

ROBUST DETECTION OF EXOTIC INFECTIOUS DISEASES IN ANIMAL HERDS: A COMPARATIVE STUDY OF THREE DECISION METHODOLOGIES UNDER SEVERE UNCERTAINTY

MATTHIAS C. M. TROFFAES AND JOHN PAUL GOSLING

ABSTRACT. When animals are transported and pass through customs, some of them may have dangerous infectious diseases. Typically, due to the cost of testing, not all animals are tested: a reasonable selection must be made. How to test effectively whilst avoiding costly disease outbreaks? First, we extend a model proposed in the literature for the detection of invasive species to suit our purpose, and we discuss the main sources of model uncertainty, many of which are hard to quantify. Secondly, we explore and compare three decision methodologies on the problem at hand, namely, Bayesian statistics, info-gap theory and imprecise probability theory, all of which are designed to handle severe uncertainty. We show that, under rather general conditions, every info-gap solution is maximal with respect to a suitably chosen imprecise probability model, and that therefore, perhaps surprisingly, the set of maximal options can be inferred at least partly—and sometimes entirely—from an info-gap analysis.

1. INTRODUCTION

This paper concerns the inspection of imported herds of animals for signs of known or unknown major exotic infectious diseases. Imports and exports of animals represent a significant contribution to UK agriculture. Even though imports are subject to strict controls at the UK border under EU and national rules, there is a real risk of animal diseases being introduced. Fèvre et al. [8] review the problems associated with animal movement and the spread of disease.

We will build further on the work of Moffitt et al. [15], who study inspection protocols for shipping containers of invasive species, employing info-gap theory [1] to model the severely uncertain number of infested items. The aim of their study is to realistically take into account economical considerations (actual costs of testing, and of invasive species passing through customs), whilst also soundly handling the enormous uncertainty.

A key feature of their, and also our, problem is that exact probabilities of the constituent events are very hard to come by [14]. This motivates the use of robust uncertainty models and decision tools, such as info-gaps [1] (i.e. robust satisficing) as in the original study, but also Bayesian statistics [3] and imprecise probabilities [19], as we will do in this paper.

Our study, using each of these decision methodologies, leads us to surmise a connection between info-gap analysis and imprecise probability theory (Γ -minimax and maximality in particular). We prove that the perceived connection is no coincidence, and we establish a rigorous theoretical link between the two approaches.

The paper is organised as follows. Section 2 introduces the problem of animal inspection, defines the model, discusses various uncertainties involved, and derives an expression for the expected loss under a simple binomial model for infection. Section 3 solves the inspection problem, first by Bayesian analysis, then using an info-gap model, and finally using an imprecise probability (or, robust Bayesian) model with maximality. These results are discussed in Section 4, where

Key words and phrases. Exotic disease, lower prevision, info-gap, maximality, minimax, robustness, inspection, protocol.

we formally define an info-gap model based on a nested set of imprecise probability models, and establish the theoretical connections between info-gap, Γ -minimax, and maximality. Section 5 concludes the paper.

2. ANIMAL HERD TESTING

In this section, we extend a model, proposed by [15] for the detection of invasive species, to suit our purpose:

- we explicitly take specificity and sensitivity into account in order to allow for imperfect testing,
- we take into account an additional cost term for terminating the herd in case an infection is detected, and
- we model the occurrence of diseased animals in the herd as a binomial process, under a worst-case assumption of independence of infections between animals.

2.1. Model Description. Consider a herd of n animals, of which m are tested—the problem is to choose m optimally. The uncertain number of diseased animals in the herd is denoted by d . The test has sensitivity—the probability that a diseased animal tests positive—equal to p , and specificity—the probability that a healthy animal tests negative—equal to q .

Testing m animals costs $c(m)$ utiles. If d diseased animals pass inspection undetected, we incur a cost of $a(d)$ utiles. When at least one diseased animal is detected, then, typically, the whole herd is terminated, costing $t(n)$ utiles.

Following [15, p. 295, Sec. 3], in the numerical examples that follow, we take

$$(1) \quad c(m) = 1000 - 2000m + 1000m^2$$

$$(2) \quad a(d) = \begin{cases} 0 & \text{if } d = 0 \\ a & \text{if } d \geq 1 \end{cases} \quad (a = 10\,000\,000)$$

Moffitt et al. [15] consider n between 250 and 2 500, do not need to consider the cost of termination ($t(n) = 0$), and assume perfect testing ($p = q = 1$). For our problem, in numerical examples that follow, we take

$$(3) \quad n = 250$$

$$(4) \quad t(n) = 400n = 100\,000$$

$$(5) \quad p = 0.9999$$

$$(6) \quad q = 0.999$$

so we assume that a diseased animal tests positive with probability 0.9999, and a healthy animal tests negative with probability 0.999. For reference, if $q = 0.999$, then probability that all animals in a healthy herd of size $n = 250$ test negative is $q^n = 0.78$.

2.2. Model Uncertainties. Obviously, many of these values are rather uncertain. The only values we are pretty certain of are the number of animals n in the herd, the cost of testing $c(n)$, and the cost of termination $t(n)$.

Due to the necessity that the herd must have valid health documentation, we would expect that the number of infected animals d would be low. Additional inspection by veterinary officials is costly and depends on the inspecting official's ability to spot signs of infectious disease like pathological lesions and abnormal behaviour. Of course, the level of experience and competency will vary from official to official, but the testing procedure should be thorough enough for us to be confident of both a high sensitivity, p , and specificity, q . In addition to this, the government would prefer the most sensitive test possible (within budgetary constraints), even if specificity

was slightly compromised, because a rare false positive would be better for the prevention of disease entry than a rare false negative. Hence, we would expect $p > q$. Further discussion of this can be found in [23].

Of course, in general, having values for p and q as high as 0.9999 and 0.999 is unrealistic. For most tests, the developers have aimed at getting a high value for the specificity and sensitivity suffers. However, there are examples in animal disease testing where both the sensitivity and specificity are this high. For example, the virus antibody test for caprine arthritis-encephalitis claims sensitivity and specificity values of over 99.5% [12] and near perfect sensitivity and specificity have been estimated for the polymerase chain reaction test for parasites in fish [6].

Regarding the cost a of an infection passing through customs, some historical data is available. For example, instances of major disease outbreaks in the last couple of decades include BSE where public spending was over £5 billion, and the foot and mouth outbreak in 2001 which costed the UK government £2.6 billion [5]. These experiences show that there is great variation in the level of costs of exotic disease outbreaks. Due to the exceptional nature of the outbreaks, there is limited evidence on which to base cost assessments. Therefore, there is great uncertainty about what may happen in the future.

Outbreaks of any particular exotic disease are generally rare or may never have occurred at all. Also, diseases change as new strains develop; consequently, the possibility of new diseases arriving can change rapidly. For example, until a few years ago, bluetongue was considered extremely unlikely in some European countries, but now outbreaks are expected every couple of years.

In late 2009, an elicitation exercise was carried out with government experts to help quantify the average annual costs to the UK government of exotic infectious disease outbreaks and the uncertainty about those estimates [11]. In that exercise, it was clear that the costs are severely uncertain even when the disease was known (for example, foot and mouth is an exotic infectious disease). A major contributor to the uncertainty about the overall cost was the possibility of an outbreak of an unknown infectious disease, which could cost anywhere from £0.5 billion to £6 billion.

The scale and costs of an outbreak will depend on the length of time between the diseased animal entering circulation and the disease's presence being confirmed, and the speed and effectiveness of the government's response. The eventual costs are influenced by any public health implications and the effects of disease controls on other industries. The main elements of the costs due to control measures include: the disposal of and payments for culled animals; the tracing, testing and diagnosis of animals; the cleaning and disinfection of infected premises; and administrative costs in managing the outbreak. The size of these costs will vary according to the scale of the outbreak with key factors being the number of infected premises, the numbers of animals culled, and the duration of the outbreak. These types of factors are considered in greater detail in [5] and [9].

A serious study of how all uncertainties involved could be taken into account in the model would of course be extremely interesting, but is beyond the goal of this paper. Instead, in this initial study, following [15] and many others, for now we will focus on the main uncertainty, that is, the number of diseased animals d , and simply assume reasonable values for the remaining parameters. Such simplistic model helps to illustrate the methodological differences and to motivate the theory.

2.3. Expected Loss. First, we derive the expected loss, in case all parameters of the problem are perfectly known, including the number of diseased animals d . Clearly, conditional on d , the expected loss is:

$$(7) \quad L(m, d, p, q, c, a, t) = c(m) + t(n) \Pr(T|d) + a(d) \Pr(T^c|d)$$




FIGURE 1. Loss as a function of the number of diseased animals for $m = 10$ and $m = 20$.

where T denotes termination of the herd, that is, the event that at least one diseased animal is detected, and T^c denotes its complement, that is, the event that the herd passes inspection.

Let us deduce $\Pr(T^c|d)$. First, if the test group of size m is sampled randomly and without replacement, then the probability of exactly z diseased animals in the test group follows a hypergeometric distribution:

$$(8) \quad \Pr(z|d) = \frac{\binom{d}{z} \binom{n-d}{m-z}}{\binom{n}{m}}.$$

Next, we calculate the probability of non-termination given z diseased animals in the test group, that is $\Pr(T^c|d, z)$. If $d = 0$, then the probability of non-termination is the probability of all healthy animals in the sample testing negative, so $\Pr(T^c|0, z) = q^m$. If $d \geq 1$, then given z diseased animals in the sample, non-termination occurs when none of the z diseased animals tests positive and all of the $m - z$ healthy animals test negative. Hence, in all cases,

$$(9) \quad \Pr(T^c|d, z) = (1 - p)^z q^{m-z}.$$

By the law of total probability,

$$(10) \quad \begin{aligned} \Pr(T^c|d) &= \sum_{z=0}^d \Pr(T^c|d, z) \Pr(z|d) \\ &= \sum_{z=0}^d (1 - p)^z q^{m-z} \frac{\binom{d}{z} \binom{n-d}{m-z}}{\binom{n}{m}}. \end{aligned}$$

Now we have all the ingredients to calculate the total expected loss if we choose to test m out of n animals:

$$(11) \quad L(m, d, p, q, c, a, t) = c(m) + t(n) + (a(d) - t(n)) \Pr(T^c|d)$$

or, if $a'(n, d) = a(d) - t(n)$ denotes the termination adjusted cost of apocalypse,

$$(12) \quad = c(m) + t(n) + a'(n, d) \Pr(T^c|d)$$

where $\Pr(T^c|d)$ is given by Eq. (10). Figure 1 depicts the expected loss for a few typical cases.

2.4. A Binomial Model for Infection. Moffitt et al. [15] consider an info-gap model directly over the number of diseased animals d , which leads to a rather tricky optimisation problem. Instead, we will consider the (highly uncertain) probability r that an animal is infected, and derive the expected loss as a function of r . Although we do not explore this topic further in this paper, this also paves the way to modelling spatial dependencies between infections in the herd, leading to more optimal testing strategies.

So, assume that each animal has a probability r of being infected; for simplicity, for now, we assume that one animal being diseased does not affect another animal being diseased. Obviously, this will generally not be satisfied, and more realistically, we would expect a positive correlation, resulting in diseased animals being clustered together in the herd. Assuming independence essentially amounts to a worst case study: at the other extreme end, if one diseased animal would immediately infect the whole herd, then it would be sufficient to test only a single animal, as $d = 0$ and $d = n$ would be the only two possibilities.

Under the worst case assumption of independence, the probability of having d out of n animals infected is:

$$(13) \quad \Pr(d|r) = \binom{n}{d} r^d (1-r)^{n-d}$$

The expected loss is:

$$(14) \quad E(L(m, \cdot, p, q, c, a, t)|r) = \sum_{d=0}^n L(m, d, p, q, c, a, t) \Pr(d|r)$$

From now onwards, we will simply write $L(m|r)$ instead of $E(L(m, \cdot, p, q, c, a, t)|r)$ in order to simplify notation. Figure 2 depicts $L(m|r)$ as a function of m for a few typical situations.

3. DECISION ANALYSIS

In this section, we explore and compare three decision methodologies, designed for severe uncertainty, on the problem at hand. In particular,

- we perform a Bayesian analysis,
- we accommodate the info-gap approach suggested by [15] to our extended model,
- we investigate possible ways of constructing sets of probabilities (i.e. imprecise probability models) which are in some sense equivalent to the proposed info-gap model, and
- we compare the decisions that these various models lead to.

3.1. Bayesian Analysis: Minimising Expected Loss. In this approach, we model the uncertainty about d probabilistically and choose the m that minimises expected loss. Of course, this is not the only strategy open to us when we use a Bayesian formulation, but it is the most common.

Using the same model for d as in Eq. (13), we can write the probability mass function for d to be

$$(15) \quad \Pr(d) = \int_0^1 \Pr(d|r) p(r) dr,$$

where $p(r)$ is a probability density function chosen to characterise our beliefs about r . For computational convenience and because the density is sufficiently flexible, we choose to model $r \sim \text{Beta}(\alpha, \beta)$, where α and β are chosen so that the Beta density matches our beliefs. Using this density for r ,

$$(16) \quad \Pr(d|\alpha, \beta) = \binom{n}{d} \frac{\text{Be}(\alpha + d, \beta + n - d)}{\text{Be}(\alpha, \beta)},$$




FIGURE 2. Expected loss $L(m|r)$ as a function of the test group size m , for $r = 0.00010$, $r = 0.00025$, $r = 0.00050$, and $r = 0.00100$, from bottom to top.

t	s	σ	α	β	m^*	$E(L \alpha, \beta)/10^6$
0.0002	199.0	0.001	0.040	198.9	2	0.316
0.0004	398.8	0.001	0.160	398.7	3	0.738
0.0008	798.4	0.001	0.639	797.7	6	1.567
0.0016	1596.4	0.001	2.554	1593.9	10	3.002

TABLE 1. Choices for parameters of the Beta distribution over r . For information, we also state the standard deviation σ , and the α and β parameters of the canonical parametrization.

where $\text{Be}(\cdot, \cdot)$ denotes the usual Beta function. In this analysis, we consider various choices of parameters for the Beta distribution, as listed in Table 1. Following Walley [22], t denotes the expectation $\frac{\alpha}{\alpha+\beta}$ of the prior, and s denotes $\alpha + \beta$, so $\alpha = st$ and $\beta = s(1 - t)$. In our analyses, we have varied the expectation of r by varying t , and we have chosen s such that the standard deviation is 0.001 throughout. We have chosen to investigate this set of distributions because they cover a range of reasonable beliefs that we may have about r when we are dealing with rare diseases.

For each Beta distribution, we can calculate the expected loss for each choice of m using

$$(17) \quad E(L|\alpha, \beta) = \sum_{d=0}^n \Pr(d|\alpha, \beta) L(m, d, p, q, c, a, t).$$




FIGURE 3. Expected loss plotted against m for each choice of parameters. Lower curves correspond to lower values of t .




FIGURE 4. Probability of exceeding loss threshold L_c for $m = 10$ and $r \sim \text{Beta}(\alpha, \beta)$ with parameters corresponding to the worst case considered (bottom row of Table 1), as a function of the critical cost L_c .

Figure 3 is a plot of the expected loss against number of tests for the three chosen distributions. The minimum expected loss occurs for $m = m^*$, given in Table 1.

Also, by taking a probabilistic viewpoint, we can derive a distribution for the possible losses given a choice of m and $p(r)$. If we choose $m = 10$ and the worst case distribution for r (bottom row of Table 1 where there is a greater than 30% chance of there being diseased animals), we can find probabilities for the losses exceeding different values. For the thresholds in Figure 4, as soon as we exceed the cost of termination and testing, we are essentially calculating the probability of disease outbreak. The jumps in this plot correspond precisely to the cost of testing (81 000), and the cost of testing plus termination (181 000).

It is worth noting that these probabilities are particularly sensitive to the choices we make for the parameters of the beta distribution. However, it is straightforward to assess the influence

	$t = 0.0002$	$t = 0.0004$	$t = 0.0008$	$t = 0.0016$	$t = 0.0032$
$s = 200$	0.030	0.058	0.113	0.211	0.371
$s = 400$	0.036	0.070	0.135	0.249	0.428
$s = 800$	0.040	0.079	0.150	0.276	0.466
$s = 1600$	0.043	0.084	0.160	0.292	0.489
$s = 3200$	0.045	0.087	0.166	0.301	0.501

TABLE 2. Probability of exceeding loss threshold $L_c = 181\,000$ for $m = 10$ and $r \sim \text{Beta}(\alpha, \beta)$ for a wide range of parameter choices. Values are shown for $L_c = 182\,000$, as the probability is discontinuous at $L_c = 181\,000$ (see Figure 4).

the parameters by tabulating the probability of L_c exceeding the cost of testing and termination (181 000 when $m = 10$) as shown in Table 2.

3.2. Info-Gap Analysis: Maximising Robustness. Another approach to solve our decision problem, under severe uncertainty about the exact probability r of a single animal being infected, is to select that decision which meets a given performance criterion, L_c , under the largest possible range of r . Given that we have almost no information about r , except that it assumes a very small but otherwise unknown value, this simple model seems to suffice for our purpose. Obviously, one could define many other more refined info-gap models—and our choice of model is just one example among many. For a much more detailed account, see [1].

Specifically, for a given value of L_c , the largest possible range $[0, h]$ of r for which we meet our performance criterion is characterised by

$$(18) \quad \hat{h}(m, L_c) = \max_{h \geq 0} \left\{ h : \underbrace{\max_{\substack{r \in [0, h] \\ r \leq 1}} L(m|r)}_{M(m, h)} \leq L_c \right\}$$

The value $\hat{h}(m, L_c)$, as a function of L_c , is called the *robustness curve*: it tells us how uncertain about r we can be for our decision m still to meet a given level of performance L_c .

A quick Poisson approximation reveals that as long as $\exp(-nh)$ is sufficiently close to 1 (and this holds for sufficiently small values of nh) the inner maximum over $r \in [0, h]$ is achieved at $r = h$ (also see Figure 2: the cost increases as r increases), so

$$(19) \quad M(m, h) = L(m|h)$$

Obviously, $M(m, h)$ increases as the horizon of uncertainty h increases, whence $\hat{h}(m, L_c)$ as a function of L_c is simply the inverse of $M(m, h)$ as a function of h . In other words, plotting $M(m, h)$ as a function of h for different values of m effectively gives us the robustness curves. Figure 5 depicts them.

The choices of m which maximise robustness, for various values of the critical cost L_c , are tabulated in Table 3. For example, at an expected cost $L(m|r)$ of at most $L_c = 3\,000\,000$, we can safeguard against any probability of infection $r \in [0, 0.001\,479]$, by testing 10 animals in the herd. For comparison, the Bayesian expected cost $E(L|\alpha, \beta)$ for $t = 0.0016$ is almost exactly equal to 3 000 000 (last row of Table 1) when $m = 10$, which is in agreement with the info-gap analysis. Of course, it is to be noted that the actual inputs into each decision model are different; nevertheless, both are assuming extreme uncertainty of a similar order if we interpret the Bayesian analysis as a worst case analysis—which we can easily do here due to the simplicity of the model. Thus the agreement is not all that surprising.




FIGURE 5. Robustness curves $\hat{h}(m, L_c)$ as a function L_c for test group sizes $m = 1$ (solid), $m = 15$ (dashed), and $m = 30$ (dotted).

$L_c/10^6$	m^*	$10^3 \hat{h}(m^*, L_c)$
0.5	2	0.207
1.0	4	0.426
1.5	5	0.661
2.0	6	0.912
2.5	8	1.184
3.0	10	1.479
3.5	11	1.803
4.0	13	2.163

TABLE 3. Info-gap choice of m , and corresponding horizon of uncertainty, for various values of the critical cost L_c .

3.3. Imprecise Probability Analysis: Maximality over a Partial Ordering. There are several ways one might go about constructing an imprecise probability model for our problem. As we have just seen, the info-gap approach hinges on the idea of satisficing: we may start out with a level of minimum performance that we hope to achieve, and the analysis tells us how much uncertainty we can tolerate, at this price. One might also interpret it conversely: for a given level of uncertainty, the analysis tells us how much we might potentially pay, if it comes to the worst.

Typical decision models for imprecise probabilities studied in the literature do not relate to satisficing, yet, they do incorporate an idea similar to the info-gap horizon of uncertainty: the imprecision of our model. Concretely, consider the set \mathcal{M}_h of all probability densities over r that are zero outside $[0, h]$.¹ We say that a choice m *dominates* a choice m' , and we write $m \succ m'$ whenever the expected loss under m is strictly less than the expected loss under m' over all densities p in \mathcal{M}_h , that is, whenever

$$(20) \quad \int_0^\infty L(m|r)p(r)dr + \epsilon \leq \int_0^\infty L(m'|r)p(r)dr$$

¹The adventurous reader may take all finitely additive probability measures μ on $[0, +\infty]$ with $\mu([0, h]) = 1$. We do without this complication: because all functions involved are continuous, those additional measures make no difference.

m	$10^3 h$							
	0.207	0.426	0.661	0.912	1.184	1.479	1.803	2.163
0	−0.9	−0.9	−0.9	−0.9	−0.9	−0.9	−0.9	−0.9
1	1.1	1.1	1.1	1.1	1.1	1.1	1.1	1.1
2	1.4	3.1	3.1	3.1	3.1	3.1	3.1	3.1
3	−0.6	2.1	4.9	5.1	5.1	5.1	5.1	5.1
4	−3.1	0.1	2.9	5.9	7.1	7.1	7.1	7.1
5	−7.7	−1.9	0.9	3.9	7.0	9.1	9.1	9.1
6	−14.3	−5.8	−1.1	1.8	5.0	8.4	11.1	11.1
7	−22.9	−11.8	−4.3	−0.2	2.9	6.3	9.9	13.1
8	−33.4	−19.7	−9.5	−2.4	0.9	4.2	7.9	11.8
9	−46.0	−29.7	−16.6	−6.7	−1.1	2.2	5.8	9.7
10	−60.6	−41.7	−25.9	−13.0	−4.3	0.1	3.7	7.6
11	−77.2	−55.6	−37.1	−21.3	−9.5	−1.9	1.7	5.6
12	−95.8	−71.6	−50.3	−31.6	−16.8	−5.9	−0.4	3.5
13	−116.4	−89.6	−65.6	−44.0	−26.1	−11.9	−2.9	1.4
14	−139.1	−109.7	−82.9	−58.4	−37.4	−20.0	−7.4	−0.7
15	−163.7	−131.7	−102.2	−74.8	−50.8	−30.1	−14.1	−3.5

TABLE 4. Result of Eq. (22) (divided by a factor 10^3 for everything to fit in the table). A positive value means that the corresponding choice of m is optimal for the given horizon of uncertainty h .

for all probability densities p in \mathcal{M}_h and some $\epsilon > 0$.

Because $L(m|r)$ is continuous in r for every m , this happens if and only if

$$(21) \quad \min_{r \in [0, h]} [L(m'|r) - L(m|r)] > 0$$

Note that the $\min_{r \in [0, h]}$ operator can be thought of as a lower expectation operator, or *lower prevision* \underline{P}_h —we will come back to this in Section 4.

One can easily prove that \succsim is a partial order, whence, a sensible way to choose m is to pick one which is not dominated by any other option, or in other words, which is *maximal*. The idea of choosing undominated options goes back at least to Condorcet [4, pp. lvj–lxix, 4.^e Exemple]; also see [17, p. 55, Eq. (1)], [21, Sections 3.7–3.9], and [19] for further discussion. Maximality has also been used in robust Bayesian models [13, §10.4], under slightly different terminology.

Given our partial order, one can easily show that an option m is maximal if and only if

$$(22) \quad \min_{m' \in \{0, 1, \dots, n\}} \max_{r \in [0, h]} [L(m'|r) - L(m|r)] \geq 0$$

The inner maximum is almost always achieved at either $r = 0$ or $r = h$, simplifying practical calculations substantially. Table 4 depicts these values for all choices of m , and varying values of h . For ease of comparison with the info-gap solution, we have chosen the same values of h as those listed in Table 3.

4. DISCUSSION

It is already well known that robust Bayesian models and imprecise probability models are, for the most part, mathematically equivalent [21, §5.9, pp. 253–258]. Therefore, in the following, we will focus on info-gap and maximality—the latter being used for both imprecise and robust Bayesian models.

Interestingly, in our example, info-gap and maximality give essentially the same result, with maximality refining the picture slightly: for a given horizon of uncertainty h , the maximal solutions are $\{1, \dots, m^*\}$, where m^* is the info-gap solution. The most notable result is that all info-gap solutions are maximal. Is this a coincidence? Formulating info-gap theory in terms of lower previsions, we show that this holds for arbitrary info-gap models and arbitrary lower previsions, subject to the mild and usually satisfied conditions of Theorem 1 (for Γ -maximin) and Theorem 2 (for maximality).

4.1. Info-Gaps for Imprecise Probabilities. Let $\omega \in \Omega$ be an uncertain parameter of interest— Ω can be an arbitrary set. We must select a decision d from a finite set D . The loss function $L(d, \omega)$ represents the loss (in utiles) if we choose d and ω obtains.

Info-gap theory starts out with a family of nested sets U_h of Ω , where h is a non-negative parameter called the *horizon of uncertainty* and $U_h \subseteq U_{h'}$ whenever $h \leq h'$. In our example, U_h was simply $[0, h]$. Following that example, we saw that a very natural way to model these nested sets U_h in terms of sets of probabilities goes by way of a *vacuous model* \mathcal{M}_h , that is, the set of all probability densities that are zero outside U_h .

If we denote the upper expectation induced by \mathcal{M}_h by \bar{P}_h (i.e. the pointwise lowest upper bound for the set of expectation operators associated to \mathcal{M}_h), then, formally, we define the info-gap solution $D^*(L_c) \subseteq D$ at satisficing level L_c as:

$$(23) \quad \hat{h}(d, L_c) := \max \{h : \bar{P}_h(L(d, \cdot)) \leq L_c\}$$

$$(24) \quad D^*(L_c) := \arg \max_{d \in D} \hat{h}(d, L_c)$$

Note that $D^*(L_c)$ will usually be a singleton (or, the empty set).

Also note that the first equation may not have a solution: this happens when $\bar{P}_0(L(d, \cdot)) > L_c$, that is, when d is infeasible even if we are as certain as can be ($h = 0$).

Now, from the point of view of imprecise probability, there is no compelling reason to restrict ourselves to vacuous models. In fact, we can allow \mathcal{M}_h to be any set of probability densities on Ω —we already picked a more general set in our Bayesian analysis in Section 3.1—under one restriction: a close inspection of the theory reveals that a crucial property that the info-gap model relies on is that the worst case cost, $\bar{P}_h(L(d, \cdot))$ is non-decreasing as the horizon of uncertainty h increases. Whence, we logically impose that $\mathcal{M}_h \subseteq \mathcal{M}_{h'}$ whenever $h < h'$.

So, instead of starting out from a family of nested subsets U_h of Ω , we start out from a family of nested sets \mathcal{M}_h of probability densities on Ω . For instance, in our example, the uncertainty was over the values of r , so \mathcal{M}_h would be some arbitrary set of probability distributions for r . In the Bayesian analysis, we restricted \mathcal{M}_h to the Beta family, for computational convenience.

One can of course interpret this again as an info-gap model, where the uncertain parameter is now the probability density over Ω —also see [2, pp. 1062–1063] for an informal discussion of this approach. The imprecise Dirichlet model [22] is an example of such family (with $h = 1/s$). For another example, see [7] for a discussion of nested sets of p-boxes and the resulting info-gap analysis.

4.2. Main Result. The next result links the info-gap solution to the so-called Γ -minimax² solution. See [2, p. 1061, Fig. 14] for an informal discussion of a very similar equivalence between info-gap and minimax. Interestingly, [18, p. 4, Table 1] constructs a maximin model that is fully equivalent to an info-gap model. The result below is quite different, as we change neither variables nor loss function.

² Γ -minimax minimises the upper expectation of the loss [20, 10].

Effectively, we show that, if certain fairly mild conditions are satisfied, the info-gap solution coincides with the Γ -minimax solution—parametrized over the horizon of uncertainty h —of exactly the same problem. In fact, such approach has already been used as a technique to solve info-gap problems (see for instance [16, pp. 1690–1691]). Below, we identify sufficient conditions for this to work, and provide two counterexamples in cases where one of these conditions fails.

For convenience, we denote the Γ -minimax loss at horizon $h \in \mathbb{R}^+$ by $L^*(h)$:³

$$(25) \quad L^*(h) := \min_{d \in D} \bar{P}_h(L(d, \cdot)).$$

Note that L^* is non-decreasing as a function of h , because each $\bar{P}_h(L(d, \cdot))$ is. We will be interested in the *right derivative* of L^* at a point h :

$$(26) \quad \partial_+ L^*(h) := \lim_{\substack{h' \rightarrow h \\ h' > h}} \frac{L^*(h') - L^*(h)}{h' - h}$$

Because L^* is non-decreasing, $\partial_+ L^*(h)$ is non-negative and well defined (possibly $+\infty$) for every $h \in \mathbb{R}^+$.

Theorem 1. *Let $h \in \mathbb{R}^+$ and $L_c \in \mathbb{R}$. The info-gap solution $D^*(L_c)$ coincides with Γ -minimax solution with respect to \bar{P}_h , that is,*

$$(27) \quad D^*(L_c) = \arg \min_{d \in D} \bar{P}_h(L(d, \cdot)),$$

whenever the following conditions are satisfied:

$$(28) \quad \partial_+ L^*(h) > 0, \text{ and}$$

$$(29) \quad L^*(h) = L_c$$

Proof. By definition, $d^* \in D^*(L_c)$ whenever, for all $d \in D$,

$$(30) \quad \hat{h}(d^*, L_c) \geq \hat{h}(d, L_c)$$

By definition of $\hat{h}(d, L_c)$, this is equivalent to saying that

$$(31) \quad \{h' : \bar{P}_{h'}(L(d^*, \cdot)) \leq L_c\} \supseteq \cup_{d \in D} \{h' : \bar{P}_{h'}(L(d, \cdot)) \leq L_c\}$$

Rewriting the above expression, we have, equivalently,

$$(32) \quad \{h' : \bar{P}_{h'}(L(d^*, \cdot)) \leq L_c\} \supseteq \left\{ h' : \min_{d \in D} \bar{P}_{h'}(L(d, \cdot)) \leq L_c \right\}$$

or, by Eq. (25),

$$(33) \quad \{h' : \bar{P}_{h'}(L(d^*, \cdot)) \leq L_c\} \supseteq \{h' : L^*(h') \leq L_c\}$$

By Eq. (28) and the non-decreasingness of L^* , it is easily seen that $L^*(h') > L^*(h)$ for all $h' > h$. Moreover, by Eq. (29), $L_c = L^*(h)$. Concluding, the set on the right hand side is a fancy way of writing $[0, h]$. Therefore, the above is equivalent to

$$(34) \quad \bar{P}_h(L(d^*, \cdot)) \leq L_c$$

Once more by Eqs. (29) and (25), this is equivalent to saying that d^* is a Γ -minimax solution with respect to \bar{P}_h . \square

³We set $\mathbb{R}^+ := \{x \in \mathbb{R} : x \geq 0\}$.

Interestingly, for given L_c such that

$$(35) \quad \min_{d \in D} \bar{P}_0(L(d, \cdot)) \leq L_c \leq \min_{d \in D} \bar{P}_\infty(L(d, \cdot))$$

it holds that Eq. (29) has a unique solution for $h \geq 0$ whenever L^* is strictly increasing and continuous in h . This means that we are effectively free to choose L_c under the additional assumption of continuity.

To see why we are not free to choose L_c when continuity is not satisfied—and violate Eq. (29)—imagine for instance that:

$$(36) \quad \bar{P}_h(L(d_1, \cdot)) = \begin{cases} h & \text{if } h \leq 1 \\ 3 + h & \text{if } h > 1 \end{cases} \quad \bar{P}_h(L(d_2, \cdot)) = \begin{cases} 1 + h & \text{if } h \leq 1 \\ 4 + h & \text{if } h > 1 \end{cases}$$

Then, for $L_c = 3$, we have that $D^*(3) = \{d_1, d_2\}$ because $\hat{h}(d, 3) = 1$ for both d_1 and d_2 , yet obviously d_1 is Γ -minimax (it could even be uniformly dominated by d_2). Effectively, this is simply a technical limitation of the info-gap model, as any reasonable person would probably agree with the Γ -minimax solution.

What happens if Eq. (28) is violated? Imagine for instance that:

$$(37) \quad \bar{P}_h(L(d_1, \cdot)) = \begin{cases} 0 & \text{if } h \leq 1 \\ h - 1 & \text{if } h > 1 \end{cases} \quad \bar{P}_h(L(d_2, \cdot)) = \begin{cases} 0 & \text{if } h \leq 2 \\ h - 2 & \text{if } h > 2 \end{cases}$$

With this choice, L^* is continuous, but not strictly increasing. For $h = 1$, we have that $\arg \min_{d \in D} \bar{P}_h(L(d, \cdot)) = \{d_1, d_2\}$ because $\bar{P}_h(L(d_1, \cdot)) = \bar{P}_h(L(d_2, \cdot)) = 0$ for $h = 1$. Yet, for $L_c = L^*(1) = 0$ we have that $\hat{h}(d_1, L_c) = 1 < \hat{h}(d_2, L_c) = 2$, so only d_2 is optimal according to the info-gap criterion. This example uncovers a technical limitation of the Γ -maximin model, as for this case, any reasonable person would probably agree with the more robust info-gap solution.

Now, it is well known that every Γ -minimax solution is also maximal (see for instance [19]), whence, we conclude:

Theorem 2. Suppose $\partial_+ L^*(h') > 0$ for all $h' \in [0, h]$. Then, for all $h' \in [0, h]$, every info-gap decision $d^* \in D^*(L^*(h'))$ is maximal with respect to \bar{P}_h :

$$(38) \quad \bigcup_{h' \in [0, h]} D^*(L^*(h')) \subseteq \{d \in D : (\forall d' \in D) (\bar{P}_h(L(d', \cdot) - L(d, \cdot)) \geq 0)\}$$

Proof. Use the preceding theorem, and note that every Γ -minimax with respect to $\underline{P}_{h'}$ is maximal with respect to \underline{P}_h , provided that $h' \in [0, h]$. \square

Again, if in addition L^* is continuous on $[0, h]$, then the range for $L^*(h')$ in the above theorem is simply an interval:

$$(39) \quad \{L^*(h') : h' \in [0, h]\} = \left[\min_{d \in D} \bar{P}_0(L(d, \cdot)), \min_{d \in D} \bar{P}_h(L(d, \cdot)) \right].$$

Summarising, Theorem 1 provides sufficient conditions for the info-gap solution, for fixed values of L_c and h , to be equal to the Γ -minimax solution: proponents of either approach are ‘observationally equivalent’ [2, Sec. 7].

Theorem 2 shows that a full fledged info-gap analysis, varying the horizon of uncertainty along an interval, yields an elegant approach to capture maximal solutions. In our example, we actually find *all* maximal options—in general this may not be the case. Still, it shows that an info-gap analysis can be of value even if maximality is the final goal:

- an info-gap analysis might give a rough idea of the size of the maximal set (in particular, it provides a lower bound for it),

- the analysis can be an appealing way to represent the maximal solution graphically (as in Figure 5), and
- as robustness curves show the trade-off between uncertainty and cost, they are also obviously useful in the process of elicitation.

5. CONCLUSION

We constructed a simple model for inspecting animal herds for dangerous exotic infections, building further on the work of Moffitt et al. [15]. We solved the problem using three popular decision methodologies suited for dealing with severe uncertainty: Bayesian analysis, info-gap analysis, and imprecise probability theory (maximality and Γ -minimax). We found that, in this example, the solutions of the info-gap and imprecise models essentially coincide, although the way they arrive at it is very different.

We explored the theoretical link between info-gap theory, Γ -minimax, and maximality. We established that, under rather general conditions, every info-gap solution is maximal. Therefore, the set of maximal options can be inferred at least partly, and sometimes wholly, from an info-gap analysis. Consequently, robustness curves also make sense in an imprecise probability (or, robust Bayesian) context, for exploring maximal options, and for elicitation, when studying the trade-off between uncertainty and cost that is often of interest to decision makers.

ACKNOWLEDGEMENTS

The authors thank Kirsty Hinchliff and Ben Powell, who have been involved with an embryonic draft of this paper. The paper has also benefited greatly from discussions with Yakov Ben-Haim, Frank Coolen, and Moshe Sniedovich, to who we extend our sincerest thanks.

REFERENCES

- [1] Yakov Ben-Haim, *Information gap decision theory: Decisions under severe uncertainty*, Academic Press, 2001.
- [2] Yakov Ben-Haim, Clifford C. Dacso, Jonathon Carrasco, and Nithin Rajan, *Heterogeneous uncertainties in cholesterol management*, International Journal of Approximate Reasoning **50** (2009), 1046–1065, doi:10.1016/j.ijar.2009.04.002.
- [3] James O. Berger, *Statistical decision theory and Bayesian analysis*, second ed., Springer, 1985.
- [4] Marquis de Condorcet, *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*, L'Imprimerie Royale, Paris, 1785.
- [5] Defra, *Impact assessment of an independent body for animal health in England*, 2009, <http://www.official-documents.gov.uk/document/cm77/7784/7784.pdf>.
- [6] Claes Enøe, Marios P. Georgiadis, and Wesley O. Johnson, *Estimation of sensitivity and specificity of diagnostic tests and disease prevalence when the true disease state is unknown*, Preventive Veterinary Medicine **45** (2000), no. 1–2, 61–81.
- [7] S. Ferson and W. T. Tucker, *Probability boxes as info-gap models*, Annual Conference of the North American Fuzzy Information Processing Society (NAFIPS 2008), 2008, Article number 4531314.
- [8] Eric M. Fèvre, Barend M. de C. Bronsvort, Katie A. Hamilton, and Sarah Cleaveland, *Animal movements and the spread of infectious diseases*, Trends in Microbiology **14** (2006), no. 3, 125–131, doi:10.1016/j.tim.2006.01.004.
- [9] M. G. Garner and M. B. Lack, *Modelling the potential impact of exotic diseases on regional Australia*, Australian Veterinary Journal **72** (1995), no. 3, 81–87.
- [10] Itzhak Gilboa and David Schmeidler, *Maxmin expected utility with non-unique prior*, Journal of Mathematical Economics **18** (1989), no. 2, 141–153.
- [11] J. P. Gosling, A. Hart, D. Mouat, M. Sabirovic, S. Scanlan, and A. Simmons, *Quantifying experts' uncertainty about the future cost of exotic diseases*, Tech. report, The Food and Environment Research Agency, 2011, doi:10.1111/j.1539-6924.2011.01704.x.
- [12] Lynn M. Herrmann, William P. Cheevers, Travis C. McGuire, D. Scott Adams, Melinda M. Hutton, William G. Gavin, and Donald P. Knowles, *Competitive-inhibition enzyme-linked immunosorbent assay for*

- detection of serum antibodies to caprine arthritis-encephalitis virus: Diagnostic tool for successful eradication*, Clin. Vaccine Immunol. **10** (2003), no. 2, 267–271.
- [13] David Ríos Insua, Fabrizio Ruggeri, and Jacinto Martín, *Bayesian sensitivity analysis*, Sensitivity Analysis (A. Saltelli, K. Chan, and E. M. Scott, eds.), Wiley, 2000, pp. 225–244.
 - [14] L. Joe Moffitt and Craig D. Osteen, *Prioritizing invasive species threats under uncertainty*, Agricultural and Resource Economics Review **35** (2006), no. 1, 41–51.
 - [15] L. Joe Moffitt, John K. Stranlund, and Craig D. Osteen, *Robust detection protocols for uncertain introductions of invasive species*, Journal of Environmental Management **89** (2008), 293–299.
 - [16] Atte Moilanen, Brendan A. Wintle, Jane Elith, and Mark Burgman, *Uncertainty analysis for regional-scale reserve selection*, Conservation Biology **20** (2006), no. 6, 1688–1697.
 - [17] Amartya Sen, *Social choice theory: A re-examination*, Econometrica **45** (1977), no. 1, 53–89.
 - [18] Moshe Sniedovich, *Eureka! Info-gap is worst case (maximin) in disguise!*, http://maximin.moshe-online.com/proof_a.pdf, December 2006, Working Paper No. MS-2-06.
 - [19] Matthias C. M. Troffaes, *Decision making under uncertainty using imprecise probabilities*, International Journal of Approximate Reasoning **45** (2007), no. 1, 17–29, doi:10.1016/j.ijar.2006.06.001.
 - [20] Abraham Wald, *Statistical decision functions which minimize the maximum risk*, The Annals of Mathematics **46** (1945), no. 2, 265–280, doi:10.2307/1969022.
 - [21] Peter Walley, *Statistical reasoning with imprecise probabilities*, Chapman and Hall, London, 1991.
 - [22] ———, *Inferences from multinomial data: Learning about a bag of marbles*, Journal of the Royal Statistical Society, Series B **58** (1996), no. 1, 3–34, URL: <http://www.jstor.org/stable/2346164>.
 - [23] D. H. Zeman, *The “best” diagnostic test*, Swine Health and Production **5** (1997), no. 4, 159–160.

DURHAM UNIVERSITY, DEPT. OF MATHEMATICAL SCIENCES, SCIENCE LABORATORIES, SOUTH ROAD, DURHAM
DH1 3LE, UNITED KINGDOM

E-mail address: matthias.troffaes@gmail.com

UNIVERSITY OF LEEDS, SCHOOL OF MATHEMATICS, LEEDS, LS2 9JT, UNITED KINGDOM

E-mail address: j.p.gosling@leeds.ac.uk