International Journal of Approximate Reasoning, 2009, 50: 1046–1065.

Heterogeneous Uncertainties in Cholesterol Management

Yakov Ben-Haim
Yitzhak Moda'i Chair in Technology and Economics
Faculty of Mechanical Engineering
Technion—Israel Institute of Technology, Haifa, Israel

Clifford C. Dacso
John S. Dunn Sr. Research Chair in General Internal Medicine
The Methodist Hospital Research Institute
Distinguished Research Professor
University of Houston, Houston, Texas

Jonathon Carrasco and Nithin Rajan
Abramson Center for the Future of Health, Houston, Texas

Contents

1 Introduction **3**

2 Clinical Guidelines **3**

3 Cholesterol Management and the LDL Risk Model **4**

4 Info-Gap Analysis with Uncertain Expected Utility **5**
4.1 Uncertain Expected Utility 5
4.2 Robustness and Opportuneness 8

5 Results: 4 Case Studies **11**
5.1 Case 1 11
5.2 Case 2 14
5.3 Case 3 15
5.4 Case 4 16

6 Discussion: Medical Guidelines **17**

7 Discussion: Info-Gap Theory and Other Methods **19**

8 References **23**

A Evaluating the Robustness and Opportuneness Functions **26**

Corresponding author: Yakov Ben-Haim, Tel: +972-4-829-3262, Fax: +972-4-829-5711. yakov@technion.ac.il.

Abstract

Physicians use clinical guidelines to inform judgment about therapy. Clinical guidelines do not address three important uncertainties: (1) uncertain relevance of tested populations to the individual patient, (2) the patient's uncertain preferences among possible outcomes, and (3) uncertain subjective and financial costs of intervention. Unreliable probabilistic information is available for some of these uncertainties; no probabilities are available for others. The uncertainties are in the values of parameters and in the shapes of functions. We explore the usefulness of info-gap decision theory in patient-physician decision making in managing cholesterol level using clinical guidelines. Info-gap models of uncertainty provide versatile tools for quantifying diverse uncertainties. Info-gap theory provides two decision functions for evaluating alternative therapies. The robustness function assesses the confidence—in light of uncertainties—in attaining acceptable outcomes. The opportuneness function assesses the potential for better-than-anticipated outcomes. Both functions assist in forming preferences among alternatives. Hypothetical case studies demonstrate that decisions using the guidelines and based on best estimates of the expected utility are sometimes, but not always, consistent with robustness and opportuneness analyses. The info-gap analysis provides guidance when judgment suggests that a deviation from the guidelines would be productive. Finally, analysis of uncertainty can help resolve ambiguous situations.

Keywords: Clinical guidelines, cholesterol management, judgment under uncertainty, patient satisfaction, info-gap decision theory.

1 Introduction

The emergence of probability theory in the early 17th century brought to the fore the diversity of categories of knowledge. On the one hand the astronomer, for instance, could aspire to certain, almost apodeictic, knowledge of the laws of celestial motion. In contrast, the physician could at best make judgments of plausible truth. The deductive certainty of Aristotelian logic was inadequate for the inferences made in the 'low' sciences of medicine, alchemy etc., and new modes of inference for uncertain induction began to emerge [1].

The twentieth century saw a diversification of models of uncertainty, reflecting a broadening diversity of categories of knowledge which underlie judgment and inference: Linguistic knowledge; observed frequencies; hunches and beliefs; microscopic or fundamental quantum randomness as distinct from the macroscopic randomness of ensembles; randomness which arises from partial knowledge of a deterministic system ("Like the queen of England, determinism reigns but does not govern." [2, p.63]); ignorance one is aware of, and ignorance one is ignorant of.

Great strides have been made in the much needed synthesis of a coherent system for the diverse types of knowledge and their attendant uncertainties [3]. Nonetheless, "As our island of knowledge grows, so does the shore of our ignorance." [4]. In particular, the integration of non-measure-theoretic models of uncertainty, such as info-gap models [5], into a system of measure theoretic formulations, has not yet been accomplished [6, p.37].

Once again, as in the 17th century, we find that medical decisions confront us with heterogeneous uncertainties for which we are incompletely prepared. Clinical trials are rendered to the practitioner as probabilistic functions, expressing outcome-likelihoods under specified conditions. However, as we will explain in section 2, these probability estimates can be subject to uncertainties for which we have no probabilistic information. Furthermore, patients are challenged to quantify their anticipated utility (or dis-utility) of future outcomes of treatments with which they have no experience. The enlightened physician assists the patient to fold these anticipated utilities into the selection of a therapy. However, these anticipations are fraught with severe uncertainty as we will discuss in section 4.1. For these uncertainties we have limited or no probability models. The challenge is to develop decision tools and evaluate the quality of these tools. The diversity of types of information, and of the associated uncertainties, has impeded the development of generic and universal decision methodologies.

This paper illustrates a methodology for quantifying severe non-probabilistic uncertainties, for combining uncertainties of different types, and for then using these results to support a medical decision. We make no claim for the optimality of this methodology, which is based on info-gap decision theory [5]. Indeed it is a formidable task to define and evaluate the optimality of a decision strategy, which is not a goal of this paper. However, the development of concepts and criteria of optimality is supported by the study of diverse methodologies. We contribute to this task by offering a methodology—info-gap theory—which is different from the many existing measure-theoretic techniques.

In section 2 we review the role of clinical guidelines in medical decision-making, and discuss the attendant uncertainties. In section 3 we present a standard probabilistic model for cholesterol risk assessment which underlies medical intervention. In section 4 we formulate the info-gap models of the relevant diverse uncertainties. We then formulate the info-gap robustness and opportuneness functions which underlie the choice of an intervention. In section 5 we apply the info-gap analysis to 4 plausible clinical cases. In section 6 we discuss the implications of our work for medical guidelines. We conclude in section 7 with a methodological comparison of info-gap theory with other methods.

2 Clinical Guidelines

The dissemination of clinical guidelines has significantly impacted medical practice. Inaugurated as algorithms for common problems in clinical medicine, guidelines now encompass all aspects of medical practice. The National Guideline Clearinghouse indexes over 2100 guidelines [7]. Mehta *et al.* [8] conclude: "Improved application of existing therapies, directed by evidence-based guidelines, may offer immediate savings of life and function to patients with cardiovascular disease." Frei *et al.* [9] report significant improvement in clinical outcomes from "guideline-concordant antibiotic therapy" in treating community-acquired pneumonia.

But how well do clinical guidelines deal with the diverse uncertainties facing both patient and clinician? How can patient and physician assess these uncertainties and incorporate this assessment

in their decisions? This paper studies the management of low density lipid (LDL) cholesterol as a framework for illustrating the construction of non-probabilistic info-gap models of uncertainty. We consider diverse types of uncertain information—some probabilistic, some not—which must be combined in assessing and choosing among therapeutic alternatives. We define and construct robustness and opportuneness functions to support these tasks.

A criticism of clinical guidelines is that they constitute "cookbook medicine" and externally applied restrictions on clinical judgment and autonomy [10]. Advocates of guidelines also stress that "it is important to tailor treatment to the needs of each individual patient" [11].

In addition, guidelines do not readily facilitate consideration of three challenging uncertainties which physicians regularly face.

First, guidelines are based on clinical trials with populations which may not reliably reflect the individual patient. Both patient and physician are sometimes quite uncertain about the relevance of the clinical trials to their specific case. Greenfield *et al.* [12] note that randomized control trials, which underlie clinical guidelines, typically enroll patients with less severe disease and exclude older patients, making the resulting guidelines of uncertain applicability to the excluded populations. Feinstein and Horwitz [13] warn against the prevalence of randomized clinical trials in which "the data do not include many types of treatments or patients seen in clinical practice". Morimoto *et al.* [14] note that clinical guidelines, developed in the U.S. for use of aspirin in primary prevention of cardiovascular events, need modification before application in Japan. McLaughlin [15] reports the conclusions of a roundtable discussion of implications of heterogeneity of treatment effects (HTE). He concludes that, due to HTE, and especially in the absence of "sound data", "care has to be individualized, using the clinician's best judgment regarding available treatment options."

The *second* challenge is that guideline recommendations do not account for the individual patient's uncertain preferences among possible outcomes of treatment, especially adverse outcomes unfamiliar to the patient.

The *third* challenge, related to the second, is that guidelines often fail to account for the cost of intervention, either the subjective cost of lifestyle change or the quantitative financial cost.

This paper presents a quantitative decision-theoretic methodology for addressing these uncertainties when using clinical guidelines. Quantitative decision theory underlies computer-based clinical decision-support. Availability of decision-support technology has resulted in substantial increase in physician compliance with clinical guidelines [16].

We apply the methodology to LDL cholesterol reduction. The guideline recommendations are sometimes, but not always, found to be consistent with the info-gap analysis of uncertainty in a series of plausible clinical examples. The examples illustrate how physician judgment to deviate from the guidelines can be supported, or refuted, by the analysis of uncertainty, as well as how patient preferences can be incorporated in the decision process.

This paper employs info-gap decision theory [5], which has been applied in a large array of decision problems under severe uncertainty, including biological conservation [17, 18], resource management [19], ethology [20], statistical testing [21], homeland security [22], engineering design [23, 24, 25, 26, 27] and fault diagnosis [28], project management [29, 30], portfolio investment [31] and conflict resolution [32].

3 Cholesterol Management and the LDL Risk Model

Alteration of cholesterol concentration is a common intervention in primary and secondary prevention of heart disease. Population studies, like the longitudinal Framingham study [33], demonstrate that cholesterol level is a major risk factor for cardiovascular events, and that alterations in cholesterol level can reduce the incidence of stroke and heart attacks.

Debate continues about altering LDL and high density lipid (HDL) cholesterol levels. Brindle *et al.* [34] review the applicability of the Framingham data to diverse populations, concluding that caution is needed in applying the Framingham score to some sub-populations such as lower socio-economic groups. Kostis [35] suggests that the 10-point cardiac risk assessment based on the Framingham study may not deal adequately with some sub-populations such as young women. Pharmacological interventions such as statins are costly and have measurable toxicity, especially for young people facing life-long treatment [35]. Some patients resist life-long intervention if it is only ameliorative. Robson [36] observes that there "is most uncertainty about treating the many people at intermediate risk ... [which] can turn large numbers of people into lifelong patients." Finally, no clinical study is ultimately definitive. For instance, Grundy *et al.* [37] report on clinical

trials which suggest some modifications of the earlier Framingham study.

The patient considering intervention is uncertain how he or she matches the study population. The Framingham study, which is the canonical prototype of all major risk models to date, examined free-living, healthy, middle-aged white suburbanites west of Boston [33], and may not be applicable to, for example, an Hispanic in Chicago or Lima. Although the application may be valid, patients and physicians are legitimately uncertain. The patient may also be uncertain regarding the future utility or disutility of the treatment outcome.

Wilson *et al.* [33] estimate the probability of a cardiac event based on the patient's condition which is characterized by a vector, c, of known, non-negative numbers. c_1 and c_2 equal the patient's age and age-squared. c_{13} and c_{14} each equal 1 if the patient has diabetes or smokes, and equal zero otherwise. Elements 3–7 specify the patient's HDL level, where $c_{2+i} = 1$ for patients in the ith HDL group. Only one of $c_3, \ldots, c_7$ is non-zero. Elements 8–12 specify the patient's blood pressure group, where $c_{7+i} = 1$ for patients in the ith group. Only one of $c_8, \ldots, c_{12}$ is non-zero. Elements 15–19 specify the patient's LDL level, where $c_{14+i} = 1$ for patients in the ith LDL group. Only one of $c_{15}, \ldots, c_{19}$ is non-zero. The non-zero element from among $c_{15}, \ldots, c_{19}$ indicates which therapeutic intervention has been chosen.

The probability of developing coronary heart disease (CHD) over a ten-year period, for members of the ℓth LDL group, is based on the Framingham study [33]:

$$P_\ell(\gamma) = 1 - S^{\mathrm{e}^{L-G}} \tag{1}$$

$S = 0.90017$ for men and $S = 0.9628$ for women, and $G = 3.00069$ for men and $G = 9.914136$ for women. $L = c^T\gamma$. The subscript ℓ indicates that $c_{14+\ell} = 1$ while the other elements of $c_{15}, \ldots, c_{19}$ equal zero. The best estimate of the vector γ is $\widetilde{\gamma}$, shown in table 1. The estimates of S, G and $\widetilde{\gamma}$ were obtained in [33] by a prospective study of 2489 men and 2856 women aged 30 to 74 at baseline with 12 years of follow-up. During these 12 years 383 men and 227 women developed CHD. Standard errors of these estimates are not available in [33]. The Framingham study is based on a middle-aged white American population. The practicing physician may be uncertain to what extent eq.(1) is relevant to a patient who belongs to another sub-population, either ethnically or due to idiosyncratic medical history.

4 Info-Gap Analysis with Uncertain Expected Utility

In section 4.1 we discuss several info-gap models for representing different types of uncertainty: uncertainty in the parameters of a probability function, uncertainty in estimated utility coefficients, and uncertainty in the shape of a utility function. In section 4.2 we discuss the info-gap robustness and opportuneness functions and explain how they are used to choose among the alternative therapies.

4.1 Uncertain Expected Utility

Many utility and quality-of-life functions are used in medical decision making [38]. We will use expected utility, which is a generally accepted, versatile, and powerful method for exploiting the probability model developed in the Framingham study, though one could use other quality-of-life functions such as quality adjusted life years (QALYs).

In an expected utility model one represents the patient's preferences as a linear combination of the utility coefficients of elementary events. In our case, the patient has positive utility u_{h} from health, negative utility (disutility) u_{e} from a cardiac event, and negative utility $u_{\mathrm{i}}(\ell)$ from the cost of intervention to bring the patient's LDL concentration to the ℓth level. While u_{h} must be positive and u_{e} must be negative, they need not have equal magnitude: many people have asymmetric utilities for good and bad outcomes. The elicitation and combination of these utilities can be done in many different ways, (see for instance Keeny and Raiffa [39] and Keeny [40]). Our discussion is independent of how this is done. In any case, u_{h}, u_{e} and u_{i} are utilities which are calibrated subjectively in linguistic terms. The patient's estimates of these utilities are based on introspection, conversation with family, friends and medical professionals, and perhaps formal supervised elicitation. The elicitation process establishes subjective meaning for the patient of different values of utility. That is, positive utility in different numerical ranges is related to various linguistic descriptors such as 'good', 'very good', etc. Likewise, negative utilities in various ranges have descriptors such as 'poor' or 'not too bad'. The formulation and elicitation of these estimates is

Symbol	Variable	Men	Women
$\widetilde{\gamma}_1$	Age, y	0.04808	0.33994
$\widetilde{\gamma}_2$	Age squared, y^2	$0^\ddagger$	−0.0027
	HDL−C, mg/dL		
$\widetilde{\gamma}_3$	< 35	0.48598	0.88121
$\widetilde{\gamma}_4$	35–44	0.21643	0.36312
$\widetilde{\gamma}_5$	45–49	$0^\star$	0.19247
$\widetilde{\gamma}_6$	50–59	−0.0471	$0^\star$
$\widetilde{\gamma}_7$	≥ 60	−0.3419	−0.35404
	Blood Pressure		
$\widetilde{\gamma}_8$	Optimal	−0.02642	−0.51204
$\widetilde{\gamma}_9$	Normal	$0^\star$	$0^\star$
$\widetilde{\gamma}_{10}$	High normal	0.30104	−0.03484
$\widetilde{\gamma}_{11}$	Stage I hypertension	0.55714	0.28533
$\widetilde{\gamma}_{12}$	Stage II–IV hypertension	0.65107	0.50403
$\widetilde{\gamma}_{13}$	Diabetes	0.42146	0.61313
$\widetilde{\gamma}_{14}$	Smoker	0.54377	0.29737
	LDL−C, mg/dL		
$\widetilde{\gamma}_{15}$	< 100	−0.69281	−0.42616
$\widetilde{\gamma}_{16}$	100–129	$0^\star$	$0^\star$
$\widetilde{\gamma}_{17}$	130–159	0.00389	0.01366
$\widetilde{\gamma}_{18}$	160–189	0.26755	0.26948
$\widetilde{\gamma}_{19}$	≥ 190	0.56705	0.33251

Table 1: Estimated γ coefficients, $\widetilde{\gamma}$, for predicting cardiac event using LDL-C categories, based on Wilson *et al.* [33]. The values marked '*' are "referent" values which are zero by definition. The value marked '‡' is zero by definition: the age-squared term is not included in the regression for men.

not easy or trivial, and these estimates are clearly highly uncertain. Our primary concern, however, is not with the process by which these estimates are formed, but in managing the considerable uncertainty which accompanies these estimates.

The expected utility is the average of $u_{\rm h}$ and $u_{\rm e}$, weighted by $P_\ell(\gamma)$, plus the fixed disutility of intervention, $u_{\rm i}(\ell)$. Denote these 3 utilities collectively with the vector u. Expected utility for the ℓth LDL level is:

$$E_\ell(\gamma, u) = P_\ell(\gamma)u_{\rm e} + [1 - P_\ell(\gamma)]u_{\rm h} + u_{\rm i}(\ell) \tag{2}$$

$u_{\rm h}$, $u_{\rm e}$ and $u_{\rm i}(\ell)$ are uncertain: people poorly predict future feelings. $P_\ell(\gamma)$ is uncertain because of uncertain relevance of the population study to the individual. These uncertainties are information gaps between what we *do know,* and *need to know,* in order to dispel reasonable doubt. Info-gap models quantify these uncertainties [5]. Roughly, $u_{\rm h}$, $u_{\rm e}$, $u_{\rm i}(\ell)$ and $P_\ell(\gamma)$ are estimated, but we don't know how wrong those estimates are. An info-gap model of uncertainty quantifies the unbounded range of possibilities which this entails, without using probability distributions or presuming knowledge of a worst case.

We now consider uncertainty quantifications for $P_\ell(\gamma)$, $u_{\rm e}$, $u_{\rm h}$ and $u_{\rm i}(\ell)$.

Uncertain probability. The probability $P_\ell(\gamma)$ is uncertain because the coefficients γ are uncertain. (One could also consider uncertainty in the form of the function, but we will not explore that here.) Our best statistical estimate of the γ-coefficients is $\widetilde{\gamma}$, table 1. However, those coefficients are estimated for a specific sub-population (middle-aged white American suburbanites) and we don't know how different the coefficients would be for some other sub-population from another ethnic group or with specific medical histories not represented by the Framingham study. We are not able to identify a worst case (greatest deviation), nor do we have a probability distribution for the error of $\widetilde{\gamma}$. We will represent this uncertainty with an info-gap model: an unbounded family

of nested sets of possible γ vectors. The nesting parameter of this family of sets, denoted by α, is referred to as the horizon of uncertainty, which is an unspecified non-negative real number. As the name implies, the horizon of uncertainty expresses the level of uncertainty. Nonetheless, it is axiomatically different from a probability, and does not obey the Kolmogorov axioms of a probability measure [41].

We now formulate an info-gap model for uncertainty in $\widetilde{\gamma}$.

In some situations one might know a standard error, σ_i, of the estimate $\widetilde{\gamma}_i$, in which case one can define $\rho_i = \sigma_i/|\widetilde{\gamma}_i|$ for all $\widetilde{\gamma}_i$'s which are included in the regression. Those terms which are zero by definition (marked '$\star$' or '$\ddagger$' in table 1) do not vary. For these terms we arbitrarily define $\rho_i = 1$, which has no impact on the numerical results. In our numerical examples, based on the Framingham study [33], we do not know standard errors. We have no information with which to differentiate between coefficients. Hence we define $\rho_i = 1$, which means that the available range of fractional error is the same for each coefficient of the terms in the linear combination, at any horizon of uncertainty. Note that we do not know the value of the horizon of uncertainty, which may be different for different coefficients.

We now define a fractional-error info-gap model [5] for uncertainty in the γ_i's:

$$\mathcal{U}(\alpha, \widetilde{\gamma}) = \{\gamma : \ |\gamma_i - \widetilde{\gamma}_i| \le \alpha \rho_i |\widetilde{\gamma}_i|, \quad i = 1, \ldots, 19\}, \quad \alpha \ge 0 \tag{3}$$

Like all info-gap models of uncertainty, $\mathcal{U}(\alpha, \widetilde{\gamma})$ is a family of nested sets of possible realizations of the uncertain quantity, γ in this case. In the absence of uncertainty (when $\alpha = 0$) the set $\mathcal{U}(0, \widetilde{\gamma})$ contains only the estimated vector, $\widetilde{\gamma}$. The set $\mathcal{U}(\alpha, \widetilde{\gamma})$ becomes more inclusive as the horizon of uncertainty, α, increases. The horizon of uncertainty, α, is unknown, so this is an unbounded family of nested sets of γ-vectors. We have estimates, $\widetilde{\gamma}$, but we are unable to specify maximum deviations of these estimates from the true values, and we know no probability distribution for the error of $\widetilde{\gamma}$.

Uncertain utility coefficients $u_{\rm e}$ and $u_{\rm h}$. The individual's personal utility coefficients, $u_{\rm e}$ and $u_{\rm h}$, reflect introspection, conversations with family, friends and physicians, social norms and constraints, and so on. (We will consider the disutility of intervention, $u_{\rm i}(\ell)$, separately.) The choice of $u_{\rm e}$ and $u_{\rm h}$ is highly uncertain. Whatever method the patient uses to estimate his or her utilities, we denote the estimates by the vector $\widetilde{u} = (\widetilde{u}_{\rm e}, \widetilde{u}_{\rm h})$. $u_{\rm e}$ is a disutility and thus must be negative, while $u_{\rm h}$ must be positive. The same holds for the best estimates, $\widetilde{u}_{\rm e}$ and $\widetilde{u}_{\rm h}$. We have no information about how differently the patient might choose the utility coefficients in other circumstances, for instance following experience of a subsequent cardiac event. Lacking more detailed information about the uncertainty in the utility coefficients, u, we adopt the following fractional-error info-gap model. Let us define the following function: $h(x) = x$ if $x \ge 0$, and $h(x) = 0$ if $x < 0$. The info-gap model for uncertainty in the utility coefficients is:

$$\begin{aligned} \mathcal{U}_{\rm u}(\alpha, \widetilde{u}) = \{ \, u : \ (1+\alpha)\widetilde{u}_{\rm e} \le \quad u_{\rm e} \quad \le h(1-\alpha)\widetilde{u}_{\rm e}, \\ h(1-\alpha)\widetilde{u}_{\rm h} \le \quad u_{\rm h} \quad \le (1+\alpha)\widetilde{u}_{\rm h} \, \}, \quad \alpha \ge 0 \end{aligned} \tag{4}$$

This info-gap model is an unbounded family of nested sets of utility vectors, $u = (u_{\rm e}, u_{\rm h})$. In the absence of uncertainty (when $\alpha = 0$) the set contains only the estimate, $\widetilde{u}$. The set $\mathcal{U}_{\rm u}(\alpha, \widetilde{u})$ becomes more inclusive as the horizon of uncertainty, α, increases, and α is unbounded so there is no known worst case.

Uncertain disutility of intervention, $u_{\rm i}(\ell)$. The disutility of intervention, $u_{\rm i}(\ell)$, is highly uncertain. The monetary cost varies with patient condition [42], full lifetime costs are difficult to identify and evaluate [43], and the impact of lifestyle change is hard to evaluate beforehand. It is hard to identify individuals who currently are free of disease but would develop disease in the future if intervention is not taken now [44]. Utility coefficients can be elicited from patients who are not diagnosed with a disease by describing the disease to the patient. However, the patient's assessments of utility may depend on knowing a positive diagnosis or not. In short, there are many sources of uncertainty in the patient's utility. One can formulate various different plausible disutility functions, and the large attendant uncertainty can be represented with an info-gap model as we now explain.

Let $\ell_{\rm c}$ denote the patient's current LDL level. The disutility, $u_{\rm i}(\ell)$, of moving to the ℓth LDL level will increase as the difference between $\ell_{\rm c}$ and ℓ increases. Furthermore, the disutility of intervention is estimated to be proportional to the spread between the (positive) utility of health,

$u_{\rm h}$, and the (negative) utility of a cardiac event, $u_{\rm e}$. A plausible (though uncertain) model for the disutility of intervention would be:

$$\widetilde{u}_{\rm i}(\ell) = |\ell - \ell_{\rm c}| u_0 \tag{5}$$

where u_0 is the following negative value:

$$u_0 = -\frac{(\widetilde{u}_{\rm h} - \widetilde{u}_{\rm e}) f}{4} \tag{6}$$

The term $\widetilde{u}_{\rm h} - \widetilde{u}_{\rm e}$ expresses the spread in estimated utility between health ($\widetilde{u}_{\rm h}$, which is positive) and disease ($\widetilde{u}_{\rm e}$, which is negative). Large damage of a cardiac event makes $\widetilde{u}_{\rm e}$ very negative and the spread very large. u_0 (which is negative) is a disutility that is large when the damage is large. The '4' in the denominator is the greatest possible value of $|\ell - \ell_{\rm c}|$. Note that if $f = 0$ then there is no disutility of intervention, while a large value of f implies large estimated disutility.

The function $\widetilde{u}_{\rm i}(\ell)$ in eq.(5) is plausible, but we don't have much actual evidence that the dependence is really linear, or for a specific value of u_0. In short, the magnitude and shape of this disutility function is highly uncertain. We will now formulate an info-gap model for uncertainty in this disutility function.

An info-gap model expresses the unknown deviation of the true function, $u_{\rm i}(\ell)$, from the plausible function $\widetilde{u}_{\rm i}(\ell)$. We require that $u_{\rm i}(\ell)$ be negative. We are considering uncertainty in the actual shape of the function, not just in its parameters. We use a "slope-bound" model, in which all the functions are negative, reach zero when $\ell = \ell_{\rm c}$, and the fractional error in the slope is unknown and unbounded. This implies that cost will increase with the magnitude of the intervention, though the slope (and magnitude) of the cost function is uncertain. $u'_{\rm i}(\ell)$ is the slope of the disutility function, which is negative for $\ell > \ell_{\rm c}$ and non-negative for $\ell \le \ell_{\rm c}$. Explicitly, the info-gap model is the following unbounded family of nested sets of functions:

$$\begin{aligned} \mathcal{U}_{\rm i}(\alpha, \widetilde{u}_{\rm i}) = \{\, u_{\rm i}(\ell) : \quad & u_{\rm i}(\ell_{\rm c}) = 0, \quad u_{\rm i}(\ell) \le 0, \quad \forall\, \ell, \\ (1+\alpha)u_0 \le \quad & u'_{\rm i}(\ell) \quad \le h(1-\alpha)u_0, \quad \ell > \ell_{\rm c} \\ -h(1-\alpha)u_0 \le \quad & u'_{\rm i}(\ell) \quad \le -(1+\alpha)u_0, \quad \ell < \ell_{\rm c} \,\}, \quad \alpha \ge 0 \end{aligned} \tag{7}$$

Recall that $u_0 < 0$, as defined in eq.(6).

4.2 Robustness and Opportuneness

We now define robustness and opportuneness functions, and explain how they assist patient and physician in choosing an LDL level.

Robustness. If we confidently knew $P_\ell(\gamma)$, $u_{\rm h}$, $u_{\rm e}$ and $u_{\rm i}(\ell)$, then we could confidently choose the LDL level, ℓ, to maximize the expected utility, $E_\ell(\gamma, u)$. However, $P_\ell(\gamma)$, $u_{\rm h}$, $u_{\rm e}$ and $u_{\rm i}(\ell)$ are highly uncertain. Consequently we evaluate how reliably we can obtain an acceptable outcome with different ℓ's. Conversely, what constitutes an acceptable outcome is influenced by how confident we are in attaining that outcome with a chosen LDL level, ℓ.

$E_{\rm c}$ denotes the lowest expected utility the patient would accept. Neither ℓ nor $E_{\rm c}$ are specified, and the patient may be unsure about what $E_{\rm c}$ value to require. We use a robustness concept [5] to choose ℓ and evaluate different $E_{\rm c}$'s.

We have estimates—likely to be off the mark—of $u_{\rm h}$, $u_{\rm e}$, $u_{\rm i}(\ell)$ and $P_\ell(\gamma)$. We wish to choose an LDL level, ℓ, for which adequate utility, $E_{\rm c}$, will be achieved even if these estimates err greatly.

The **robustness** of the ℓth LDL level is the greatest horizon of uncertainty, α, up to which expected utility $E_\ell(\gamma, u)$ is guaranteed to be no less than the critical utility, $E_{\rm c}$:

$$\widehat{\alpha}(\ell, E_{\rm c}) = \max \left\{ \alpha : \left(\min_{\substack{\gamma \in \mathcal{U}(\alpha, \widetilde{\gamma}) \\ u \in \mathcal{U}_{\rm u}(\alpha, \widetilde{u}) \\ u_{\rm i} \in \mathcal{U}_{\rm i}(\alpha, \widetilde{u}_{\rm i})}} E_\ell(\gamma, u) \right) \ge E_{\rm c} \right\} \tag{8}$$

The robustness function $\widehat{\alpha}(\ell, E_{\rm c})$ establishes preferences over the options. ℓ is preferred over ℓ', at critical aspiration $E_{\rm c}$, if we are more robust to the uncertainty with ℓ than with ℓ'. Succinctly:

$$\ell \succ_{\rm r} \ell' \quad \text{if} \quad \widehat{\alpha}(\ell, E_{\rm c}) > \widehat{\alpha}(\ell', E_{\rm c}) \tag{9}$$

The decision strategy implied by this preference ordering *satisfices* the expected utility and ranks the alternatives based on the *robustness.* For brevity, this strategy is called *robust-satisficing.*

We show how to conveniently evaluate the robustness function in appendix A.

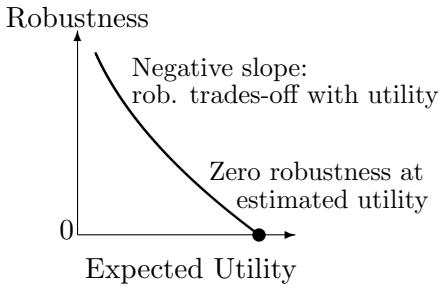

Figure 1: Properties of robustness curves.

Interpreting robustness curves. We will use plots of robustness, $\widehat{\alpha}(\ell, E_{\rm c})$, vs. expected utility, $E_{\rm c}$, to choose therapeutic intervention. We now explain how to interpret these plots, and illustrate that a choice of ℓ can usually be made with only a rough idea of the required critical utility, $E_{\rm c}$. Two points are characteristic of all robustness curves, illustrated in fig. 1.

First, the negative slope of the robustness curve in fig. 1 expresses the trade-off of robustness to uncertainty, against the expected utility which can be reliably anticipated. *Greater* expected utility is invariably associated with *lower* robustness to error in the underlying models. This intuitive idea—that high aspirations are more vulnerable to surprises or errors than low aspirations—is quantified by the negative slope.

Second, the robustness curve reaches the horizontal axis at some value of expected utility, as shown by the solid dot in fig. 1. The robustness is zero at this value of utility. Significantly, this value of expected utility is precisely the value obtained with the estimated models and data. In other words, the best estimate of the expected utility has no robustness against error in the models and data upon which that estimate depends. Best estimates exploit all available information, without regard to their potential error, so the estimated utility has no immunity against error in this information. Only less optimistic estimates—for lower expected utility—have positive robustness, which returns us again to the previous point: the trade-off between robustness and utility.

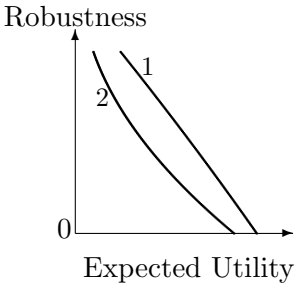

Figure 2: Preference for intervention 1.

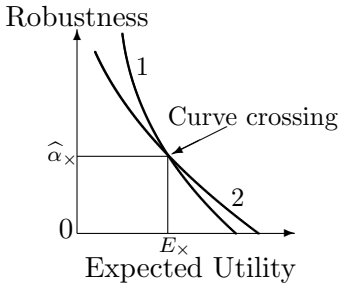

Figure 3: Preference reversal between therapies.

We now examine figs. 2 and 3 to understand how robustness curves are used to evaluate and choose among alternative interventions.

Fig. 2 shows the robustness curves of two different therapeutic interventions, denoted '1' and '2'. Option '1' is more robust than '2' at all levels of expected utility with positive robustness. Since more robustness is preferable over less, option '1' is unambiguously favored over option '2'. Note that this is the same choice as the guideline recommendation based on the estimated utilities of these two options (represented by the points at which the curves meet the horizontal axis).

Fig. 3 shows a different situation, in which the robustness curves for two different therapeutic interventions intersect at some value of expected utility, $E_\times$, and robustness, $\widehat{\alpha}_\times$. More robustness is better than less robustness. However, when utility is sacrificed for robustness, with it goes the aspiration for the higher results promised by the model. If utility in excess of $E_\times$ is needed, then

option '2' is preferred. This is the same choice as the guideline recommendation based on the estimated utilities of these two options. We must recognize, however, that the trade-off between robustness and utility implies that utility in excess of $E_\times$ will be obtained only with robustness in the lower part of the robustness scale, below robustness $\widehat{\alpha}_\times$. If greater confidence (greater robustness) is needed, and the patient and physician are willing to obtain this robustness premium by relaxing the aspiration for utility to a value below $E_\times$, then alternative '1' is preferred. In this case the crossing of the robustness curves has caused the decision maker to change preference from the guideline option. The decision maker does not need to choose a specific value of either critical utility, E_c, or robustness, $\widehat{\alpha}_\times$, in order to gain the advantage of this formulation. Only a choice of the required range of these values is needed.

Opportuneness. Robustness against failing to achieve acceptable utility is only one part of managing uncertainty. Uncertainty also entails opportunities for windfall, that is, a benefit or outcome better than anticipated [5]. A more opportune choice is preferred, though this may disagree with the robustness preferences.

E_w denotes a large and highly desirable level of utility, a windfall, larger than the critical utility E_c and larger than the utility based on the estimated models. The **opportuneness** of the ℓth LDL level is the lowest horizon of uncertainty, α, at which expected utility $E_\ell(\gamma, u)$ is possibly (though not necessarily) as large as the windfall value E_w:

$$\widehat{\beta}(\ell, E_\mathrm{w}) = \min\left\{ \alpha : \left(\max_{\substack{\gamma\in\mathcal{U}(\alpha,\widetilde{\gamma}) \\ u\in\mathcal{U}_\mathrm{u}(\alpha,\widetilde{u}) \\ u_\mathrm{i}\in\mathcal{U}_\mathrm{i}(\alpha,\widetilde{u_\mathrm{i}})}} E_\ell(\gamma, u) \right) \geq E_\mathrm{w} \right\} \tag{10}$$

The opportuneness function is the lowest level of uncertainty which enables better-than-anticipated results, E_w. If this level of uncertainty is large, then better-than-anticipated results (windfalls) will require extraordinary circumstances; if this level of uncertainty is small, then windfall is possible (though not guaranteed) even in nearly ordinary situations. Thus, a small value of $\widehat{\beta}(\ell, E_\mathrm{w})$ means that windfall is feasible, and decision ℓ is opportune. A large value of $\widehat{\beta}(\ell, E_\mathrm{w})$ means that great uncertainty is needed in order to enable windfall as large as E_w. We can summarize this by saying that $\widehat{\beta}(\ell, E_\mathrm{w})$ assesses the degree to which intervention ℓ is immune to windfall outcomes: large $\widehat{\beta}(\ell, E_\mathrm{w})$ implies high immunity to windfall and low opportuneness; small $\widehat{\beta}(\ell, E_\mathrm{w})$ implies low immunity and high opportuneness.

In short, the opportuneness function, $\widehat{\beta}(\ell, E_\mathrm{w})$, is the *immunity against windfall.* Since windfall is desirable, small values of $\widehat{\beta}(\ell, E_\mathrm{w})$ (low immunity to windfall) are preferable over large values. Hence the opportuneness function establishes preferences over the options. ℓ is preferred over ℓ', at windfall aspiration E_w, if we are more prone to favorable uncertainty with ℓ than with ℓ'. Thus ℓ is preferred over ℓ' if $\widehat{\beta}(\ell, E_\mathrm{w})$ is less than $\widehat{\beta}(\ell', E_\mathrm{w})$:

$$\ell \succ_\mathrm{o} \ell' \quad \text{if} \quad \widehat{\beta}(\ell, E_\mathrm{w}) < \widehat{\beta}(\ell', E_\mathrm{w}) \tag{11}$$

The decision strategy implied by this preference ordering attempts to facilitate *windfalls* (which we will refer to as 'windfalling'), and ranks the alternatives based on the *opportuneness.* For brevity, this strategy is called *opportune-windfalling.*

The preference rankings in eqs.(9) and (11) may or may not agree.

We show how to conveniently evaluate the opportuneness function in appendix A.

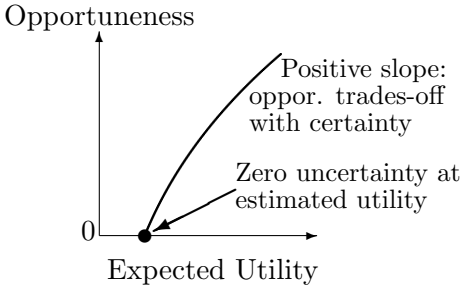

Figure 4: Properties of opportuneness curves.

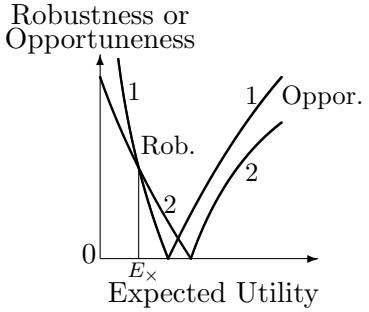

Figure 5: Using opportuneness and robustness.

Interpreting opportuneness curves. Fig. 4 shows a typical opportuneness curve, illustrating the analog of the two robustness properties described in fig. 1. The positive slope in fig. 4 expresses the trade-off between opportuneness and certainty: high windfall aspiration (for larger-than-expected utility) is feasible only at great ambient uncertainty. Likewise, the opportuneness curve hits the horizontal axis precisely at the estimated utility: no surprise is needed in order to enable (though not guarantee) the expected outcome.

Since decision makers are usually risk averse, the opportuneness curves are usually used to resolve situations where robustness is ambiguous, as illustrated in fig. 5. The decreasing curves on the left of the figure are the robustness curves for two alternative interventions. These curves intersect at expected utility $E_\times$. Consequently, if utility around $E_\times$ is required, consideration of robustness does not resolve the choice between these options, since the two options are equally robust. The opportuneness curves—with positive slope—do not intersect one another in this case. Since a small value for $\widehat{\beta}(\ell, E_\mathrm{w})$ is preferred, we see that option 2 is more opportune than option 1, suggesting a resolution of the robustness-ambiguity in favor of option 2.

5 Results: 4 Case Studies

In this section we illustrate the info-gap decision analysis with uncertainty in the probability of a cardiac event, the patient's utility coefficients for disease and health, and the cost of intervention. We consider 4 hypothetical, but realistic, case studies.

Cases 1 and 2, sections 5.1 and section 5.2, compare the decision analysis for a middle-aged male with difference in HDL level, blood pressure, smoking and diabetes.

Case 3, section 5.3, looks at the effect of age, with other variables unchanged.

Case 4, section 5.4, considers a female of age 55.

5.1 Case 1

Group	HDL-C, mg/dL	Blood pressure	LDL-C, mg/dL
1	< 35	Optimal	< 100
2	35–44	Normal	100–129
3	45–49	High normal	130–159
4	50–59	Stage I hypertension	160–189
5	≥ 60	Stage II–IV hypertension	≥ 190

Table 2: HDL-C, blood pressure and LDL-C groups. The systolic and diastolic pressures defining the blood pressure groups are defined in table 1 of Wilson [33].

Example 1 is a male, aged 55, group-5 HDL, group-1 blood pressure (see table 2), non-diabetic non-smoker whose pre-intervention LDL level is group 5. The info-gap analysis examines various estimated (but uncertain) costs of intervention, together with uncertainty in the probability of a cardiac event and uncertainty in the patient's utility coefficients for disease and health.

We will discuss robustness and opportuneness curves (figs. 6 and 7) for transition to various LDL levels ℓ. Before doing so we explain how these curves are constructed.

The patient's utility coefficients, $\widetilde{u}_\mathrm{h}$, $\widetilde{u}_\mathrm{e}$ and $\widetilde{u}_\mathrm{i}(\ell)$ are estimated by one or another method as described in section 4.1. One outcome of this estimation process is that the patient develops subjective understanding of the meaning, in terms of linguistic descriptors such as 'poor' or 'pretty good', of different ranges of utility on the horizontal axes of these figures. The central task of the uncertainty analysis is to address the following questions. First, for any given intervention, what levels of utility can be reliably attained; this is based on the robustness function. Second, for any given intervention, what levels of utility can be aspired to as potential windfall; this exploits the opportuneness function.

The robustness and opportuneness functions are evaluated for the individual patient based on the patient's estimated utility coefficients. Thus the curves are individualized to the patient, expressing the preferences, and the uncertainties, for that individual. The evaluation of robustness

and opportuneness is specified mathematically in the appendix. A different robustness and opportuneness curve is calculated for each possible intervention. Thus, since 5 LDL levels are possible, there are 5 robustness and 5 opportuneness curves.

The disutility of intervention, $u_{\mathrm{i}}(\ell)$, for moving to the ℓth LDL level is estimated to increase as the difference between ℓ and the patient's current LDL level, ℓ_{c}, increases. Furthermore, the disutility is estimated to be proportional to a fraction f of the utility of health u_{h} minus the disutility of a cardiac event u_{e}. These subjective estimates, however, are highly uncertain.

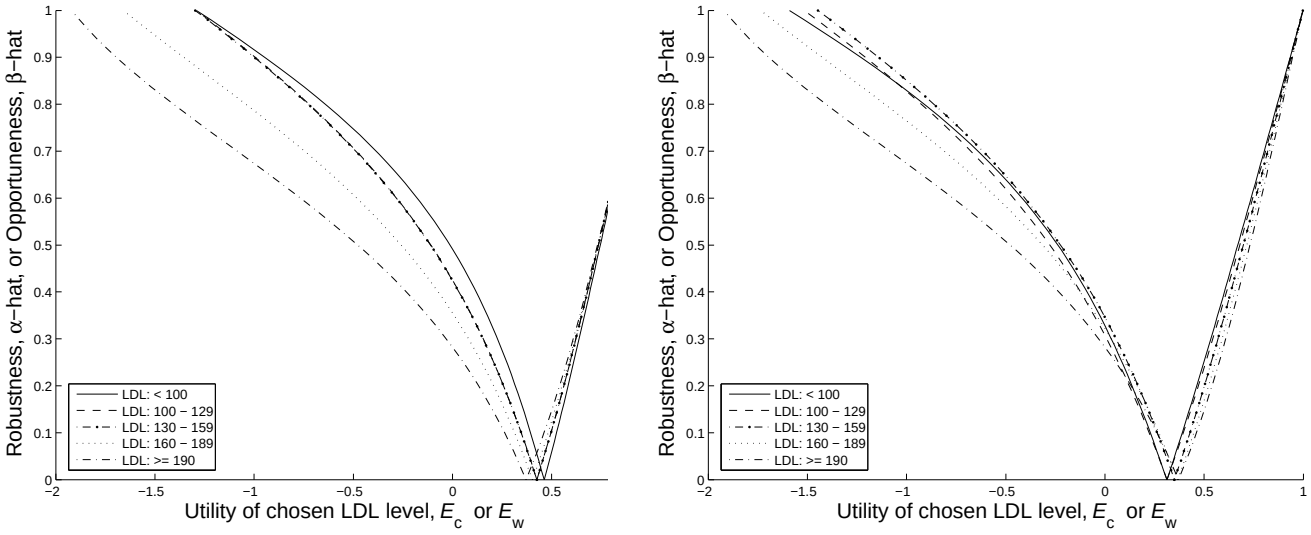

Figure 6: Case 1. Robustness and opportuneness vs. desired utility for case in section 5.1. Patient's estimated utilities are $\widetilde{u} = [-1,\ 0.5]$. ρ is the unit vector. $\ell_{\mathrm{c}} = 5$. $f = 0$ (left), $f = 0.1$ (right).

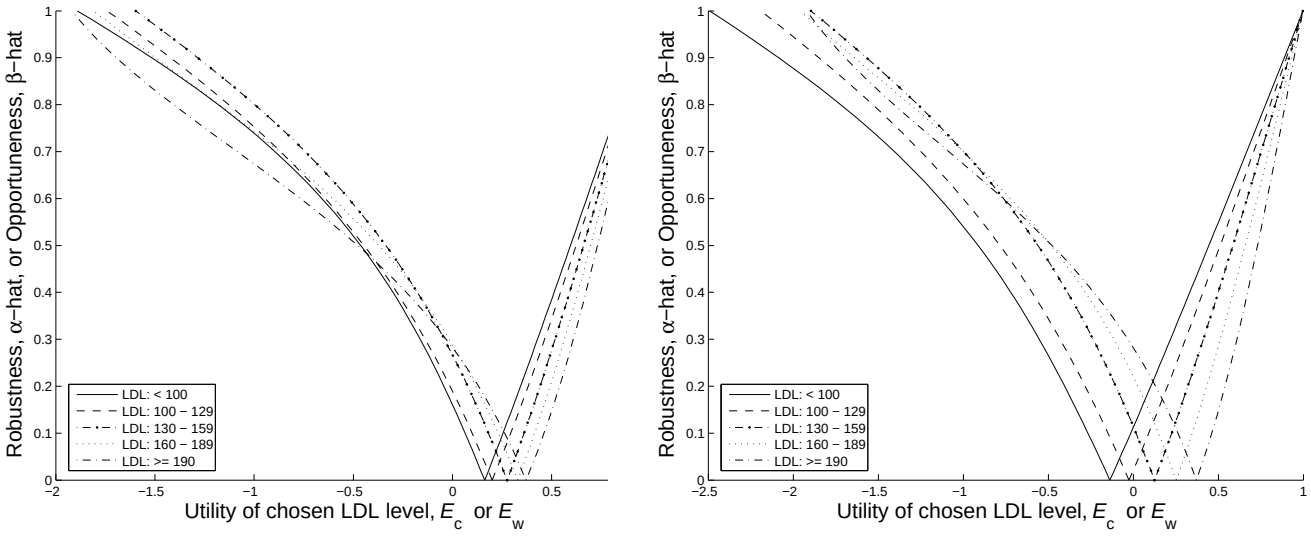

Figure 7: Case 1. Robustness and opportuneness vs. desired utility for case in section 5.1. Patient's estimated utilities are $\widetilde{u} = [-1,\ 0.5]$. ρ is the unit vector. $\ell_{\mathrm{c}} = 5$. $f = 0.2$ (left), $f = 0.4$ (right).

Figs. 6 and 7 show robustness and opportuneness curves for transition to various LDL levels ℓ. The left frame of fig. 6 is the case of no cost at any level of uncertainty ($f = 0$), while the right frame of fig. 7 shows an uncertain cost estimated as 40% of the benefit ($f = 0.4$). The other frames of figs. 6 and 7 show intermediate values ($f = 0.1$ and $f = 0.2$) for the estimated cost of intervention.

The robustness curves in the left frame of fig. 6 are for the case of no cost of intervention. The

nominal preferences, based on the best estimates—at which the robustness is zero as explained in connection with fig. 1—are for decreasing the LDL level as much as possible. Level 1 is preferred over level 2 (represented by "$\succ$"). Levels 2 and 3 are identical (represented by "$\sim$"), since their estimated γ coefficients are essentially the same (table 1). However, they are both preferred over level 4, which is preferred over level 5. We can succinctly represent these preferences as:

$$1 \succ 2 \sim 3 \succ 4 \succ 5 \tag{12}$$

The robustness to uncertainty for each of these options, for attaining the corresponding estimated utility, is zero, as explained in fig. 1. Consequently these options can not be depended on to reliably result in the corresponding utilities. However, the preference ranking is the same at every level of robustness up to 100% ($\widehat{\alpha} = 1$), since the robustness curves have not crossed one another up to robustness of 1. The situation here is like fig. 2, in which the robustness curves do not cross one another. It is noteworthy, though, that levels 4 and 5 have substantially less utility than lower LDL levels at robustness of 0.5 and greater. This is unlike the situation at zero robustness, in which the utility-margin between the different options is small. At robustness of $\widehat{\alpha} = 1$ the robustness curve for $\ell = 1$ crosses the curves for levels 2 and 3, thus reversing the preference among these levels if larger robustness is required, noting that the corresponding utility is quite low compared to the nominal utility.

The opportuneness curves in the lefthand frame of fig. 6 are ordered nominally, and rapidly converge. This indicates that considerations of opportuneness do not promote one intervention over another.

The situation is very different in the righthand frame of fig. 6, for which the estimated (though uncertain) cost of intervention is 10% of the estimated benefit, so $f = 0.1$. Now the nominal (zero-robustness) preferences are the reverse of the no-cost preferences in eq.(12):

$$5 \succ 4 \succ 3 \succ 2 \sim 1 \tag{13}$$

The curve for $\ell = 5$ has not shifted between the two frames of fig. 6 because the patient's pre-intervention LDL level is $\ell_c = 5$. However the other curves have shifted to the right, and changed in shape, as a result of the cost of intervention. These nominal, zero-robustness preferences are weak. However, there is substantial crossing of robustness curves at low robustness (around $\widehat{\alpha} \approx 0.2$), causing preference-reversal among the interventions as explained in connection with fig. 3. At 50% robustness ($\widehat{\alpha} = 0.5$) there is strong preference against level 5, despite the intervention cost, while the preferences among the other levels is not large. More precisely, the preferences at 50% robustness are:

$$(3 \succ 1) \succ 2 \succ 4 \succ 5 \tag{14}$$

where the parentheses indicate that the preferences between levels 3 and 1 are weak.

A similar picture emerges when we consider higher estimated cost of intervention, $f = 0.2$ and $f = 0.4$ in fig. 7. The results of these 4 figures are summarized in table 3.

Before examining table 3, consider the righthand frame of fig. 7 at moderate robustness ($\widehat{\alpha} = 0.5$). The robustness curves for the two most attractive interventions, levels 4 and 5, cross at $\widehat{\alpha} = 0.5$, indicating indifference between these two options, where level 5 implies no intervention since this is the patient's pre-intervention LDL level. However, examining the opportuneness curves in this figure shows that remaining at level 5 is more opportune for windfall outcome than moving to level 4. This is an instance where the opportuneness analysis can be used to resolve a robustness-tie between options, as discussed in connection with fig. 5.

A clear picture emerges from table 3 for this individual, whose pre-intervention LDL concentration is at level 5. If there is no cost for intervention ($f = 0$), then the nominal, zero robustness ($\widehat{\alpha} = 0$), preference is for reducing the LDL concentration to level 1. If there is any cost to intervention ($f \geq 0.1$), then the nominal (no-robustness) preferences call for no intervention. However, requiring moderate robustness to uncertainty ($\widehat{\alpha} = 0.5$) changes the picture, and intervention to lower the LDL to the 3rd level is indicated at moderate cost ($f = 0.1$ or $f = 0.2$). At high cost, $f = 0.4$, levels 4 and 5 are equivalent in robustness, and opportuneness mitigates for no intervention (level 5). The enhanced robustness resulting from these decisions is obtained at the expense of guaranteeing lower utility, as understood from the trade-off between robustness and utility (fig. 1). The picture remains the same at high robustness ($\widehat{\alpha} = 1$), and correspondingly low utility.

Cost f	No robustness $\widehat{\alpha} = 0$	Moderate robustness $\widehat{\alpha} = 0.5$	High robustness $\widehat{\alpha} = 1$
0.0	$\mathbf{1} \succ 2 \sim 3 \succ 4 \succ 5$	$\mathbf{1} \succ 2 \sim 3 \succ 4 \succ 5$	$1 \sim 2 \sim \mathbf{3} \succ 4 \succ 5$
0.1	$\mathbf{5} \succ 4 \succ 3 \succ 2 \sim 1$	$(\mathbf{3} \succ 1) \succ 2 \succ 4 \succ 5$	$\mathbf{3} \succ 2 \succ 1 \succ 4 \succ 5$
0.2	$\mathbf{5} \succ 4 \succ 3 \succ 2 \succ 1$	$\mathbf{3} \succ 4 \succ (2 \succ 1 \succ 5)$	$\mathbf{3} \succ 2 \succ 4 \succ (1 \succ 5)$
0.4	$\mathbf{5} \succ 4 \succ 3 \succ 2 \succ 1$	$\mathbf{5} \sim 4 \succ 3 \succ 2 \succ 1$	$3 \sim \mathbf{5} \succ 4 \succ 2 \succ 1$

Table 3: Preference ranks of LDL options for case 1, based on robustness, from figs. 6 and 7.

5.2 Case 2

Let's consider a case whose prognosis is quite different from the case in section 5.1: a male, aged 55, group-1 HDL, group-5 blood pressure, diabetic smoker whose current (pre-intervention) LDL level is group 5. We will again consider the info-gap analysis for various estimated (but uncertain) costs of intervention, together with uncertainty in the probability of a cardiac event and uncertainty in the patient's utility coefficients for disease and health.

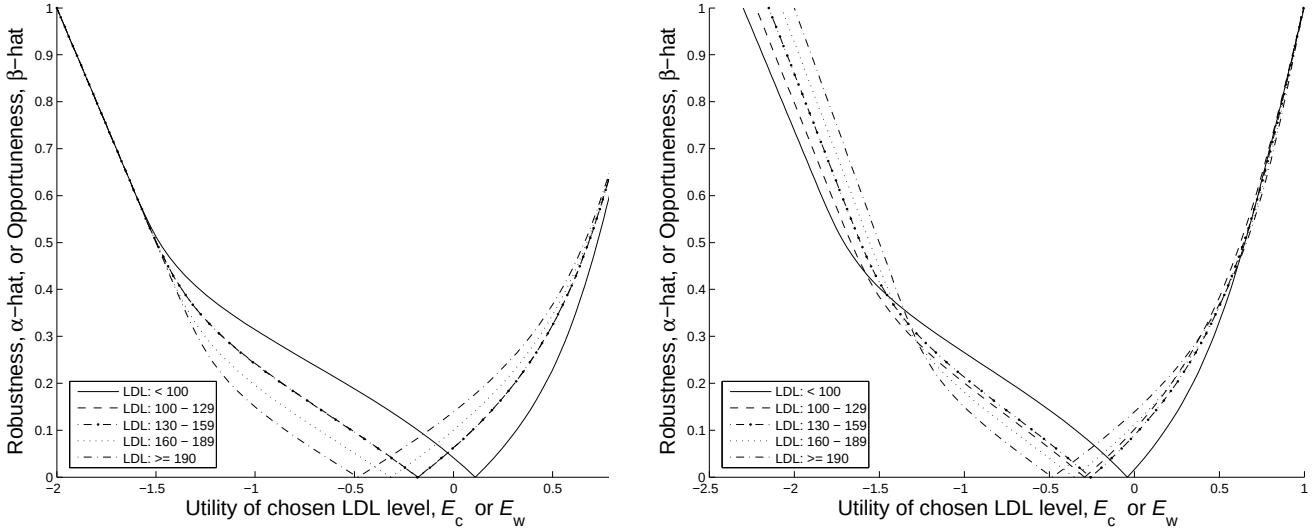

Figure 8: Case 2. Robustness and opportuneness vs. desired utility for case in section 5.2. Patient's estimated utilities are $\widetilde{u} = [-1,\ 0.5]$. ρ is the unit vector. $\ell_c = 5$. $f = 0$ (left), $f = 0.1$ (right).

Results for 4 different estimated costs of intervention are shown in figs. 8 and 9. The nominal best-estimated utilities of the options—at zero robustness—are lower than in case 1 (section 5.1, figs. 6 and 7) at all estimated costs. Likewise, at any given level of utility, the robustnesses are lower in case 2 than in case 1 for all estimated costs. The crossing of robustness occurs as in case 1, though at lower robustness and lower utility.

The picture which emerges for this case, in table 4, is somewhat different than for case 1 in table 3. In case 2, as in case 1, if there is no cost for intervention ($f = 0$), then the nominal, zero robustness ($\widehat{\alpha} = 0$) preference is for reducing the LDL concentration to level 1. However, the robustness curves rapidly converge at $f = 0$, so that at moderate and high robustness, $\widehat{\alpha} = 0.5$ and $\widehat{\alpha} = 1$, there is no preference among the options in terms of robustness, and very little preference based on opportuneness, though the opportuneness curves are separated and do not cross one another. If there is no cost of intervention, then one presumably would not intervene in the absence of additional considerations. At low or moderate cost of intervention, $f = 0.1$ or $f = 0.2$, the nominal, zero-robustness preference is for minimizing the LDL level. At high cost, $f = 0.4$, the zero-robustness preference is indifferent between all levels except $\ell = 2$ which is less preferred; again one would presumably not intervene. However, at moderate or high robustness ($\widehat{\alpha} = 0.5$ or $\widehat{\alpha} = 1$) and positive cost ($f \geq 0.1$), the preference is for $\ell = 5$, which is the patient's current state,

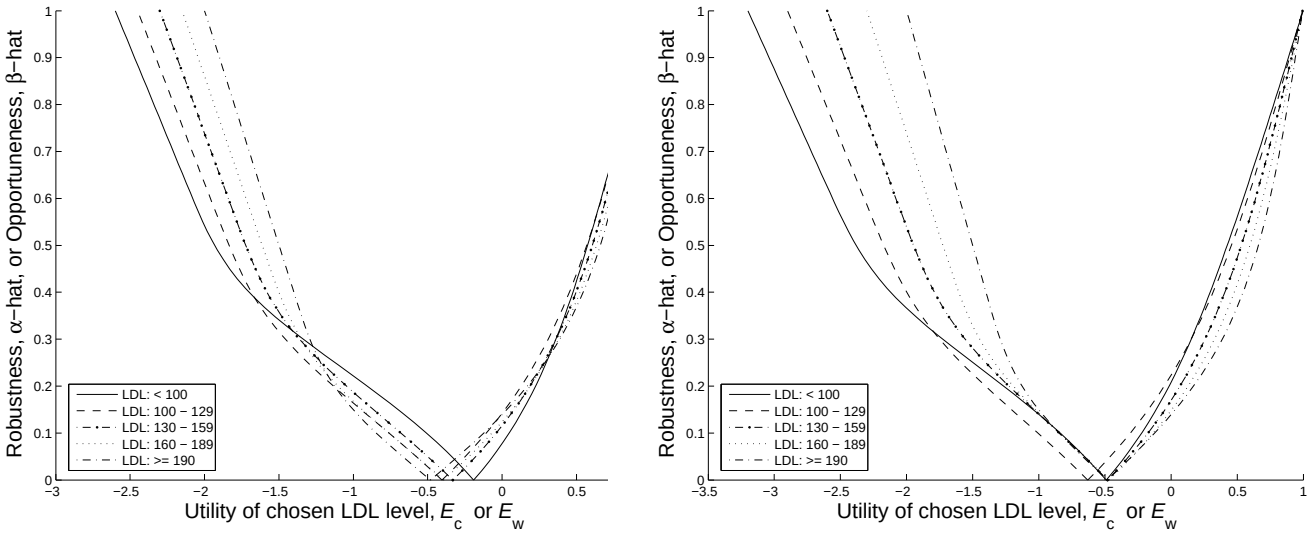

Figure 9: Case 2. Robustness and opportuneness vs. desired utility for case in section 5.2. Patient's estimated utilities are $\widetilde{u} = [-1,\ 0.5]$. ρ is the unit vector. $\ell_c = 5$. $f = 0.2$ (left), $f = 0.4$ (right).

Cost f	No robustness $\widehat{\alpha} = 0$	Moderate robustness $\widehat{\alpha} = 0.5$	High robustness $\widehat{\alpha} = 1$
0.0	$\mathbf{1} \succ 2 \sim 3 \succ 4 \succ 5$	$1 \sim 2 \sim 3 \sim 4 \sim \mathbf{5}$	$1 \sim 2 \sim 3 \sim 4 \sim \mathbf{5}$
0.1	$\mathbf{1} \succ (3 \succ 2) \succ 4 \succ 5$	$\mathbf{5} \succ 4 \succ 3 \succ 2 \succ 1$	$\mathbf{5} \succ 4 \succ 3 \succ 2 \succ 1$
0.2	$\mathbf{1} \succ 3 \succ 2 \sim 4 \succ 5$	$\mathbf{5} \succ 4 \succ 3 \succ 2 \succ 1$	$\mathbf{5} \succ 4 \succ 3 \succ 2 \succ 1$
0.4	$1 \sim 3 \sim 4 \sim \mathbf{5} \succ 2$	$\mathbf{5} \succ 4 \succ 3 \succ 2 \succ 1$	$\mathbf{5} \succ 4 \succ 3 \succ 2 \succ 1$

Table 4: Preference ranks of LDL options for case 2, based on robustness, from figs. 8 and 9.

so intervention is not indicated here either.

We see here, as in case 1, the strong effect of considering robustness to uncertainty in cost and the other factors of the estimated decision model.

5.3 Case 3

In this example we explore the effect of age and estimated cost of intervention. Case 3 is a non-smoking non-diabetic male with group-5 HDL and group-1 blood pressure, at pre-intervention LDL level 5. We compare ages 35, 55 and 75.

Consider first the robustness curves without intervention cost, fig. 10 (ages 35 and 75), and the left frame of fig. 6 (age 55). The main effects of advancing age are to reduce the zero-robustness nominal utility, which shifts the curves to the left, and to substantially reduce the robustness at any lower utility, which lowers the robustness curves.

Age (y)	No robustness $\widehat{\alpha} = 0$	Moderate robustness $\widehat{\alpha} = 0.5$	High robustness $\widehat{\alpha} = 1$
35	$\mathbf{5} \succ 4 \succ 3 \succ 2 \succ 1$	$\mathbf{5} \succ 4 \succ 3 \succ 2 \succ 1$	$(3 \succ \mathbf{4}) \succ (2 \succ 5) \succ 1$
55	$\mathbf{5} \succ 4 \succ 3 \succ 2 \succ 1$	$\mathbf{3} \succ 4 \succ (2 \succ 1 \succ 5)$	$\mathbf{3} \succ 2 \succ 4 \succ (1 \succ 5)$
75	$(\mathbf{5} \succ 4 \succ 3) \succ (1 \succ 2)$	$1 \sim 3 \sim (2 \succ \mathbf{4}) \succ 5$	$\mathbf{5} \succ 4 \succ 3 \succ 2 \succ 1$

Table 5: Preference ranks of LDL options for case 3, based on robustness and figs. 10 and 11, and the left frame of fig. 7. Results for $f = 0.2$.

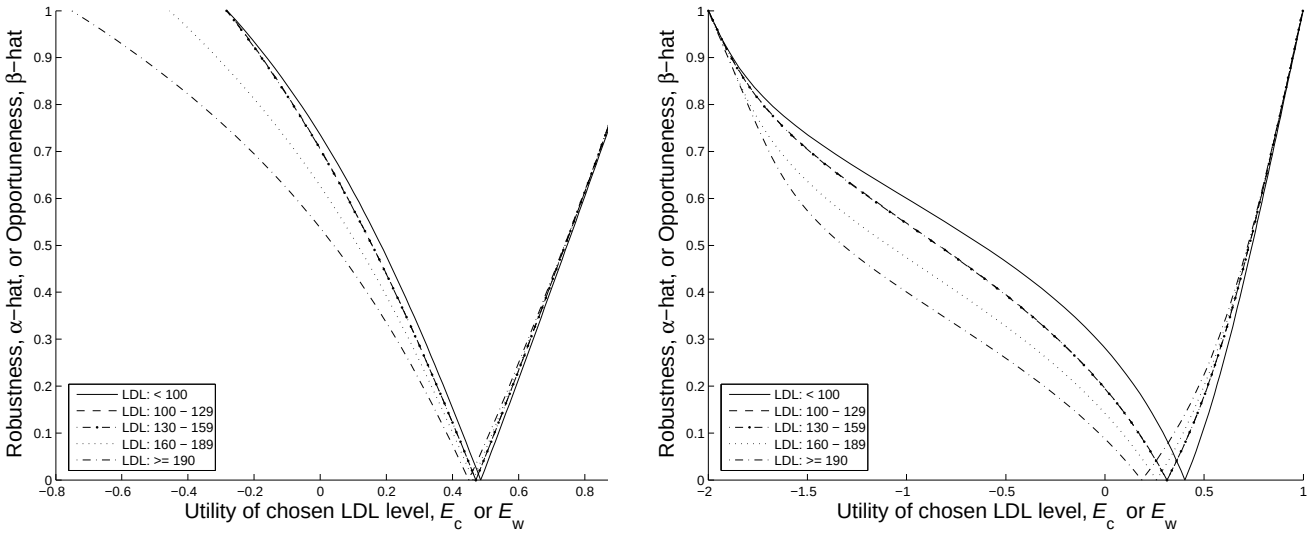

Figure 10: Case 3. Robustness and opportuneness vs. desired utility for case in section 5.3. Patient's estimated utilities are $\widetilde{u} = [-1,\ 0.5]$. ρ is the unit vector. $\ell_c = 5$, $f = 0$. Age 35 (left), age 75 (right).

Similar conclusions hold at moderate estimated cost of intervention, $f = 0.2$, as seen in fig. 11 and the left frame of fig. 7. Furthermore, the curves have moved with respect to one another sufficiently so that the preference rankings are different at the different ages, as illustrated in table 5, part of which (age 55) is reproduced from table 3. The nominal, zero-robustness recommendation is the same for all three ages: no intervention. At moderate robustness, $\widehat{\alpha} = 0.5$, the recommendations are: no intervention at age 35, lower LDL to level 3 at ages 55 and 75. At high robustness, $\widehat{\alpha} = 1$, the young and mid-age groups move to LDL level 3, while the oldest age group does not mandate intervention. We see in this example, as before, the strong effect of considering uncertainty in the underlying data and models.

5.4 Case 4

This example is a non-smoking non-diabetic aged-55 woman with HDL level 5, blood pressure group 1, and pre-intervention LDL level 5. This is the same as the case in section 5.1 except for the sex of the patient, so the figures for comparison are figs. 6 and 7 and table 3.

Figs. 12 and 13 shows results for four different estimated costs of intervention. The most striking result is the far lower robustness than for the similar male patient, figs. 6 and 7.

Cost f	No robustness $\widehat{\alpha} = 0$	Moderate robustness $\widehat{\alpha} = 0.08$	High robustness $\widehat{\alpha} = 0.2$
0.0	$\mathbf{1} \succ 2 \sim 3 \succ (4 \succ 5)$	$\mathbf{1} \succ 2 \sim 3 \succ (4 \succ 5)$	$1 \sim 2 \sim 3 \sim 4 \sim \mathbf{5}$
0.1	$\mathbf{5} \succ 4 \succ 3 \succ 2 \succ 1$	$\mathbf{1} \succ 3 \succ (2 \succ 5 \succ 4)$	$\mathbf{5} \succ 4 \succ 3 \succ 2 \succ 1$
0.2	$\mathbf{5} \succ 4 \succ 3 \succ 2 \succ 1$	$\mathbf{5} \succ (3 \succ 4) \succ 1 \succ 2$	$\mathbf{5} \succ 4 \succ 3 \succ 2 \succ 1$
0.4	$\mathbf{5} \succ 4 \succ 3 \succ 2 \succ 1$	$\mathbf{5} \succ 4 \sim 3 \succ 2 \succ 1$	$\mathbf{5} \succ 4 \succ 3 \succ 2 \succ 1$

Table 6: Preference ranks of LDL options for case 4, based on robustness and figs. 12 and 13.

We summarize the results of figs. 12 and 13 in table 6, which should be compared to table 3. 'Moderate' and 'High' robustnesses are lower for female than male patients, but the interventions are roughly similar. At zero robustness, both male and female cases call for no intervention if there is even low cost of intervention ($f \geq 0.1$), and lowering LDL to level 1 if $f = 0$. At moderate robustness ($\widehat{\alpha} = 0.08$) the recommendation for females is for no intervention at high cost ($f = 0.2$ or $f = 0.4$), and lowering to level 1 at zero or low cost ($f = 0$ or $f = 0.1$). At high robustness

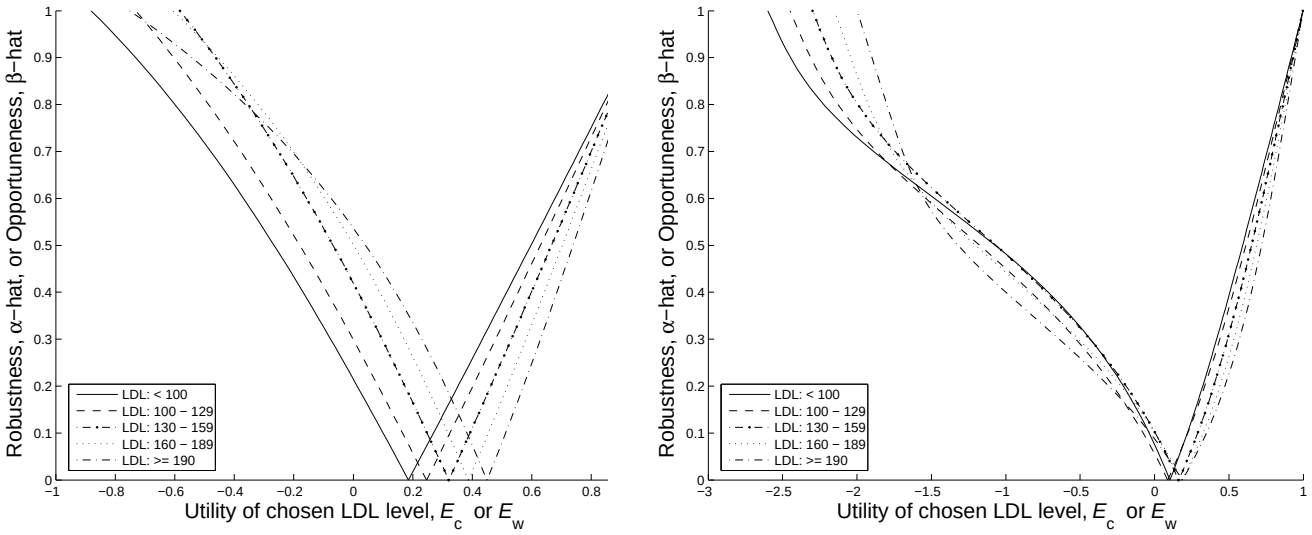

Figure 11: Case 3. Robustness and opportuneness vs. desired utility for case in section 5.3. Patient's estimated utilities are $\widetilde{u} = [-1,\ 0.5]$. ρ is the unit vector. $\ell_c = 5$, $f = 0.2$. Age 35 (left), age 75 (right).

($\widehat{\alpha} = 0.2$) the recommendation for females is strong intervention at zero cost, and no intervention otherwise.

6 Discussion: Medical Guidelines

Clinical guidelines have become central to medical practice. The physician is charged with evaluating a guideline's relevance for each patient, and knowing when and how to recommend a deviation from the guideline. Furthermore, the patient should be able to evaluate the guideline and express individual preferences. Missing from the process of considering clinical guidelines is the diverse and partly non-probabilistic uncertainty facing both physician and patient. Info-gap analysis explicitly incorporates this severe uncertainty into the decision making process, thus respecting individual judgment, variation in preference, and autonomy. A methodology is presented here for employing guidelines, in light of uncertain relevance of population studies to individual patients, uncertainty in patient preferences among outcomes, and uncertain cost of intervention. The uncertainties represented by info-gap models are both in the *values* of parameters and in the *shapes* of functions.

Info-gap robustness and opportuneness functions are useful for evaluating and selecting therapies. Robustness is defined as the greatest uncertainty at which acceptable outcomes will occur. Large robustness is preferred over small robustness, generating preferences among available therapies. Opportuneness is defined as the lowest uncertainty which is needed in order to enable a windfall outcome. Small values for the opportuneness function are preferred over large values, also generating preferences among options which need not agree with robustness-based preferences. The opportuneness preferences can be used to resolve ambiguous robustness preferences.

We have illustrated the use of robustness and opportuneness for selecting LDL levels in conjunction with guidelines, for plausible clinical situations. The following conclusions are suggested.

- Decisions using guidelines and best estimates of expected utility are sometimes supported by analysis of robustness to, and opportuneness from, uncertainty. That is, ranking the options in terms of estimated expected utility is sometimes the same as ranking the options in terms of robustness and opportuneness. When this occurs, the robustness and opportuneness analysis provides additional support for the guideline-based decision.
- Combined with guidelines, info-gap analysis can increase confidence in the choice of treatment. The degree of confidence in attaining an outcome is quantified by robustness: the numerical evaluation of the robustness indicates how wrong the estimates can be without jeopardizing one's aspiration for quality of outcome. When the robustness is low, as occurs in some cases, the decision maker will perhaps look further before deciding. Or, less risk-averse decision makers may use the opportuneness function to identify opportune therapies.

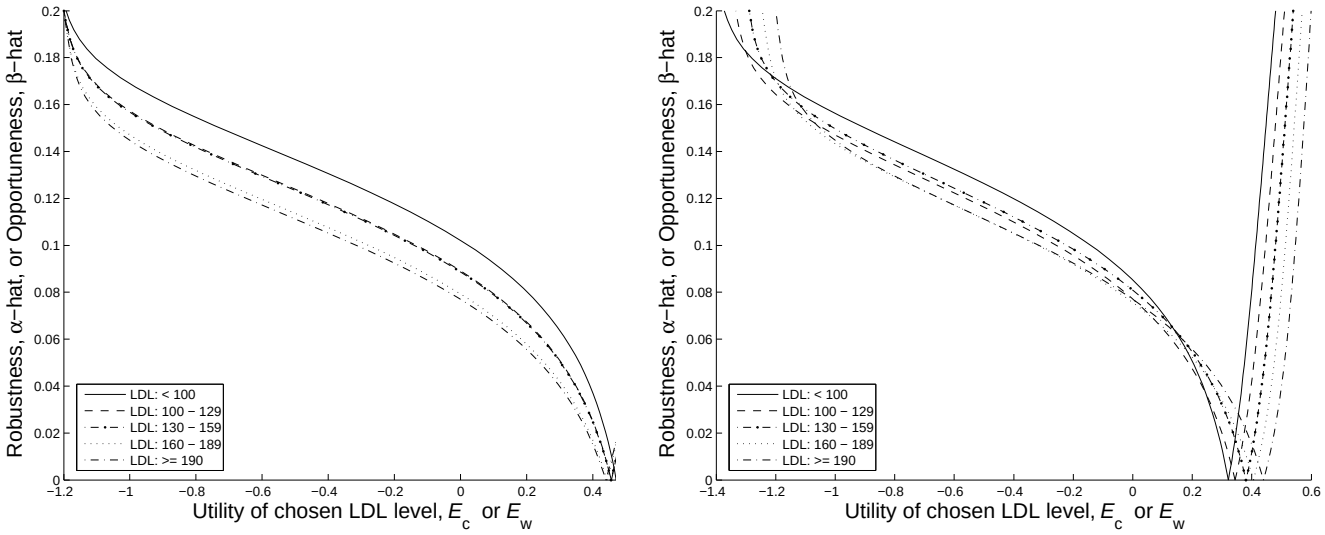

Figure 12: Case 4. Robustness and opportuneness vs. desired utility for case in section 5.4. Patient's estimated utilities are $\widetilde{u} = [-1,\ 0.5]$. ρ is the unit vector. $\ell_c = 5$, $f = 0$ (left), $f = 0.1$ (right).

- The analysis helps resolve ambiguity. In some situations the best-estimated outcome is essentially the same for several therapies (righthand frame of fig. 6 or fig. 9, or the left hand frame of fig. 10). One may choose the most robust option, which may differ from the guideline recommendation. For example, the nominal, guideline, recommendation in the righthand fig. 6 is for no intervention ($\ell = 5$), but the nominal preference for this choice is weak. The robustness curves cross one another at low robustness in this figure and the $\ell = 5$ option rapidly becomes substantially less robust than the other options, suggesting that the nominal recommendation is not suited to this case.
- The opportuneness function can be used to resolve situations where robustness does not establish unambiguous preferences. The left hand frame of fig. 13 is a case in point. The robustness curves are very close to one another over much of the utility-range, resulting in ambiguous robustness-preferences. However, the opportuneness curves are well separated, indicating strong opportuneness-preference for LDL level 5. The decision makers may choose to resolve the robustness-ambiguity by selecting the greater opportuneness of non-intervention (level 5) in this situation.
- The analysis provides guidance when judgment suggests deviation from the guideline. The left hand frame of fig. 9 illustrates this. The best estimates indicate guideline-based preference for the lowest LDL level. However, as always, best-estimated outcomes have zero robustness against uncertainty (fig. 1). The left hand frame of fig. 9 shows that LDL level 5 has far greater robustness than level 1, over much of the utility range. When the decision makers feel substantial uncertainty, they may opt for level 5, rather than 1, due to greater robustness of this option. Additionally, exogenous considerations of collateral medical impacts may be folded into preferring level 5, supported by greater robustness of this option.

Conventional application of clinical guidelines is often appropriate. A robust approach expands the usefulness of guidelines to individual patients and provides increased confidence in outcomes when dealing with diverse and substantial uncertainty. Further, coupling opportuneness to robustness reveals potential for better-than-expected outcomes. Additionally, uncertainty analysis may provide support when the physician's judgment indicates deviation from the guidelines. Finally, the value of clinical guidelines is enhanced by the availability of quantitative decision-theoretical tools, suggesting that guidelines are amenable to computer-based decision-support technology.

We make no claim of optimality for the info-gap analysis developed here. For decisions of such complexity, and with uncertainties of such diversity, we know of no feasible criterion for optimality of a decision methodology. We claim only that diverse and clinically realistic information and uncertainty can be efficiently and meaningfully managed with info-gap theory. Furthermore, just as we have combined non-probabilistic info-gap tools with probabilistic ones (expected utility theory)

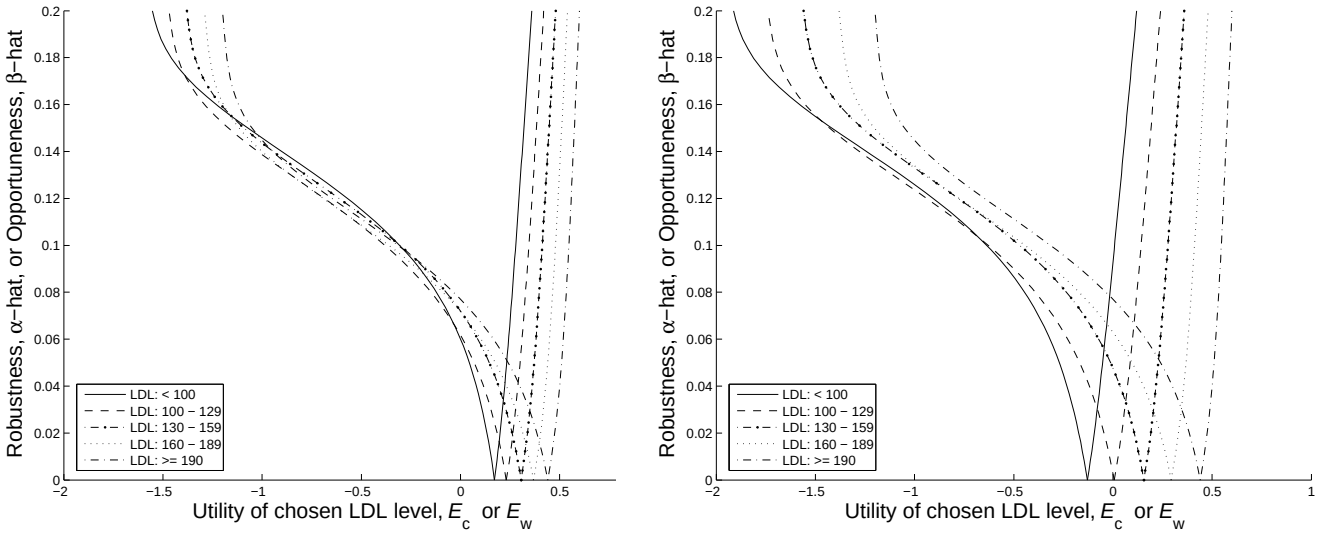

Figure 13: Case 4. Robustness and opportuneness vs. desired utility for case in section 5.4. Patient's estimated utilities are $\widetilde{u} = [-1,\ 0.5]$. ρ is the unit vector. $\ell_c = 5$, $f = 0.2$ (left), $f = 0.4$ (right)

our contention is that the combination of tools can prove useful in situations such as clinical decision making. The search for criteria of optimality of a decision tool is of prime importance. We suggest that methodological pluralism is a possible guide in this search.

7 Discussion: Info-Gap Theory and Other Methods

In the spirit of methodological pluralism we will briefly compare info-gap robust-satisficing and opportune-windfalling to four different decision methodologies: robust Bayesian analysis, min-max, probability bounds analysis, and coherent lower previsions. Our comparison will concentrate on three questions. The first question is *epistemic:* what information is needed in order to implement the methods. The second question is *behavioral:* what aspirations or requirements does the decision maker bring to the decision process, and how is the outcome of the decision evaluated? The third question is: how can the methods can be combined in a *hybrid* analysis, which is often a fruitful approach. We will not explicitly discuss axiomatic questions (what is the foundational structure of the method?), or questions of normative rationality (what philosophy of good decision making does the method entail?), though our discussion will sometimes intrude on these issues. It is hoped that this discussion will both clarify the nature of info-gap theory, and encourage researchers to study its relation to other theories. Klir writes that "It is not clear ... whether the *information-gap* conception of uncertainty ... can be formalized within [Generalized Information Theory]. This is an open research question." (italics in the original) [3, p.418].

Robust Bayesian Decisions. In a nutshell, robust Bayesian methods embed ordinary Bayesian methods in an analysis where some or all of the functions involved are uncertain. More specifically, a function—usually a prior probability or sometimes the loss function or the conditional probability—is replaced by a set of possible functions. The robustness analysis then evaluates the stability of the decision, or of some function of the decision such as a Bayesian risk. This can be done in many different ways [45, 46]; Jack Good has suggested that there are 46,656 types of Bayesian [47, chap.4]. One particularly Bayesian version of robust Bayesian analysis is to put a prior probability on the set of functions, and then perform an ordinary Bayesian analysis on this extended problem.

Epistemics. Choosing a set of possible realizations of the uncertain function is similar to the info-gap approach of defining an unbounded family of nested sets of possible realizations. The difference is knowing the size of this set, which is either explicit or implicit in the Bayesian approach. Explicit knowledge of the size is based on judgment: one chooses the size to be reasonable or plausible. Implicit knowledge is manifested in pragmatically varying the size to see how the robustness changes. This is not all that different in spirit from the info-gap robustness and opportuneness functions, which ask: how large can the uncertain set be without jeopardizing a critical

requirement (robustness), or how large must the set be in order to enable a windfall aspiration (opportuneness).

The strictly Bayesian version of robust Bayesian analysis—choose a prior probability distribution on the set of uncertain functions—is epistemically very different from the info-gap approach, in which uncertainty is represented by sets of functions without any measure functions at all. As Levi pointed out: "Strict Bayesians are legitimately challenged to tell us where they get their numbers." [48, p.387].

Behavior. Bayesians are very concerned with making good decisions, and thus pay careful attention to the choice of the loss or gain function. When a risk, or regret, or other loss function is used, it is minimized; when a benefit, or utility, or other gain function is used, it is maximized. This differs from both the robust-satisficing and the opportune-windfalling approaches.

Robust-satisficing is motivated by a behavioral orientation which says: a particular quality of outcome is essential, and any decision which achieves at least this critical level is acceptable even if it is less than the putative maximum. Satisficers do not attempt to optimize the outcome. Psychologically, 'more' can be 'less' at the end of the day [49]. If the critical level is less than the maximum, then there will usually be a multiplicity of actions which achieve the critical level. Furthermore, a satisficing action may be more likely than a purportedly optimizing action to achieve the critical outcome [5, section 11.4; 50]. The robust-satisficing approach seeks an action which satisfies the critical requirement with maximum immunity to uncertainty. What is maximized is robustness, not outcome.

In the opportune-windfalling approach the behavior is more ambitious, seeking to facilitate favorable, better-than-anticipated outcomes. The decision maker aspires to windfalls and chooses an action which enables (though cannot guarantee) a wonderful outcome at the lowest possible level of uncertainty. What is maximized is the opportuneness of the decision to exploit favorable uncertainty, and not the outcome.

Hybridization. The Bayesian uses judgment to choose a set of uncertain functions. Such judgments may be subject to uncertainty, implying that the set of uncertain functions may itself be uncertain. An info-gap model can represent uncertainty in this set. That is, the elements of the info-gap model are sets. The info-gap robustness analysis asks: how wrong can the prior choice of the uncertainty set be, and the outcome requirements are satisfied? The info-gap opportuneness analysis asks: how wrong must the initial set be in order for windfall to be possible?

Min-Max Decisions. The min-max approach identifies a set of possible contingencies or models or relevant functions and seeks the decision which minimizes the worst (maximal) loss on this set. (The max-min approach maximizes the lowest (minimal) gain when considering benefits rather than losses.) This concept is implicit in many robust Bayesian realizations, and many of our comments there apply here as well. *Epistemically,* the info-gap and min-max approaches are similar in representing uncertainty without measure functions, though the min-max approach requires the choice of a specific set. The *hybridization* of a min-max with an info-gap approach is often attractive, as discussed in connection with Bayesian methods. Indeed, Wald's work in the early 1940s on min-max considers sets of uncertain probability distributions [51].

Behavior. We will discuss two concepts: the observational equivalence of min-max with info-gap robust-satisficing, and the behavioral difference of these methods [52].

Observational equivalence: Suppose a robust-satisficing decision maker must choose between two options, D_1 and D_2, and requires an outcome no worse than L_c in fig. 14. This leads the robust-satisficer to choose decision D_1, which is more robust than D_2 at this requirement. An observer can describe this behavior by supposing the decision maker to be an min-maxer who believes that the horizon of uncertainty is α_1, because at this level of uncertainty the maximum potential loss from D_1 is less than from D_2. Conversely, a min-maxing decision maker who believes that α_1 is the true horizon of uncertainty would likewise choose D_1 over D_2. An observer could describe this by supposing the decision maker is a robust-satisficer whose requirement is L_c. In short, either strategy can be used to describe observed behavior by ascribing particular beliefs to the decision maker. In other words, the modelling of decision-behavior under uncertainty is under-determined in choosing between robust-satisficing and min-maxing.

Behavioral difference: Suppose a min-maxing decision maker believes that the horizon of uncertainty can be as large as α_2 in fig. 14, but no larger. The min-maxer will prefer D_2, whose loss can be as large as L_m, but less than the maximum potential loss of D_1. Suppose a robust-satisficing decision maker can accept a loss as large as L_c, but no larger. This robust-satisficing decision maker will prefer D_1 over D_2 since D_1 can tolerate greater uncertainty for achieving this requirement. The robust-satisficer will choose D_1 over D_2 even if the min-maxer has convinced the

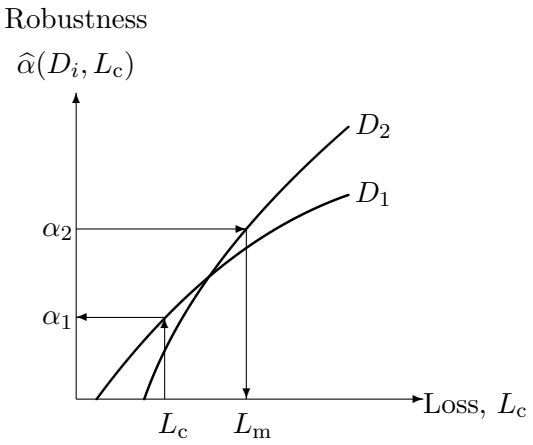

Figure 14: Crossing robustness curves, illustrating the observational equivalence and behavioral difference between min-maxing and robust-satisficing.

robust-satisficer that α_2 is the true horizon of uncertainty. In short, when the robustness curves for two decisions cross one another, a min-maxer and robust-satisficer may disagree about the decision, depending on their beliefs and requirements.

In conclusion, the *observational equivalence* between min-maxing and robust-satisficing means that modellers can use either strategy to describe observed behavior of decision makers. In contrast, the *behavioral difference* means that actual decision makers will not necessarily be indifferent between these strategies, and will choose a strategy according to their beliefs and aspirations.

The concepts of observational equivalence and behavioral difference have been noted before, in different terms. Walley writes [53, p.10]:

> Every [Dempster-Shafer] belief function can be represented as a lower envelope of a set of probability measures. This is merely a mathematical representation, however; it is misleading and unnecessary to regard a belief function as a lower bound for some unknown probability measure. In the same way, every coherent lower prevision can be represented as a lower envelope of a set of linear previsions, but this is no reason to regard the lower prevision as a model for partial information about an unknown linear prevision.

The observational equivalence of min-maxing and robust-satisficing asserts that either can be used as a mathematical representation of the other. The behavioral difference between these methods asserts that real decision makers with specific beliefs and requirements need not be indifferent between these methods.

Probability Bounds Analysis. A probability box, or P-box for short, is an interval of cumulative (probability) distribution functions (cdf's). In its simplest form a P-box is specified by a left-bounding and a right-bounding cdf. Uncertainty is analyzed by propagating the P-box through the equations which describe the system, resulting in probability bounds on outcomes. P-boxes are useful when the analyst is unsure about the precise cdf to use. Furthermore, convenient generic software is available for implementing P-box analyses [54].

The P-box method belongs to the family of non-measure-theoretic set-based representations of uncertainty, and thus is *epistemically* similar to info-gap theory. Since an info-gap model of uncertainty is an unbounded family of nested sets of functions, one can formulate an info-gap model for uncertainty in cdf's as a family of nested P-boxes [55]. An ordinary P-box analysis requires choosing a specific size for the P-box, while the info-gap model does not specify the size. Nonetheless it is clear that the methods can be combined in a *hybrid* analysis, as has been done in a study of portfolio investment [31]. The P-box method is *behaviorally* neutral or non-committal, unlike Bayesian analysis which usually entails some type of risk function, or info-gap analysis which supports robust-satisficing and opportune-windfalling. A P-box analysis can be integrated into either an outcome-optimizing analysis (e.g. maximizing expected utility or minimizing Bayesian risk) or an info-gap satisficing or windfalling analysis.

Coherent Lower Previsions. Much has been written on lower previsions; a sketchy mention of central ideas will have to suffice. A gamble is a utility function, $X(\omega)$, which depends on an uncertain state of the world ω. That is, utility $X(\omega)$ results from state ω. A lower prevision $\underline{P}(X)$

can be interpreted as the supremum buying price that an agent is willing to pay for the gamble X. An upper prevision, $\overline{P}(X)$ is the infinum selling price for the gamble, and is related to the lower prevision as $\overline{P}(X) = -\underline{P}(-X)$ [53]. One measure of the imprecision, incompleteness, or internal conflict of the agent's information is the difference between the upper and the lower previsions, $\overline{P}(X) - \underline{P}(X)$. The lower prevision is 'coherent' if it satisfies several properties, whose meaning is that (1) sure losses are not acceptable, (2) a gamble is acceptable if it is sure to be better than another acceptable gamble, and (3) a positive linear combination of acceptable gambles is also acceptable [56].

Epistemics. Any coherent lower prevision can be expressed as the lower envelope of a closed convex set of linear previsions (which are ordinary statistical expectations) [53, p.12]. This connects the lower prevision to a set of uncertain alternatives, thus establishing a similarity to info-gap models of uncertainty which are families of sets. However, an info-gap model of uncertainty is not a single set, but rather an unbounded family of nested sets of uncertain alternatives. While a closed and bounded set has a worst case, an info-gap model does not (except a vacuous case which bounds the entire universe of possibilities, if such a case exists at all). Thus the set uncertainty in info-gap theory is more extreme than the set uncertainty to which coherent lower previsions can be related.

In addition, the epistemic starting point for coherent lower previsions is (usually) a set of "judgments and expressions of uncertainty" [53, p.8] such as 'A is more likely than B' or 'If A then probably B', and so on. The process of 'natural extension' then constructs the coherent lower prevision from the uncertainty statements. This can be done by finding the set of linear previsions which are consistent with the uncertainty statements. In contrast, a typical starting point for an info-gap analysis includes specification of the unbounded family of sets which make up the info-gap model. (This suggests a possible hybridization which we will mention shortly.) In short, uncertainty is represented by real-valued functions in the theory of coherent lower previsions, while in info-gap theory this task is assumed by a set-valued function (the info-gap model of uncertainty).

Another epistemic distinction between the theory of lower previsions and info-gap theory has to do with imprecision. In the former theory imprecision is evaluated as the difference between the upper and lower previsions. This difference equals zero in the Bayesian case when probabilities are known precisely. The imprecision that is quantified by the lower and upper previsions represents "incompleteness or conflict in the available information" [53, p.10], which would have to be resolved before implementing a Bayesian analysis. An info-gap model of uncertainty represents incompleteness or conflict in the available information much less informatively, by defining a family of sets of possibilities. For instance, suppose we have evidence provided by experts. The evidence is perhaps in part conflicting, and other experts might give different evidence. The incompleteness and conflict of the evidence can be represented by an info-gap model whose elements are sets of possible evidence, of which the actual evidence in hand is one such set. In both theories the imprecision is propagated through to the decision, though the nature of that propagation, and the impact on the decision, can be quite different.

Behavior. The behavior of the decision maker is a central concern in the theory of lower and upper previsions, where the interpretation "in terms of buying and selling prices for gambles or in terms of their implications in other decision problems ... is sufficient to justify the axioms and calculus of the theory." [53, p.10] This behavioral interpretation is normative. It identifies the characteristics of a decision procedure which is rational in the sense of conforming to rules of consistency for achieving goals desired by the agent: a rational agent shouldn't accept a sure loss, etc. This approach to the behavioral formulation of a decision theory is in the tradition of Ramsey, von Neumann and Morgenstern, Savage, and other early thinkers for whom rationality is manifested in the concept of ordered preferences together with some form of logical consistency of preferences (not to imply that these thinkers agreed on all points).

Stated differently, the behavioral interpretation of coherent lower previsions is foundational. This differs from the behavioral interpretation of info-gap robust-satisficing and opportune-windfalling which put major emphasis on evaluating the confidence in, or feasibility of, required or aspired outcomes, and revising these requirements or aspirations as a result of analysis. Robust-satisficing and opportune-windfalling can certainly be related to a foundational logic of preferences. Likewise, a decision analysis based on coherent lower previsions can enable the decision maker to start over or revise judgments when outcomes or implications are unacceptable. However, one attraction of info-gap satisficing and windfalling for actual decision makers is that info-gap theory directly and immediately incorporates outcome-aspirations in an iterative analysis of uncertainty.

Hybridization. Info-gap theory hybridizes easily and naturally with many disparate decision

theories, and can be done in many different ways. We will suggest here two possible hybridizations with lower previsions.

An analysis with lower previsions typically starts by identifying judgments of uncertainty which the decision maker believes. The procedure of 'natural extension' enables the computation of new lower previsions based on these judgments of uncertainty. One can imagine situations in which the list of uncertainty judgments is itself uncertain. This may be due to "framing effects": the way verbal judgments are phrased can be influenced by the way in which information is presented. Or it may be due to subtle distinctions which escape the decision maker. For instance, the statement 'If A then probably B' is not the same as the statement 'Probably A, hence B'. However, the decision maker may not distinguish between them or consider them both. Or, uncertainty in the list of uncertainty judgments may result from uncertainty about the situation. For instance, in the Cuban missile crisis Kennedy's advisors may have been honestly uncertain if they should tell the president 'Russia probably *does* want war between Cuba and the US' or 'Russia probably *does not* want war between Cuba and the US'. (See [32] for an info-gap analysis of uncertain knowledge about the preferences of an adversary.) In short, the set $\mathcal{M}$ may be uncertain. We may not know what alternative sets the decision maker might induce if the information is presented or cogitated differently. One can formulate an info-gap model for uncertainty in the natural extension of the uncertainty judgments. This info-gap model is a family of nested sets whose elements are themselves sets of probability distributions. One then asks the info-gap robustness question: how wrong can the decision maker's actual specified set, $\mathcal{M}$, be, and any proposed inference or decision will lead to an acceptable outcome? The analogous windfalling question is also relevant.

The other hybridization of info-gap theory with coherent lower previsions is more direct. One can construct an info-gap model for uncertainty in the lower prevision function itself, $\underline{P}(X)$. The lower prevision inference is again embedded in an info-gap robustness or opportuneness analysis.

Acknowledgment

The authors gratefully acknowledge the support of the Abramson Center for the Future of Health. The authors are indebted to comments by Scott Ferson, Malka Gorfine, Matthias Troffaes and Miriam Zacksenhouse.

8 References

1. Hacking, Ian, 1975, *The Emergence of Probability: A Philosophical Study of Early Ideas About Probability, Induction and Statistical Inference,* Cambridge University Press.
2. Ekeland, Ivar, 1988, *Mathematics and the Unexpected,* University of Chicago Press.
3. Klir, George J., 2006, *Uncertainty and Information: Foundations of Generalized Information Theory,* Wiley Publishers.
4. John A. Wheeler, quoted in *Scientific American,* Dec. 1992, p20.
5. Ben-Haim Y., 2006, *Info-Gap Decision Theory: Decisions Under Severe Uncertainty,* 2nd edition. London: Academic Press.
6. Klir, George J., 2004, Generalized information theory: aims, results and open problems, *Reliability Engineering and System Safety,* 85: 21–38.
7. NGC. 2007, National Guideline Clearinghouse, http://www.guideline.gov/browse/guideline_index.aspx.
8. Mehta RH, Peterson ED, Califf RM. Performance measures have a major effect on cardiovascular outcomes: A review. *Am. J. Med.* 2007;120:398–402.
9. Frei CR, Restrepo MI, Mortensen EM, Burgess DS. Impact of guideline-concordant empiric antibiotic therapy in community-acquired pneumonia. *Am. J. Med.* 2006;119:865–871.
10. Cabana MD, Rand CS, Powe NR *et al.* Why don't physicians follow clinical practice guidelines? A framework for improvement. *J. Amer. Med. Assoc.,* 1999;282:1458–1465.
11. Hening WA. Current Guidelines and Standards of Practice for Restless Legs Syndrome. *Am. J. Med.* Jan. 2007, Vol. 120, Issue 1, Supplement 1, pp.S22–S27.

12. Greenfield S, Kravitz R, Duan N, Kaplan SH. Heterogeneity of treatment effects: Implications for guidelines, payment, and quality assessment. *Am. J. Med.* April 2007, Vol. 120, Issue 4, Supplement 1, pp.S3–S9.

13. Feinstein AR, Horwitz RI. Problems in the "evidence" of "evidence-based medicine". *Am. J. Med.* 1997;103:529–535.

14. Morimoto T, Fukui T, Lee TH, Matsui K. Application of U.S. guidelines in other countries: Aspirin for the primary prevention of cardiovascular events in Japan. *Am. J. Med.* 2004;117:459–468.

15. McLaughlin MJ, for the members of the HTE Policy Roundtable Panel. Healthcare policy implications of heterogeneity of treatment effects. *Am. J. Med.* 2007;120:S32–S35.

16. Lobach DF, Hammond WE. Computerized decision support based on a clinical practice guideline improves compliance with care standards. *Am. J. Med.* 1997;102:89–98.

17. Burgman M. Risks and Decisions for Conservation and Environmental Management. Cambridge: Cambridge University Press; 2005.

18. Regan HM, Y. Ben-Haim, B. Langford, W.G. Wilson, P. Lundberg, S.J. Andelman, M.A. Burgman, Robust decision making under severe uncertainty for conservation management. *Ecological Applications.* 2005;15:1471–1477.

19. McCarthy MA, Lindenmayer DB. Info-gap decision theory for assessing the management of catchments for timber production and urban water supply. *Environmental Management.* 2007;39:553–562.

20. Carmel Y, Ben-Haim Y. Info-gap robust-satisficing model of foraging behavior: Do foragers optimize or satisfice? *American Naturalist.* 2005;166:633–641.

21. Fox DR, Ben-Haim Y, Hayes KR, McCarthy M, Wintle B, Dunstan P. An info-gap approach to power and sample size calculations. *Environmetrics.* 2007;18:189–203.

22. Moffitt LJ, Stranlund JK, Field BC. Inspections to Avert Terrorism: Robustness Under Severe Uncertainty. *Journal of Homeland Security and Emergency Management.* 2005. http://www.bepress.com/jhsem/vol2/iss3/3

23. Kanno Y, Takewaki I. Robustness analysis of trusses with separable load and structural uncertainties. *International Journal of Solids and Structures.* 2006;43:2646–2669.

24. Kanno Y, Takewaki I. Sequential semidefinite program for maximum robustness design of structures under load uncertainty. *Journal of Optimization Theory and Applications.* 2006;130:265–287.

25. Pantelides CP, Ganzerli S. Design of trusses under uncertain loads using convex models, *ASCE J. Structural Engineering.* 1998;124:318–329.

26. Ganzerli S, Pantelides CP. Optimum structural design via convex model superposition. *Computers and Structures.* 2000;746:639–647.

27. Lindberg HE. Dynamic response and buckling failure measures for structures with bounded and random imperfections. *ASME Journal of Applied Mechanics.* 1991;58:1092–1094.

28. Pierce SG, Ben-Haim Y, Worden K, Manson G. Evaluation of neural network robust reliability using information-gap theory. *IEEE Transactions on Neural Networks.* 2006;17:1349–1361.

29. Ben-Haim Y, Laufer A. Robust reliability of projects with activity-duration uncertainty. *ASCE Journal of Construction Engineering and Management.* 1998;124:125–132.

30. Regev S, Shtub A, Ben-Haim Y. Managing project risks as knowledge gaps. *Project Management Journal.* 2006;37:17–25.

31. D. Berleant, L. Andrieu, J.-P. Argaud, F. Barjon, M.-P. Cheong, M. Dancre, G. Sheble and C.-C. Teoh, 2008, Portfolio management under epistemic uncertainty using stochastic dominance and information-gap theory, *International Journal of Approximate Reasoning,* vol.49, issue 1, pp.101–116.

32. Ben-Haim Y, Hipel KW. The graph model for conflict resolution with information-gap uncertainty in preferences. *Applied Mathematics and Computation.* 2002;126:319–340.

33. Wilson PWF, D'Agostino RB, Levy D, Belanger AM, Silbershatz H, Kannel, WB. Prediction of Coronary Heart Disease, *Circulation,* 12 May 1998; 97(18): 1837–1847.

34. Brindle PM, McConnachie A, Upton MN, Hart CL, Smith GD, Watt, GCM. The accuracy of the Framingham risk-score in different socioeconomic groups: A prospective study. *British Journal of General Practice.* 2005;55:838–845.

35. Kostis JB. A New Approach to Primary Prevention of Cardiovascular Disease. *Am. J. Med.* 2007;120:746–747.

36. Robson J. Information needed to decide about cardiovascular treatment in primary care. *British Medical Journal.* 1997;314:277–283.

37. Grundy SM, Cleeman JI, Merz CNB; Brewer HB, Jr, Clark LT, Hunninghake DB, Pasternak RC, Smith SC Jr, Neil J. Stone NJ, Implications of Recent Clinical Trials for the National Cholesterol Education Program Adult Treatment Panel III Guidelines, *Circulation,* 13 July 2004, 110: 227–239.

38. Gómez M, Bielza C, Fernàndez del Pozo JA, Ríos-Insua, S. A graphical decision-theoretic model for neonatal jaundice. *Medical Decision Making.* 2007;27:250–265.

39. Keeney, RL and Raiffa, H, 1993, *Decisions with Multiple Objectives: Preferences and Value Tradeoffs,* New York: Cambridge University Press.

40. Keeney, RL, 1992, *Value Focussed Thinking,* Boston: Harvard University Press.

41. Ben-Haim, Y, 1999, Set-models of information-gap uncertainty: Axioms and an inference scheme, *Journal of the Franklin Institute,* 336:1093–1117.

42. Chan PS, Nallamothu BK, Gurm HS, Hayward RA, Vijan S. Incremental benefit and cost-effectiveness of high-dose statin therapy in high-risk patients with coronary artery disease. *Circulation.* 2007;115:2398–2409.

43. Davie AP, McMurray JJV. Editorial. *European Heart Journal.* 1996;17:974–975.

44. Mark DB, Hlatky MA. Medical Economics and the Assessment of Value in Cardiovascular Medicine: Part II. *Circulation.* 2002;106:626–630.

45. Berger, JO, 1980, *Statistical Decision Theory: Foundations, Concepts, and Methods.* Springer Verlag, New York.

46. Keren, C, 2008, *Info-Gap Bayesian Classification,* MSc thesis, Technion—Israel Institute of Technology.

47. Good, IJ, 1983, *Good Thinking: The Foundations of Probability and Its Applications,* University of Minnesota Press, Minneapolis.

48. Levi, I, 1982, Ignorance, probability and rational choice, *Synthese,* 53: 387–417.

49. Schwartz, B, 2004, *Paradox of Choice: Why More Is Less,* Harper Perennial, New York.

50. Ben-Haim, Y, 2008, Info-gap forecasting and the advantage of sub-optimal models, to appear in *European Journal of Operational Research.*

51. Wald, A. 1945, Statistical decision functions which minimize the maximum risk, *The Annals of Mathematics,* 46(2): 265–280.

52. Ben-Haim, Y, Akram, QF, and Eitrheim, Ø, 2007, Monetary policy under uncertainty: Min-max vs robust-satisficing strategies, Norges Bank Working Papers, ANO 2007/6.

53. Walley, P, 1996, Measures of uncertainty in expert systems, *Artificial Intelligence,* vol.83, pp.l–58.

54. Ferson, S, 2002, *RAMAS Risk Calc 4.0 Software: Risk Assessment with Uncertain Numbers,* Boca Raton: Lewis Publishers.

55. Ferson, S and Tucker, WT, 2008, Probability boxes as info-gap models, Annual Conference of the North American Fuzzy Information Processing Society (NAFIPS) 2008, Article number 4531314.

56. Miranda, E, 2008, A survey of the theory of coherent lower previsions, *International Journal of Approximate Reasoning,* vol.48, pp.628–658.

A Evaluating the Robustness and Opportuneness Functions

Robustness. Let $\mu(\alpha,\ell)$ denote the inner minimum in the definition of the robustness function, eq.(8). The robustness is the greatest horizon of uncertainty, α, up to which $\mu(\alpha,\ell) \geq E_{\rm c}$. The uncertainty sets, $\mathcal{U}(\alpha,\widetilde{\gamma})$, $\mathcal{U}_{\rm u}(\alpha,\widetilde{u})$ and $\mathcal{U}_{\rm i}(\alpha,\widetilde{u}_{\rm i})$, become more inclusive as α increases. This means that $\mu(\alpha,\ell)$ (which is a minimum on the uncertainty set at horizon α) decreases monotonically as α increases. Hence the robustness is the greatest value of α at which $\mu(\alpha,\ell) = E_{\rm c}$. Finally, this implies that $\mu(\alpha,\ell)$ is the inverse of the robustness function: a plot of $\mu(\alpha,\ell)$ on the horizontal axis, versus α on the vertical axis, is identical to a plot of $E_{\rm c}$ horizontally versus $\widehat{\alpha}(\ell,E_{\rm c})$ vertically. This is the basis for evaluating the robustness function.

One can readily derive the following expression for $\mu(\alpha,\ell)$:

$$\mu(\alpha,\ell) = \left[1 - S^{\rm e}{-G + c^T\widetilde{\gamma} + \alpha c^T b}\right](1+\alpha)\widetilde{u}_{\rm e} + S^{\rm e}{-G + c^T\widetilde{\gamma} + \alpha c^T b}\, h(1-\alpha)\widetilde{u}_{\rm h} + (1+\alpha)|\ell-\ell_{\rm c}|u_0 \tag{15}$$

The dependence of the righthand side on the chosen LDL level, ℓ, arises through c. b is the vector whose elements are $b_i = |\widetilde{\gamma}_i|\rho_i$.

Opportuneness. Let $M(\alpha,\ell)$ denote the inner maximum in the definition of the opportuneness function in eq.(10). By an argument analogous to the one about $\mu(\alpha,\ell)$ we conclude that $M(\alpha,\ell)$ is the inverse of the opportuneness function. That is, a plot of $M(\alpha,\ell)$ on the horizontal axis versus α on the vertical axis, is identical to a plot of $E_{\rm w}$ horizontally versus $\widehat{\beta}(\ell,E_{\rm w})$ vertically. This is the basis for evaluating the opportuneness. One finds:

$$M(\alpha,\ell) = \left[1 - S^{\rm e}{-G + c^T\widetilde{\gamma} - \alpha c^T b}\right] h(1-\alpha)\widetilde{u}_{\rm e} + S^{\rm e}{-G + c^T\widetilde{\gamma} - \alpha c^T b}\,(1+\alpha)\widetilde{u}_{\rm h} + h(1-\alpha)|\ell-\ell_{\rm c}|u_0 \tag{16}$$