a range of frequencies f_1, f_2, \dots, f_n that is thought to include the actual frequency of some event. The second level is composed of subjective probabilities p_1, p_2, \dots, p_n , each of which expresses a degree of belief that some frequency is the actual frequency of the event under

³² National Research Council (1975): 112.

³³ Shrader-Frechette (1991): 189.

³⁴ Lewis et al. (1978): viii.

³⁵ Aven (2008): 156.

³⁶ Kaplan, Garrick (1981): 18–22.

consideration. Hence ' p_1f_1 ' might represent the subjective probability p_1 that the actual frequency of the event is f_1 . As Kaplan and Garrick point out, this two-level procedure provides "the ability to explicitly include uncertainty in the calculation of risk."³⁷

Whenever Kaplan and Garrick's procedure can be followed, I submit that it should be. Unfortunately, it cannot always be followed. The frequencies at the first level may be unknown or, if the event is a single case, nonexistent. The probabilities at the second level may be unavailable as well. While subjective probabilities can be elicited from an agent's betting preferences using procedures introduced by Ramsey and by Anscombe and Aumann,³⁸ these preferences must be well-defined to begin with, and this condition may not be met.³⁹ As Savage pointed out, "all [elicitation] subjects report, or otherwise reveal, that they do not know their own preferences; they experience wavering and indecision that cannot be identified with mere indifference."⁴⁰

Consequently, we need a fallback procedure for representing second-order uncertainty under conditions of information poverty. My proposal is to appeal to the generalization of probability as plausibility described in Section 3. Since plausibility measures can be defined for members of any partially ordered set, imprecise terms such as 'nearly certain', 'doubtful', and 'highly unlikely' can be treated as plausibilities. Imprecise plausibilities might be pressed into service in three ways, corresponding to different scenarios of information poverty.

The first scenario is realized whenever objective frequencies are available but subjective probabilities of these frequencies are not. In such cases, the available information might nevertheless be rich enough to permit an adaptation of Kaplan and Garrick's procedure. If we could substitute imprecise but informative plausibility values for subjective probabilities, we might represent second-order uncertainty as plausibility of frequency. For example, a frequency of .6 might be assigned a plausibility of medium (M). Then the plausibility of frequency might be expressed as $M.6$.

The second scenario is the converse of the first. Suppose that objective frequencies are not available but that imprecise plausibilities and subjective probabilities of these plausibilities are. Then second-order uncertainty might be expressed as probability of plausibility. If an agent judges the plausibility of a given state of

³⁷ Ibidem: 21.

³⁸ Ramsey (1931); Anscombe, Aumann (1963).

³⁹ Gilboa (2009): 130–132.

⁴⁰ Savage (1971): 795.

the world to be low (L) but can assign this plausibility a subjective probability of $.3$, the agent's probability of plausibility could be expressed as $.3L$.

The third scenario reflects a greater degree of information poverty, one that is all too frequent in ethical decision making. In these situations, neither objective frequencies nor subjective probabilities of these frequencies are at hand. Nevertheless, spotty information might still allow us to work with imprecise plausibilities. Second-order uncertainty could then be expressed as plausibility of plausibility. For example, if an agent assigns the plausibility high (H) to a given state of the world and has medium confidence in this assignment, then the agent's uncertainty could be expressed as MH .

There is no blinking the fact that representations of uncertainty such as $M.6$, $.3L$, and MH are woefully imprecise. But imprecise information can still be precise enough to guide action. Each of these representations can be compared to others. For example, $M.6 > M.3$; $.3L < .5L$; and $MH > LM$. Unhappily, these representations are not always comparable, for $M.6$ cannot be compared to $H.3$; nor can $.3L$ to $.1M$; nor can MH to HM . But where the representations are comparable, they can be plugged into the plausibility column of Table 2 in Section 3. So employed, they can play key roles in comparative decision-theoretic choice.

As a toy example of the aforementioned third scenario, let us return for a moment to Kant's life-saving lie, which was briefly described in Section 3. There we considered the acts of telling the truth (t) and lying (l), where the relevant states of the world were the murderer believing the agent (b) and the murderer not believing the agent ($-b$), and the possible outcomes were the friend being murdered (m) and the friend not being murdered ($-m$). *Pace* Kant, the outcomes were described in consequentialist terms (though alternative outcomes could be described in deontological terms, as noted in the final objection of Section 5). If the agent has a plausibility measure π that underwrites $\pi(b) = L$ and $\pi(-b) = H$, and if the agent is moderately confident of both plausibilities, then plausibility of plausibility values Π for the two states can be expressed as $\Pi(b) = ML$ and $\Pi(-b) = MH$. The agent's decision could then be represented in preliminary fashion by Table 3.

	$\Pi(b) = ML$	$\Pi(-b) = MH$
t	m	$-m$
l	$-m$	m

Table 3. Kant's life-saving lie (preliminary version)

Table 3 is clearly provisional since the inputs in the interior of the table represent outcomes rather than utilities. But we might be tempted to assign utilities to these outcomes by identifying utilities with numbers of deaths: a utility of -1 for murder and a utility of 0 for no murder, for example. This would be plainly unsatisfactory, however, for the outcome of the friend not being murdered should surely be assigned some positive utility instead of 0 ; in addition, there is no warrant for thinking that the utility of the friend being murdered is exactly one unit less than the utility of the friend not being murdered. Still, the scanty available information would permit the agent to subscribe to a utility function v that assigns utilities U and $-U$ such that $U > -U$. Then utilities of the outcomes of Table 3 can be expressed in comparative terms: $v(m) = -U$ and $v(-m) = U$. Now if the agent's plausibility function assigns plausibilities P and p such that $P > p$, then the plausibility of plausibility values of Table 3 can also be restated in comparative terms: $\Pi(b) = ML = p$, and $\Pi(-b) = MH = P$. As a result, the decision on the life-saving lie can be given a final description in Table 4.

	$\Pi(b) = p$	$\Pi(-b) = P$
t	$v(m) = -U$	$v(-m) = U$
l	$v(-m) = U$	$v(m) = -U$

Table 4. Kant's life-saving lie (final version)

As already noted, the plausibilistic expectations of the two acts under consideration can be calculated. But we can see how the calculation must go without actually carrying it out. If the agent tells the truth, the outcome will be either murder with utility $-U$ or no murder with utility U . But if the agent lies, the outcome will be either no murder with utility U or murder with utility $-U$. Since the possible outcomes of the two acts and the respective utilities of these outcomes are completely symmetrical, utility considerations offer an advantage to neither act. But the relevant plausibilities are asymmetrical. Obtaining a favorable result by telling the truth has a higher plausibility than obtaining a favorable result by lying. Hence if telling the truth is act 1 and lying is act 2, the decision is an instance of case 6 from Table 2. The agent should choose act 1: tell the truth.

5. Possible Objections

Sections 3 and 4 have proposed an approach to coping with the uncertainty endemic to most ethical decision making. Typically, this approach produces action-guiding outputs on the basis of a bare minimum of imprecise inputs. Hence it

offers a lot of bang for the buck. We can anticipate several objections, however. I will mention just three.

The first is that the comparative decision theory of Section 3 is limited to binary choice. The objection is well-taken, but the limitation is not as great as it might first appear. Provided the ethical options under consideration are finite, comparative consideration of a list of options can proceed two by two; that is, option 1 confronts option 2; the winner then confronts option 3; the winner of that comparison then confronts option 4; and so on. Granted, this procedure depends on the assumption that decision-theoretic preference is a transitive relation. I will not attempt to defend this assumption here, but transitivity is common to both the Anglo-American and Franco-European schools of decision theory,⁴¹ and it seems to be as widely accepted as any normative principle of rational choice. Even if it should turn out that the transitivity assumption does not hold without exception, it could still be invoked in those situations where it does hold. The issue is treated more fully elsewhere.⁴²

The second objection focuses on the interpretation of the very mixed bag of plausibility values. How are they to be understood? Because probabilities are plausibilities, probabilistic plausibilities can be interpreted in the usual ways: objectively, as relative frequencies, or subjectively, as degrees of belief. Non-probabilistic plausibilities can be interpreted analogously. For example, Section 4 posited the plausibility values *L*, *M*, and *H*. They could be interpreted objectively by correlating them with relative frequencies such as 0 – 30% for *L*; 31 – 69% for *M*; and 70 – 100% for *H*. Alternatively, they could be interpreted subjectively as degrees of belief, manifest in betting behavior. Suppose we are faced with a choice among three gambles for the same valuable prize: one on an event with plausibility *L*, another on an event with plausibility *M*, and a third on an event with plausibility *H*. The gamble on the *H*-event is strictly preferred to that on the *M*-event, and the gamble on the *M*-event is strictly preferred to that on the *L*-event. Imprecise though these plausibilities are, there is nothing mysterious about them. They can be invoked to cope with the most challenging forms of ethical uncertainty.

The final objection is that applying decision theory to ethical decisions biases deliberation in a consequentialist direction. To address this objection adequately in a few lines is probably impossible, but I will venture a demurral nonetheless. The demurral departs from the observation that the question ‘What makes this

⁴¹ Fishburn (1991): 115.

⁴² Welch (2012): 563–565.

action moral?' and the question 'What makes this action rational?' require different answers. For a Kantian, the answer to 'What makes this action moral?' is 'It is motivated by the good will'. But a Kantian's answer to 'What makes this action rational?' should be different, I submit. In the language of this paper, an action is rational if, and only if, it maximizes plausibilistic expectation.

Let us develop this last point a bit. An action, in order to be an action, must be aimed at a good, as we saw Aristotle claim in Section 1. A Kantian's action should be aimed at the good will, at acting from duty. Hence if an agent chooses to act from a sense of duty and this sense of duty coincides with the legislation of reason, the agent has attained her good. But if an agent chooses to act without being motivated by a sense of duty, or if she chooses to act from a sense of duty that does not coincide with the legislation of reason, the agent has failed to attain her good. Thus an action can be morally good according to Kantian criteria, because motivated by the good will, and decision-theoretically rational, because it maximizes the agent's chance of attaining her good. But the Kantian who manages to be decision-theoretically rational in aiming at the good will is not a consequentialist, for her moral good is not desirable consequences.

The foregoing line of argument may be somewhat controversial, but I will close by noting that Charles Larmore makes a related point in discussing expected utility:

In fact, within the moral domain itself it [maximizing expected utility] does not, strictly speaking, privilege "consequentialist" over "deontological" ways of reasoning, despite the common perception of an elective affinity between maximization and consequentialism (which holds that one is to act so as to bring about the most good overall). For the deontologist maximizes too when he conforms as best he can to the moral principle he holds supreme, which is that one is to respect certain rights, whatever the consequences [...].⁴³

Larmore's point can be generalized from expected utility to plausibilistic expectation. Although the deontologist and the consequentialist aim at different goods, they both attempt to attain some good, and they both can be understood as maximizers of plausibilistic expectation.

6. Versatility of Plausibility

This paper has drawn on the qualitative possibilities of plausibility measures as defined by Friedman, Halpern, and Chu. Specifically, it has made two

⁴³ Larmore (2008): 102.

proposals for dealing with uncertainty in ethical decision making. The first proposal (in Section 3) was that a comparative form of decision theory be employed in conditions of information poverty. The second proposal (in Section 4) was that second-order uncertainty in information-poor situations be represented via a two-level approach adapted from Kaplan and Garrick's probability of frequency.

Like other forms of poverty, information poverty comes in degrees. A relatively mild form permits relevant states of the world and outcomes of the actions under consideration to be characterized by sharp probabilities and utilities. A greater degree of poverty imposes less precise descriptions in the form of probability and utility intervals. Still greater poverty provides just enough information to compare imprecise plausibilities and utilities. The paper's proposals concern only this last form of information poverty.

Lack of information is hardly restricted to ethical decision making. When agents faced with other sorts of decisions must choose despite lack of critical information, plausibilistic reasoning can be brought to bear. This can be done in different ways, as Section 2 has shown. But even if we restrict our reliance on plausibility to the sense employed by Friedman, Halpern, and Chu, resources analogous to those deployed here can be used to guide nonethical decisions of various kinds, including cognitive choice,⁴⁴ real-life decisions,⁴⁵ and theory diagnosis.⁴⁶

References

- Allen R.J. (1994), "Factual Ambiguity and a Theory of Evidence," *Northwestern University Law Review* 88 (2): 604–640.
- Allen R.J., Pardo M.S. (2007), "The Problematic Value of Mathematical Models of Evidence," *Journal of Legal Studies* 36 (1): 107–140.
- Anscombe F.J., Aumann R.J. (1963), "A Definition of Subjective Probability," *The Annals of Mathematical Statistics* 34 (1): 199–205.
- Aristotle (1984), *The Complete Works of Aristotle*, J. Barnes (ed.), Princeton University Press, Princeton.
- Aven T. (2008), *Risk Analysis: Assessing Uncertainties beyond Expected Values and Probabilities*, Wiley, Chichester, England and Hoboken (NJ).
- Chu F.C., Halpern J.Y. (2004), "Great Expectations. Part II: Generalized Expected Utility as a Universal Decision Rule," *Artificial Intelligence* 159 (1–2): 207–229.

⁴⁴ Welch (2011).

⁴⁵ Welch (2012).

⁴⁶ Welch (2013).

- Chu F.C., Halpern J.Y. (2008), "Great Expectations. Part I: On the Customizability of Generalized Expected Utility," *Theory and Decision* 64 (1): 1–36.
- Fishburn P.C. (1991), "Non-transitive Preferences in Decision Theory," *Journal of Risk and Uncertainty* 4 (2): 113–134.
- Friedman N., Halpern J.Y. (1995), "Plausibility Measures: A User's Guide," [in:] *Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence (UAI '95)*: 175–184, URL = <https://dslpitt.org/uai/papers/95/p175-friedman.pdf> [Accessed 26.8.2017].
- Friedman N., Halpern J.Y. (2001), "Plausibility Measures and Default Reasoning," *Journal of the ACM* 48 (4): 648–685.
- Gilboa I. (2009), *Theory of Decision under Uncertainty*, Cambridge University Press, Cambridge.
- Gowans C.W. (ed.) (1987), *Moral Dilemmas*, Oxford University Press, New York.
- Greenspan P. (1983), "Moral Dilemmas and Guilt," *Philosophical Studies* 43 (1): 117–125.
- Halpern J.Y. (2003), *Reasoning about Uncertainty*, The MIT Press, Cambridge (MA).
- Jeffrey C.R. (1965), "Ethics and the Logic of Decision," *The Journal of Philosophy* 62 (19): 528–539.
- Kant I. (1996), *On a Supposed Right to Lie from Philanthropy*, trans. M.J. Gregor, [in:] I. Kant, *Practical Philosophy*, M.J. Gregor (ed.), Cambridge University Press, Cambridge: 605–615.
- Kaplan S., Garrick B.J. (1981), "On the Quantitative Definition of Risk," *Risk Analysis* 1 (1): 11–27.
- Keats J. (2009), "To George and Tom Keats, December 21, 27 (?) December 1817," [in:] *Selected Letters of John Keats: Based on the Texts of Hyder Edward Rollins*, rev. ed., G.F. Scott (ed.), Harvard University Press, Cambridge (MA): 59–61.
- Klir G.J. (2006), *Uncertainty and Information: Foundations of Generalized Information Theory*, John Wiley & Sons, Hoboken.
- Larmore C. (2008), *The Autonomy of Morality*, Cambridge University Press, Cambridge.
- Lewis H.W. et al. (1978), *Risk Assessment Review Group Report to the U. S. Nuclear Regulatory Commission*, NUREG/CR-0400, URL = http://www.iaea.org/inis/collection/NCLCollectionStore/_Public/10/452/10452296.pdf [Accessed 26.8.2017].
- National Research Council (1975), *Reactor Safety Study: An Assessment of Accident Risks in U. S. Commercial Nuclear Power Plants*, National Academy Press, Washington (DC).
- Pardo M.S. (2000), "Juridical Proof, Evidence, and Pragmatic Meaning: Toward Evidentiary Holism," *Northwestern University Law Review* 95 (1): 399–442.
- Paris J.B. (1994), *The Uncertain Reasoner's Companion: A Mathematical Perspective*, Cambridge University Press, Cambridge.
- Peirce C.S. (1986), "The Fixation of Belief," [in:] *Writings of Charles S. Peirce: A Chronological Edition*, vol. 3, C.J.W. Kloesel (ed.), Indiana University Press, Bloomington: 242–257.
- Pólya G. (1954), *Mathematics and Plausible Reasoning*, 2 vols., Princeton University Press, Princeton.

- Ramsey F.P. (1931), *Truth and Probability*, [in:] *The Foundations of Mathematics and Other Logical Essays*, R.B. Braithwaite (ed.), Kegan, Paul, Trench, Trubner & Co., London; and Harcourt, Brace and Company, New York: 156–198.
- Rescher N. (1976), *Plausible Reasoning: An Introduction to the Theory and Practice of Plausibilistic Inference*, Van Gorcum, Assen/Amsterdam.
- Rott H. (2003), “Basic Entrenchment,” *Studia Logica* 73 (2): 257–280.
- Rott H. (2014), “Two Concepts of Plausibility in Default Reasoning,” *Erkenntnis* 79 (6): 1219–1252.
- Savage L.J. (1971), “Elicitation of Personal Probabilities and Expectations,” *Journal of the American Statistical Association* 66 (336): 783–801.
- Shafer G. (1987), *Belief Functions and Possibility Measures*, [in:] *Analysis of Fuzzy Information*, vol. I, J.C. Bezdek (ed.), CRC Press, Boca Raton: 51–84.
- Shrader-Frechette K. (1991), *Risk and Rationality: Philosophical Foundations for Populist Reforms*, University of California Press, Berkeley.
- Van Horn K.S. (2003), “Constructing a Logic of Plausible Inference: A Guide to Cox’s Theorem,” *International Journal of Approximate Reasoning* 23 (1): 3–24.
- Walton D. (2001), “Abductive, Presumptive and Plausible Arguments,” *Informal Logic* 21 (2): 141–169.
- Walton D. (2004), *Abductive Reasoning*, University of Alabama Press, Tuscaloosa (AL).
- Welch J.R. (2007), “Vagueness and Inductive Molding,” *Synthese* 154 (1): 147–172.
- Welch J.R. (2011), “Decision Theory and Cognitive Choice,” *European Journal for Philosophy of Science* 1 (2): 147–172.
- Welch J.R. (2012), “Real-Life Decisions and Decision Theory,” [in:] *Handbook of Risk Theory*, S. Roeser, R. Hillerbrand, P. Sandin, M. Peterson (eds), Springer, Dordrecht: 545–573.
- Welch J.R. (2013), “New Tools for Theory Choice and Theory Diagnosis,” *Studies in History and Philosophy of Science Part A* 44 (3): 318–329.
- Welch J.R. (2014), *Moral Strata: Another Approach to Reflective Equilibrium*, Springer, Cham.
- Zimmerman M.J. (1996), *The Concept of Moral Obligation*, Cambridge University Press, Cambridge.