ELSEVIER

Contents lists available at ScienceDirect

Reliability Engineering and System Safety

journal homepage: www.elsevier.com/locate/ress



Information avoidance and overvaluation under epistemic constraints: Principles and implications for regulatory policies



Matteo Pozzi^{a,*}, Carl Malings^b, Andreea Minca^c

- ^a Carnegie Mellon University, Pittsburgh, PA, United States
- b OSU-Efluve, Laboratoire Interuniversitaire des Systèmes Atmosphériques (LISA), Centre National de la Recherche Scientifique (CNRS) UMR 7583, Université Paris-Est Créteil, France
- ^c Cornell University, Ithaca, NY, United States

ARTICLE INFO

Keywords: Information avoidance Pre-posterior bayesian analysis Reliability-based constraints Value of information

ABSTRACT

The Value of Information (VoI) assesses the impact of data in a decision process. A risk-neutral agent, quantifying the VoI in monetary terms, prefers to collect data only if their VoI surpasses the cost to collect them. For an agent acting without external constraints, data have non-negative VoI (as free "information cannot hurt") and those with an almost-negligible potential effect on the agent's belief have an almost-negligible VoI. However, these intuitive properties do not hold true for an agent acting under external constraints related to epistemic quantities, such as those posed by some regulations. For example, a manager forced to repair an asset when its probability of failure is too high can prefer to avoid collecting free information about the actual condition of the asset, and even to pay in order to avoid this, or she can assign a high VoI to almost-irrelevant data. Hence, by enforcing epistemic constraints in the regulations, the policy-maker can induce a range of counter-intuitive, but rational, behaviors, from information avoidance to over-evaluation of barely relevant information, in the agents obeying the regulations.

This paper illustrates how the structural properties of VoI change depending on such external epistemic constraints, and discusses how incentives and penalties can alleviate these induced attitudes toward information.

1. Introduction

Information collected by sensors and inspectors can significantly reduce the uncertainty in decision-making in many fields of engineering. In the management of urban systems, for example, pervasive integration of sensing technologies can be the base for a novel quantitative urban science that, in turn, would allow for a better control of these systems. However, such integration would follow from the decisions of managers and stakeholders of the urban assets (whom we will hereafter refer to as "agents"), who act under societal regulations which are often developed neglecting their effects on information collection.

Design, operation and maintenance of these assets can be formulated as a decision making process, under uncertainty on demands, capacities and long-term evolution. Agents take these "exploitative" decisions (e.g. repairing or replacing some assets) with the aim of optimizing their own utilities, or minimizing their own losses. As the consequences of these actions can potentially affect safety and economic prosperity of communities at a broader level, the society usually imposes regulations and public policies to affect or even control them.

Specifically, as agents may be prone to accept risks higher than the society can tolerate, possibly because they do not include all societal costs in their analysis, society can impose constraints on the available actions, depending on the circumstances. For example, a building code can prevent a structure from being open to the public when the probability of its failure is too high, despite the owner's will to do so. Through these constraints, society is able to indirectly implement the decisions that it considers optimal, balancing costs for construction, maintenance, operation and renovation with risks related to failures and malfunctioning.

If agents are free to take explorative actions for supporting their exploitative strategy, e.g. by collecting information using sensors and inspectors, they can base their behavior on Value of Information (VoI), an utility-based metric to assesses the impact of data in a decision process introduced by the seminal work of Raiffa and Schlaifer [16]. A risk-neutral agent, quantifying the VoI in monetary terms, prefers to collect data only if their VoI surpasses the cost to collect them. When an agent acts without external constraints, data have non-negative VoI (as free "information cannot hurt" [5]) and those with an almost-negligible

E-mail address: mpozzi@cmu.edu (M. Pozzi).

^{*} Corresponding author.

potential effect on the agent's belief have an almost-negligible VoI. However, these intuitive properties do not hold true for an agent acting under external constraints related to epistemic quantities, such as those posed by some regulations. For example, a manager forced to repair an asset when its probability of failure is too high will prefer to avoid collecting free information about the actual condition of the asset, and even to pay in order to avoid this, or she can assign a high VoI to almost-irrelevant data, if her economic perspective does not agree with that enforced by regulations. Hence, by enforcing epistemic constraints in the regulations, the policy maker can induce a range of counter-intuitive, but rational, behaviors, from information avoidance to overevaluation of barely relevant information, in the agents obeying the regulations.

Information is related to the legal distinction between "negligence", i.e. the failure to exercise the care against a risk which might be expected of a prudent person in the same circumstances, and "recklessness", i.e. knowing and willful exposure of others to a risk [3]: implicit in this distinction is whether or not the agent had knowledge of the risk in question. Between these, recklessness is typically considered to be the more severe transgression, carrying with it a heavier penalty for the agent. From a legal standpoint, agents may prefer to avoid information to be merely responsible for "negligence", hedging against the heavier penalties of "recklessness". This provides a general example of how well-intentioned regulations can prompt seemingly counter-intuitive and counter-productive "willful ignorance" in agents acting under these constraints. A related example in the engineering domain, presented by Rayner [17], is an environmental remediation project where few field samples were taken, so as to minimize the possibility of these samples contradicting a computer model which predicted that the project was on course for success.

In other contexts such as non-cooperative games, where agents compete against each other, revealing a piece of information to all agents may have a negative impact to some of them, as the negative effect of the competitors being informed and adjusting their policies surpasses the direct VoI. Being aware of this, some agents prefer to avoid having certain information collected, when it must be shared with others, as the overall VoI is negative for them. For example, an entrepreneur sharing a market with a competitor may find a piece of information irrelevant for her, but key for her competitor, who can improve his strategy and reduce her share of the market [2]. In that case, the impact of information is clearly negative, when assessed by that agent. These mechanisms are related to the topic of "information avoidance", extensively studied by social sciences, including psychology and behavioral economics [7, 10, 21].

Recent interest in VoI analysis for civil infrastructure systems is testified to by works such as that of Pozzi and Der Kiureghian [13], Srinivasan and Parlikad [19], Straub [20], Zonta et al. [24], Qin et al. [15], Goulet et al. [8], Thons [22]. VoI is used as an objective function to be maximized for optimizing information collection by Malings and Pozzi [11] and Memarzadeh and Pozzi [12]. A discussion of the effects of the discrepancies between agents' preferences in relation to civil infrastructure management is in Pozzi et al. [14], Tonelli et al. [25] and Verzobio et al. [26].

In this paper, we illustrate how the structural properties of the VoI, of observing the state of an engineering system, change depending on the external epistemic constraints, and discuss how regulatory design can alleviate undesired attitudes toward information. Our motivating question for writing this paper was: "why do rational agents sometimes prefer not to know?" i.e. "how can the VoI be negative, in some conditions?" To clarify our scope, we note that some phenomena that can also be described as cases of negative VoI are not the core of this paper. Experience suggests that sometimes it is better to neglect or refuse irrelevant information because "too much information can harm". This can happen because supposedly free information is not actually free, when considering all costs related for collecting and processing it, and so agents should neglect information with nil VoI. Also, it is a common

experience that an agent can take a decision, then revise it based on noisy measures, while the prior decision was actually correct and the information has misled her. However, we point out that "information cannot hurt" is a principle holding in the expected prior prediction while, in a given single empirical realization, it may not hold true. We can also argue that if the processing model is incorrect, then the impact of information can be detrimental. For example, consider an agent overconfident in the precision of a sensor, or unaware of its systematic bias. That agent can be misled by the information, so that she would have done better without it. Again, previous results hold under model consistency: after all, probability models ignorance, and if an agent suspects that a model may be inappropriate, she should extend it until it captures the complete uncertainty in the relation between the system's state and measures.

2. Problem formulation

We focus on a simple and paradigmatic case of reliability assessment and control. Consider a system exposed to the risk of failure, for example a structural system prone to deterioration and collapse. If the failure occurs, the agent controlling the system must sustain a significant loss, but the failure also has consequences at the societal level. The probability of such failure is assessed considering the current available information on the uncertain system condition, the stochastic evolution of capacity and demand. Expensive maintenance actions are available to the agent, to mitigate the risk. Hence, she is facing a decision-making problem under uncertainty: should she repair the system or not? We assume the agent is rational, so that her behavior is modeled by the principle of minimizing expected loss, and risk-neutral. To decide the best course of action, the agent compares the expected loss of doing nothing, accepting the failure risk, with that related to maintenance actions. However, the agent must also follow rules defined by a societal regulation: in practice, these rules may prescribe specific actions to be undertaken in certain circumstances. We simplify and generalize such regulations into a constraint to the agent's decision: the agent must take a maintenance action when the probability of system failure exceeds a threshold of societally-accepted risk. In the context of structural systems, for example, such regulation is motivated by the societal need of reducing the rate of failure events: a building code prevents owners from opening unsafe buildings to the public.

In this setting, we consider that information can be collected for reducing the uncertainty on the system condition state, for example by inspecting it. The agent has to assess whether it is worth inspecting, trading off the inspection cost with the cost reduction due to a more appropriate decision based on the information. Fig. 1 depicts the corresponding decision graph. Chance node x indicates the uncertain system's state, node y the inspection outcome, the two decision nodes are controlled by the agent: a indicates the maintenance action and I the execution of the inspection. The maintenance action is subject to the societal constraint. Diamonds indicated the costs and losses for the agent and for society, related to the inspection and maintenance action (the latter including the failure risk): the agent aims at minimizing the expected sum of L and L_D , while society should calibrate the constraint

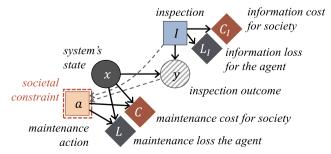


Fig. 1. Decision graph for the analyzed problem.

to minimize the expected sum of C and C_I .

Analyzing this problem, our goal is twofold: first, we aim at assessing whether inspecting is convenient for the agent, by quantifying the VoI under societal constraint; second, we aim at evaluating the overall impact of the constraint, to discuss the pros and cons of such constraint from the societal perspective. We outline the properties of unconstrainted VoI in the next Section, the effect of epistemic societal constraints in Section 4, we discuss regulatory design in Section 5 and draw some conclusions in Section 6.

3. Why "information never hurts" in unconstrained decision-making

3.1. Unconstrained decision making under uncertainty

We start modeling the problem of information collection when the agent is not subject to any external constraint (i.e., there is no societal regulation). In this section we illustrate why, in this case, the expected prior loss is a concave and thus continuous function in the convex domain of possible beliefs; hence, why the VoI is always non-negative in unconstrained decision-making.

As outlined in Fig. 1, consider the agent controlling a system and facing a one-shot decision, aiming at minimizing her expected loss. She has to select action a among a finite set of options $A = \{1, 2, ..., |A|\}$ of size |A|, while her loss L(x, a) depends on the selected action and on the state x of the system, defined on finite set $X = \{1, 2, ..., |X|\}$ of size |X|. Set A includes maintenance actions, as repairing or replacing a component, and the do-nothing option. List X can include specific malfunctions, damages, up to the failure of the system, together with the system well-functioning. The agent's belief about the state of the system can be summarized in vector \mathbf{b} , of length |X|, with entry i defined as $b_i = \mathbb{P}[x = i | G]$, and it is evaluated considering all relevant background information G. Hence stochastic vector b lists non-negative entries and it has unitary norm-1: $\sum_{i=1}^{|X|} b_i = 1$. These belief components include the probability of damage and failure of the system. If she exactly knows that the system is in state j, than $\mathbf{b} = \mathbf{e}_j$, where \mathbf{e}_j is the j unit vector in the standard basis (i.e. all entries of b are zeros except for a unitary entry at position *j*). The domain Ω_B of the possible beliefs is a convex set of dimension (|X| - 1), resulting from the intersection of the hypercube where all |X| components are between zero and one and the hyperplane where the norm-1 is unitary.

We define $l_a(\mathbf{b}) = \mathbb{E}_{\mathbf{b}} L(x, a)$ as the expected loss selecting action a under belief \mathbf{b} (where $\mathbb{E}_{\mathbf{v}} f(i,j)$ indicates the expectation of function f with vector \mathbf{v} assigning the distribution of variable i). Introducing, for each action a, |X|-dimensional vector λ_a of possible losses depending on the system state, whose entry x is $\lambda_a(x) = L(x, a)$, we derive by the definition of expectation that the expected loss under action a is a scalar product, a linear function of belief \mathbf{b} : $l_a(\mathbf{b}) = \lambda_a^T \mathbf{b}$.

The decision problem can be expressed as the minimization on the set of functions $\{l_1, l_2, ..., l_{|A|}\}$. The optimal loss l^* , as a function of belief **b**, is:

$$l^*(\mathbf{b}) = \min_{a \in A} l_a(\mathbf{b}) = \min_{a \in A} \lambda_a^{\mathrm{T}} \mathbf{b}$$
(1)

And the corresponding optimal policy Π^* , mapping the current belief into the optimal action, is defined by using "argmin" instead of "min" in the previous equation. As is clear from Eq. (1), the optimal loss l^* is the lower envelope of the set of |A| linear functions defined by "lambda" vectors $\{\lambda_1, \lambda_2, ..., \lambda_{|A|}\}$, and therefore, it is a concave, and hence continuous function. We can also define the optimality basin for action a as the subset $B_a \subseteq \Omega_B$ of the belief domain where that action is optimal: $\mathbf{b} \in B_a \Leftrightarrow \Pi^*(\mathbf{b}) = a$. Each basin is a convex set as it is the subset of a convex set resulting from imposing linear inequalities (of course, some basins can also be empty). Also, the set of basins are disjoint (defining, if necessary, an arbitrary criterion to pick one action among those with equal expected losses) and it completely covers Ω_B .

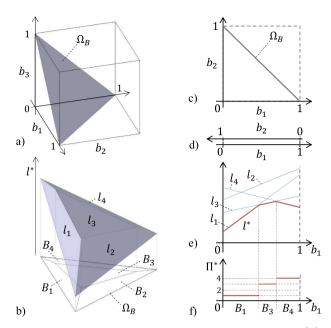


Fig. 2. Convex domain Ω_B of possible belief when the number of states |X| is 3 (a). Example of expected loss functions when the number of actions |A| is 4, with optimal loss function l^* and convex domain of optimality for each action (b). Domain Ω_B for the binary state, when |X| is 2 (c). Re-parametrized belief domain of size |X|-1=1, as a unit segment (d). Example of expected loss functions for |A| equal to 4, with optimal loss function l^* (e). Optimal policy and convex domain of optimality for each action (f).

Fig. 2 provides two examples of decision problems. Picture (a) shows the triangular domain Ω_B of possible beliefs when the number of states |X| is 3, i.e. a convex figure of |X|-1=2 dimensions: a triangle. On that domain, partitioned in the set of convex optimality basins, picture (b) illustrates an example set of expected loss functions related to 4 actions, and the convex optimal loss function l^* . Picture (c) refers to the binary state, when |X| is 2, and Ω_B is the diagonal of a square, that can be described in |X|-1=1 dimension: as $b_1+b_2=1$, any of these two coordinates can completely define that domain, as indicated in picture (d). In this case the entire belief can be described by the probability of the system being in one state, e.g. the probability of failure. Picture (e) shows the expected losses for 4 actions, and the function l^* . Picture (f) reports the optimal policy and the optimality basins.

3.2. Belief updating after information processing

In this section we show how the measurement can be related to a distribution of possible posterior beliefs, and how the expected posterior belief is equal to the prior one. Bayes' formula is the key for processing information. We assume a measure or observation \mathbf{y} , that can be a multi-dimensional vector of arbitrary dimension, defined on domain Y, is related, directly or indirectly, to the system state x. The relation between the system and the measure is captured by emission function $p(\mathbf{y}|x)$, that defines the probability density of \mathbf{y} (or its discrete probability distribution, if \mathbf{y} is a discrete variable), when the system state is x. We define \mathbf{b}_{π} as the prior belief, before considering the measure (i.e., when information G mentioned in the previous section is the empty set \emptyset).

By combining prior belief and the emission function for a given realized measure \mathbf{y} , we derive the posterior belief $\mathbf{b}_{\omega}(\mathbf{y})$ by Bayes' formula, whose entry i is:

$$b_{\omega,i}(\mathbf{y}) = \mathbb{P}\left[x = i|\mathbf{y}\right] \propto p(\mathbf{y}|i)b_{\pi,i} \tag{2}$$

Eq. (2) defines a deterministic map between measure y and

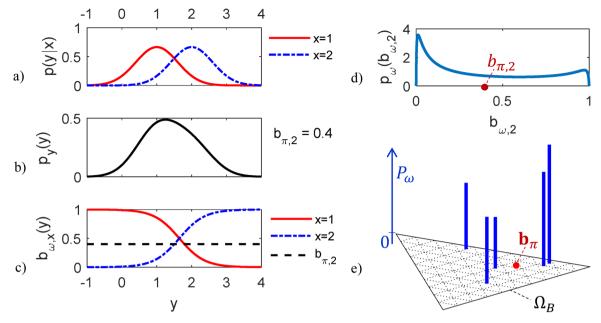


Fig. 3. Example of the inference process. Continuous emission functions (a), probability of measure (b) and posterior belief (c), as a function of the measure value y, for the binary state, when |X| is 2. Corresponding distribution of the posterior belief (d), in terms of the second component. Example with discrete observation, with |Y| equal to 5, when |X| is 3.

posterior belief \mathbf{b}_{ω} . It shows how, under a specific emission function, the prior belief is moved to the posterior one, following a realized value of the measure. In the prior condition, measure \mathbf{y} is a random variable, modeled by predictive distribution:

$$p_{\mathbf{y}}(\mathbf{y}) = \sum_{i=1}^{|X|} p(\mathbf{y}|i) b_{\pi,i}$$
(3)

which serves as a normalizing constant for Eq. (2). Consequently, in the prior condition the posterior belief \mathbf{b}_{ω} is also a random variable, whose distribution p_{ω} derives from the transformation of distribution p_y via the map of Eq. (2). Variables \mathbf{y} and \mathbf{b}_{ω} are of the same type: if the set of possible measures is discrete and finite, so is that of posterior beliefs while, if the former set is continuous, so is the latter one. Moreover, by the consistency of probability calculus, as:

$$\mathbb{E}_{\mathbf{y}}[p(x|\mathbf{y})] = \int p(x|\mathbf{y})p_{\mathbf{y}}(\mathbf{y})d\mathbf{y} = p(x)$$

we see that the expected posterior belief, evaluated in the prior condition, has to be identical with the prior one: $\mathbb{E}_{\omega}\mathbf{b}_{\omega}=\mathbf{b}_{\pi}$ (where \mathbb{E}_{ω} indicates the expectation with respect to distribution p_{ω}).

Fig. 3(a–d) shows an example of the inference process, when |X| is 2 and there is a univariate continuous measure y. Emission functions are with mean $\mathbb{E}[y|x=i]=i$ (a), and $\mathbb{V}[y|x=i] = \sigma_{\varepsilon}^2 = 1$, thus distribution $p_{\mathbf{v}}$ is the mixture of two Gaussian components (b), depending on prior belief $\mathbf{b}_{\pi} = [0.6 \ 0.4]^{\mathrm{T}}$. The posterior belief, as a function of the observed value y, is as plotted in graph (c). Graph (d) shows the corresponding distribution of the posterior belief, in terms of its second component. Collecting the measure can be seen as trading the prior belief for a random realization of the posterior belief, generated from a distribution whose expected value is identical to the prior belief. In this trading, the belief can move toward the extreme regions of the domain, where the state x is perfectly known. Graph (e) illustrates an example when |X| is 3, and only |Y| = 5 values of observation y are possible. Hence, the posterior belief can assume only 5 values, and discrete distribution P_{ω} defines the probability of each possible outcome.

3.3. Pre-posterior analysis and value of information

The Value of Information (VoI) is the expected loss reduction due to the availability of a measurement. In the prior condition, an agent would select action $\Pi^*(\mathbf{b}_\pi)$, obtaining expected loss $l_\pi^* = l^*(\mathbf{b}_\pi) = l^*(\mathbb{E}_\omega \mathbf{b}_\omega)$. Posterior optimal action $\Pi^*[\mathbf{b}_\omega(\mathbf{y})]$, using the same policy identified above, and loss $l^*[\mathbf{b}_\omega(\mathbf{y})]$ depend on the realized measure \mathbf{y} . The expected posterior loss is $l_\omega^* = \mathbb{E}_\omega l^*(\mathbf{b}_\omega)$. The VoI can be defined as the difference between prior and posterior expected loss:

$$VoI = l_{\pi}^* - l_{\omega}^* = l^*(\mathbb{E}_{\omega} \mathbf{b}_{\omega}) - \mathbb{E}_{\omega} l^*(\mathbf{b}_{\omega})$$

$$\tag{4}$$

We now define a subset of the belief domain, Ω_Y , as the union of the prior belief and the set of possible posterior beliefs for any possible realization of the measure, so that p_ω is nil outside that subset. Only the value of function l^* on Ω_Y is relevant for assessing the VoI. Hence, the VoI can be computed by combining two functions defined on Ω_Y : the optimal loss function l^* and the predictive distribution p_ω , whose mean value is \mathbf{b}_π .

As the expected loss function l^* is a concave function, we derive from Eq. (4) and Jensen's inequality [9] that VoI is non-negative for any prior belief \mathbf{b}_{π} , any loss function L and any emission function $p(\mathbf{y}|x)$. This is the "information never hurts" principle.

Also, consider the case when function l^* is linear in Ω_Y : in that case, we can commute function l^* and the expectation in Eq. (4), the difference vanishes and the VoI is zero. A trivial limit case is when the measure does not affect the belief, so \mathbf{b}_ω is identical to \mathbf{b}_π for all possible realized measures, and Ω_Y contains just one point. This can happen if variables x and y are independent: clearly, statistically negligible information has no value in decision making. A more general case is that when $\Omega_Y \subseteq B_{a_*}$, that is, for any possible realized measure the posterior optimal action is equal to the prior one, so that l^* is a linear function in Ω_Y . Again, if no possible measure output is able to change the prior action, then the VoI is zero. Conversely, if there is a positive probability that the posterior optimal action is different than a_* , because the corresponding posterior optimal loss is strictly less than l_{a_*} , then the VoI is strictly positive.

In the Appendix, we prove another property of the VoI, related to the continuity of l^* : information with almost-negligible effect on the belief has almost negligible VoI. Intuitively, if the posterior belief is

close to the prior one for all possible realizations of the measures, the expected posterior loss tends to be similar to the prior loss, and the VoI tends to be zero.

3.4. Posterior loss and value of perfect information

In this section, we show how posterior loss l_{α}^* , seen as a function of prior belief \mathbf{b}_{π} , is also a concave function, never above l_{π}^* . That is, we show that the expected loss after taking a measurement is never higher than the expected loss before measuring. How should a rational agent react to information? As seen before, she should update her belief using Bayes' formula, and then act as if the updated belief was her prior one, identifying optimal action a_* and getting corresponding expected loss l*. However, the posterior decision can also be described in terms of a "conditional plan", describing how to react to any realized observation, by selecting a specific action [18]. If we now assume a finite set Y of |Y|observable values, the entire plan is a set of |Y| instructions such as "if observation is y = j, then take action a = k". For example, a plan prescribes to execute the first action no matter what the observation is, another to execute that action only after some observed values, and do the second otherwise. Plan j, φ_i : $Y \to A$, is therefore a map from the set of observable values Y to the set of actions A. The set M of all possible conditional plans contains $|M| = |A|^{\wedge}|Y|$ members. We can compute the expected loss $l'_{i,j}$, executing plan φ_i when the state x is equal to i, combining the emission function and the loss function:

$$l'_{i,j} = \sum_{y=1}^{|Y|} p(y|i)L[i, \varphi_j(y)]$$
(6)

So $l'_{i,j}$ is a weighted average of the values of function L related to the system being in state i, and so it is bounded by the minimum and the maximum of these values. We can define a new loss matrix L', of size |X| by |M|, populated following the formula in Eq. (6). Now, the conditional plan plays the role played by the action in the original setting: depending on her prior belief, the agent selects the conditional plan that minimizes the expected loss. Following the same approach outlined above, we define linear function $l'_j(\mathbf{b}) = \alpha_j^T \mathbf{b}$ as the expected loss following conditional plan j, where vector \mathbf{a}_j , of size |X|, lists the expected losses for all states, as $\alpha_j(i) = l'_{ij}$. Using Eq. (6), we can derive vector \mathbf{a}_j from the set of lambda vectors, as $\alpha_j(i) = \sum_{y=1}^{|Y|} p(y|i)\lambda_{\varphi_j(y)}(i)$. Starting from a specific belief \mathbf{b} , the optimal policy Π^* can be translated into the specific plan related to the minimum loss. So optimal posterior loss l^*_{ω} can be expressed as the lower envelope of a set of |M| linear functions, defined by the set of "alpha" vectors $\{\alpha_1,...,\alpha_{|M|}\}$:

$$l_{\omega}^{*}(\mathbf{b}) = \min_{j \in M} l_{j}^{\prime}(\mathbf{b}) = \min_{j \in M} \alpha_{j}^{\mathrm{T}} \mathbf{b}$$
(7)

Hence, l_{ω}^* is a concave, continuous function. While we have showed that this property holds for a discrete set of |Y| observable values, the same property holds for an infinite number of possible observations, when |Y| goes to infinity, and even when the set of observable values is uncountable.

Hence, we can define the VoI as a function on Ω_B : $VoI(\mathbf{b}) = l_\pi^*(\mathbf{b}) - l_\omega^*(\mathbf{b})$, depending on the loss matrix and emission probabilities.

Also, among the |M| conditional plans, |A| of them assign the same action for all possible realized measures: those plans can be executed in the prior condition, as they do not require access to the information, and they are equivalent to the prior selection of an action, independent of the observation. Hence, the set of alpha vectors includes that of lambda vectors as a subset. By comparing Eqs. (1) and (7), we conclude that $l_{\omega}^* \leq l_{\pi}^*$, and this provides another proof that the function VoI, defined as the difference between l_{π}^* and l_{ω}^* , is non-negative everywhere. In light of this, we can read the VoI as the benefit of adopting a more flexible reaction plan: while the expected loss l_{π}^* derives from selecting a plan among the |A| possible rigid ones, independent of the measure, l_{ω}^* derives from the selection among the broader set of |M| adaptive ones

that react to the measures.

The VoI is thus the difference of two concave functions: it is continuous (as the difference between two continuous functions is continuous), but it is not necessarily concave. Also, at the corners of the belief domain, where the state is known (vector \mathbf{b} lists only zeros except for one unitary entry), then the posterior belief is surely equal to the prior one, so functions l_0^* and l^* are identical, and the VoI is nil.

In the limit case of perfect information, observable variable y is equivalent to state variable x, so that the posterior belief collapses on one unit vector of the standard basis: $\mathbf{b}_{\omega}(i) = \mathbf{e}_i$, with probability b_{π_i} . For those posterior beliefs, the losses are minima of function L, that we list in vector λ .:

$$l^*(\mathbf{e}_i) = \lambda_*(i) = \min_{a} L(i, a)$$
(8)

So that the expected loss with perfect information is linear in the belief: $l_{\nu}^{*}(\mathbf{b}) = \lambda_{\mathbf{k}}^{\mathrm{T}}\mathbf{b}$. As $\lambda_{\mathbf{k}}(i)$ is a lower bound for each possible alpha vector component $a_{j}(i)$, function l_{ν}^{*} is a lower bound for each posterior function l_{ω}^{*} , for any observation model. In summary, inequality $l_{\pi}^{*} \geq l_{\nu}^{*}$ holds in the belief domain Ω_{B} , and these functions have the same value at the vertices of the domain. The VoI for perfect information, usually called Expected Value of Perfect Information (EVPI), is given by $EVPI(\mathbf{b}) = l_{\pi}^{*}(\mathbf{b}) - l_{\nu}^{*}(\mathbf{b})$: being the difference between a concave function and a linear one, function EVPI is also concave. As l_{ω}^{*} dominates l_{ν}^{*} , the EVPI provides an upper-bound for the VoI: $0 \leq VOI \leq EVPI$.

In Fig. 4 we illustrate an example where |X|=2, i.e. for the binary state, and |A|=3. The loss matrix is given by the 3 lambda vectors $\lambda_1=[0\ 6]^T$, $\lambda_2=[1\ 4]^T$ and $\lambda_3=[4\ 2]^T$. Emission functions are both Gaussian, similar to Fig. 2, again with mean $\mathbb{E}[y|x=i]=i$ and variance $\mathbb{V}[y|x=i]=\sigma_\varepsilon^2$ for both states. In graph (a), posterior loss function l_ω^* is reported for σ_ε equal to 1, 0.5 and 0 (the latter being the case of perfect information). Clearly, the posterior loss is lower if σ_ε is lower, down to linear function l_ν^* . For a high value of σ_ε , the posterior loss is indistinguishable from the prior one, and the VoI is nil. Graph (b) shows how the VoI is affected by σ_ε : function EVPI is concave, while VoI has local maxima for higher values of σ_ε .

3.5. Interpretation and use of the VoI depending on the attitude toward risk

In our definition of the VoI, no cost or loss for the very act of collecting information is explicitly included. Information only has an effect on the epistemic belief of the agent, and not on the cost she has to pay. In the previous section, all loss functions and VoI can be intended as being in "loss units", corresponding to the utility unit [23]. As such, the

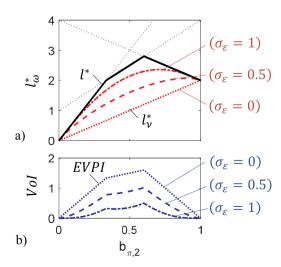


Fig. 4. Example of posterior expected loss, related to continuous observations (a), corresponding VoI (b).

VoI represents the expected loss reduction due to free information. We have seen that this is always non-negative, indicating that free "information never hurts". For a version of the same principle from the standpoint of information theory, see Cover and Thomas [5]. In terms of decision-making, this shows that the rational agent should always accept free information. But is the actual numerical VoI of any use? If the agent had to decide between collecting free information and an alternative action whose effect can be summarized by an expected loss reduction, then the agent should opt for the former option only if its VoI is above the reduction due to the latter. However, we may be more interested in a different setting, where information is expensive, and the agent has to decide on buying it before selecting an action. Information cost may impact the overall agent's loss. On one hand, the agent can assess the optimal expected loss without information, as we have seen above; on the other, to predict the overall effect of information, she can update the original loss function L including the impact of information cost for any pair of state and action and, using this updated loss function, she can assess the expected posterior loss. She should buy the information only if that posterior loss is less than that without information. This procedure, that we can call "the long route" is affected by the sensor cost from the beginning and the VoI, as defined before, plays no role in it. If we now consider a different information cost, we have to repeat the procedure from scratch. We may be tempted to follow a different and shorter route, first computing the VoI in terms of losses, and then converting this quantity in a monetary value, to compare with the sensor cost. However, if the conversion between monetary value is non-linear, as it is for risk-seeking or risk-adverse agents, it is unclear where this conversation between VoI and monetary cost should be made: on what interval along the cost to loss conversion? To find a monetary value V_M consistent with the long route, one has to equate the prior loss with the posterior one as a function of unknown V_M , and solve that non-linear equation to identify V_M (this procedure is illustrated by [4]). However, there is no guarantee that a single solution does exist. As an extreme case, consider a peculiar agent that loves monetary costs of even amounts of dollars, and hates costs of odd amounts. Also, consider that all costs related to each pair of state and action are even. Clearly, for this agent the VoI (related to free information) is not of much use for deciding about buying an information whose cost is an odd amount of dollars, and not a single solution for V_M may exist.

However, for risk neutral agents the VoI, as defined above, is also directly useful for deciding whether to buy expensive information. Indeed, for those agents all losses, and so the VoI itself, can be directly expressed in monetary costs. Moreover, due to the superposition of effects, the expected posterior loss buying an information for cost L_I (as outlined in Fig. 1) is simply the sum of L_I and the expected posterior loss with free information. So it is worth buying the information when $l_{\omega}^* + L_I \leq l_{\pi}^*$, i.e. when $L_I \leq VoI$. Hence, by invoking the superposition of effects, we can interpret the VoI as the maximum cost the rational risk neutral agent should pay for getting that information, as paying more would induce an overall negative effect. Let us call $D \subseteq \Omega_B$ the subset of the belief domain where information should be collected. If cost L_I is zero then D coincides with Ω_B , as the VoI is non-negative. When L_I is positive, then D is a proper subset of Ω_B , that does not include its corners corresponding to perfect knowledge, where the VoI is zero. As the VoI is not generally concave, set *D* is generally not convex, or even connected. However, as function EVPI is concave, set D corresponding to perfect information is convex. If cost L_I is above the maximum value of the VoI, set *D* is empty.

Fig. 5 is related to the example in Fig. 4, when $\sigma_{\varepsilon} = 0.5$ and $L_{l} = 0.5$. It shows that the set D where $l_{\omega}^{*} + L_{l}$ is above l_{π}^{*} is also that where the VoI is above L_{l} .

We briefly summarize the properties of the VoI in unconstrained decision making, discussed above. The VoI is a continuous non-negative function in the belief's domain. Information with almost-negligible effect on the belief has almost negligible VoI.

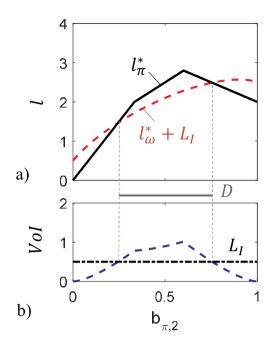


Fig. 5. Comparison between prior and posterior loss for information at cost L_I (a), VoI and information cost (b).

4. Vol when acting under external constraints

4.1. External epistemic constraints

We now consider the effects of external constraints, e.g. societal regulations, in decision-making. Most external constraints can be captured by appropriate adjustments of loss matrix L: e.g., if one action a is forbidden when the state is x, this can be modeled by assigning an infinite loss to this pair, and similarly if a cost exceeds a budget constraint, to claim that those combinations cannot be accepted. Hence, those constrains can be embedded in the pre-processing that defines function L and, given that the structural properties that we have listed above hold for any loss function, they hold under any of these constraints

A more complex setting occurs when the constraint restricts the available actions depending on the belief, and not depending on the state. We call these "epistemic constraints", as the belief is an epistemic descriptor. As outlined in Section 2, our motivating example is that of an agent following a set of regulations forbidding some actions when the risk is too high. These constraints cannot be embedded in any equivalent loss matrix. The agent's loss is modeled by matrix L but, following those constraints, the agent cannot generally implement her optimal policy Π^* , as this may violate them. Hence, to investigate the exploitative and explorative behavior of an agent acting under external epistemic constraints we assume she has to follow a sub-optimal policy: the best one compatible with the constraints. In the next section, we show how the structural properties of VoI are not preserved in this setting

We add a clarification about the nature of epistemic constraints. The decision-making problem under an epistemic constraint can be translated into an equivalent unconstrained problem where the constraint is represented by a change occurring in the external world the agent is interacting with. For example, the activation of the external epistemic constraint forbidding the agent to open the asset to the public when the probability of failure is above a threshold can be equivalently intended as a physical impediment (say a team of policemen), preventing that opening, which is activated following a mechanism consistent with that of activating the epistemic constraint. As a result of this exercise in the

art of the analogy, the agent may think that now she has re-formulated the problem as one without any epistemic constraint, where the general properties of the unconstrained VoI hold. Now, while this is formally correct, we note that, in this setting, collecting information not only affects the agent's belief, but also triggers reactions of the "external world", and the properties described in Section 3 may not hold true. For example, even if an agent decides to maintain her posterior action identical to the prior one (and so neglecting the information), her expected posterior loss is not necessarily identical to the prior one, if others react to the information. So our analysis of external epistemic constraints applies to both equivalent settings: when the constraint is intended as an "internal" rule or when it models an external force acting on the agent.

4.2. Acting under external epistemic constraints

We consider that an agent adopts the best policy $\Pi^+\colon \Omega_B \to A$ consistent with the external epistemic constraints, that is generally different with respect to the optimal unconstrained one Π^* . The constraints, however, are inactive for perfect beliefs, at the corners of domain Ω_B , as constraints acting on those cases can be directly embedded in the loss matrix L. The corresponding expected loss l^+ is:

$$l^{+}(\mathbf{b}) = l_{\Pi^{+}(\mathbf{b})}(\mathbf{b}) \tag{9}$$

While function l^* is concave and continuous, as it results from a minimization, function l^+ is not necessarily concave or even continuous. Clearly, it cannot be below l^* : $l^+ \ge l^*$, as the latter represents the minimum loss. So the non-negative expected loss increment due to constraints, Δl , is:

$$\Delta l(\mathbf{b}) = l^{+}(\mathbf{b}) - l^{*}(\mathbf{b}) \ge 0 \tag{10}$$

Increment Δl is also not necessarily continuous. We can express the prior loss as a function of the expected prior increment $\Delta l_{\pi} = \Delta l(\mathbf{b}_{\pi})$, as:

$$l_{\pi}^{+} = l^{+}(\mathbf{b}_{\pi}) = l_{\pi}^{*} + \Delta l_{\pi} \tag{11}$$

And the expected posterior loss as related to the posterior increment $\Delta l_{\omega} = \mathbb{E}_{\omega} \Delta l(\mathbf{b}_{\omega})$, as:

$$l_{\omega}^{+} = \mathbb{E}_{\omega} l^{+}(\mathbf{b}_{\omega}) = \mathbb{E}_{\omega} l^{*}(\mathbf{b}_{\omega}) + \mathbb{E}_{\omega} \Delta l(\mathbf{b}_{\omega}) = l_{\omega}^{*} + \Delta l_{\omega}$$
 (12)

4.3. Value of information acting under an external epistemic constraints

In this section, we show how, when acting under an external constraint (e.g. a regulation), the VoI for a piece of information can be different from the VoI evaluated without considering the constraint. That is, we show how external constraints can cause an agent to overvalue or under-value information. We define VoI^+ to be the value of information if acting following policy Π^+ . This quantity is related to the unconstrained VoI by increment $\Delta VoI = \Delta l_{\pi} - \Delta l_{\omega}$, as:

$$VoI^{+} = l_{\pi}^{+} - l_{\omega}^{+} = VoI + \Delta VoI$$

$$\tag{13}$$

Both Δl_{π} and Δl_{ω} are non-negative, and ΔVoI can be either positive or negative. As functions of the belief, Δl_{π} is not necessarily continuous, nor is ΔVoI , while Δl_{ω} is continuous under certain assumptions as we will see in a later example. At the corners of belief domain Ω_B , information has no impact and both ΔVoI and VoI^+ are zero, as is VoI.

Due to the non-negativity of the increments, we can bound ΔVoI as $-\Delta l_{\omega} \leq \Delta VoI \leq \Delta l_{\pi}$. However, in general Δl_{π} and Δl_{ω} are arbitrarily large, and so ΔVoI and VoI^+ can be positive or negative, and their modulus arbitrary large. Negativity of the VoI can explain information avoidance for rational agents: actually, if VoI^+ is negative they should be ready to pay up to $-VoI^+$ to avoid collecting information.

Fig. 6(a) shows a simple example of those limit cases. The first three actions are related to a low loss while the other two are related to a loss δ higher. From prior belief b_B , the posterior belief can be b_A or b_C , depending on the realized observation, from prior belief b_C , the posterior

belief can be b_B or b_D . Thus, from prior belief b_B , the agent would like to collect the information, as the corresponding VoI^+ is δ . From prior belief b_C , instead, the agent prefers to avoid information, as VoI^+ is $-\delta$. This shows how VoI^+ can be positive or negative, as δ can be arbitrarily large. Fig. 6(b,c) shows the example of Figs. 4 and 5, under an external epistemic constraint forcing action 2 for belief 15% $\leq b_{\pi, 2} \leq$ 85%. The corresponding function l^+ is discontinuous. We consider the same emission functions introduced above, with $\sigma_{\varepsilon} = 1$. Function Δl_{π} is discontinuous and it is zero where policies Π^+ and Π^* agree with each other. Function l_{ω}^+ (and so function Δl_{ω}) is continuous, as the continuous distribution of the posterior belief acts as a smoother of function l^+ . Where l_s^+ is continuous, a jump of function l^+ is identical of that of VoI^+ . Both functions ΔVoI and VoI^+ are discontinuous, and they assume positive and negative values in Ω_B . As illustrated above for the general case, we note that when $b_{\pi, 2}$ is just below 15% the VoI^+ is indeed negative. Another interesting belief is at $b_{\pi,2} = 85\%$, when the expected loss is maximum: any update increasing belief component $b_{\pi, 2}$, even for an arbitrary small quantity (i.e., even for almost negligible information), will significantly reduce the loss, hence VoI+ is not only positive, but much higher than VoI. Generally, VoI+ can be higher, equal or lower than VoI, depending on the prior belief. In the Appendix, we show how almost-negligible information can have arbitrarily large value, in the constrained setting.

We close this Section by focusing on two special cases. If, at the prior belief \mathbf{b}_{π} , the two policies agree with each over (i.e., $\Pi^+(\mathbf{b}_\pi) = \Pi^*(\mathbf{b}_\pi)$) then Δl_π is zero, so $\Delta VoI = -\Delta l_\omega$ cannot be positive, and VoI^+ is not higher than VoI. Hence the constrained agent gives less value (or the same value) to the information with respect to the unconstrained one, due to the sub-optimality of the posterior action. Indeed, for such a belief, the constraint is inactive in the prior condition, and the information exposes her to fall under the constraints: the non-positive ΔVoI quantifies this effect. Conversely, it can be the case that, while Π^+ and Π^+ are different at \mathbf{b}_π , they are identical for all reachable posterior beliefs in Ω_Y : in that case, Δl_ω is zero and VoI^+ is not lower than VoI: the agent is penalized by the constraint only in the prior setting, and the information has the additional value of letting the agent escape the constraint. This latter case happens, for example, in the case of perfect information.

5. Regulatory design to promote well-balanced information collection

5.1. Reasons for external epistemic constraints and their impact on information collection

Why are external epistemic constraints imposed on the decisionmaker? We assume they are imposed by society, through regulations, with the purpose of influencing and controlling private decisions towards a common good ([6] presents a recent analysis in the context of structural design). The actions taken by the agent have public economic consequences that society can assess. Once those consequences are summarized in cost function C, as illustrated in Fig. 1, society can identify the basins of optimality corresponding to the optimal policy according to societal evaluation, following the approach outlined above. If societal cost C differs from agent's cost L, the optimal policy identified by society will differ from that identified by the decisionmaker. Assessing function L and predicting the agents' behavior, society can calibrate a regulation. This regulation can forbid certain actions to be taken under specific circumstances, namely under certain epistemic beliefs. By appropriate regulation, society can enforce the adoption of its optimal policy by the agent. However, what is the effect of this enforcement on information collection? Society assigns a value to any possible information, and it wishes information to be collected when this value is above its cost. For example, society would like relevant information to be always collected, as, from their perspective, it "never hurts". However, the VoI as assessed by the decision-maker acting

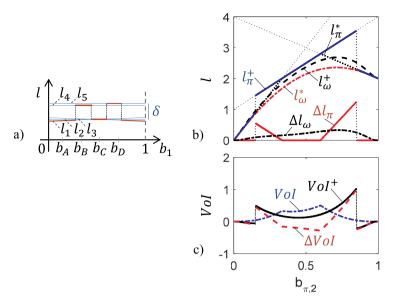


Fig. 6. Example of loss function for illustrating the lack of sub-optimality bounds when acting under external constraints (a), Example of decision making problem under external constraints: loss function (b) and VoI (c).

under external epistemic constraints can be different, much higher of that assessed by society, or even negative, as we have seen above. So, from the standpoint of society, agents collect information that is too expensive to be collected, and avoid other relevant information.

5.2. Incentives and penalties for influencing decision making

Policy makers can influence the behavior of agents in many ways, e.g. by funding an education system promoting respect to other citizens and to specific societal values. Here we are considering a more direct enforcement method: a regulation enforcing penalties and incentives to make loss function L, summarizing the economic cost the agent has to pay depending on the system condition and the adopted action, consistent with societal function C, modeling the corresponding societal costs. For example, society could enforce an economic penalty in case of failure of the asset, or pass laws to increase the liability of the agent, thereby transferring some costs from C to L. Let Z be the economic penalty (if positive) or incentives (if negative) society poses to the agent. Then, the overall loss for the agent L'(x, a) = L(x, a) + Z(x, a). Clearly, if Z(x, a) = C(x, a) - L(x, a), agent's and societal losses coincide. After the implementation of these corrections, epistemic constraints are redundant for the agent, as her optimal policy is aligned with that preferred by society. We note that in order to induce an agent's behavior consistent with societal desiderata, it is not necessary that C and L' are identical: it is sufficient that they differ for an arbitrary positive scaling factor $\gamma > 0$ and an arbitrary offset β . If so, the VoI assessed by the agent is just scaled by factor γ with respect to that assessed by society, and it has the same intuitive properties of the unconstrained VoI. The same effect can also be achieved by providing epistemic incentives and penalties, depending on the belief, instead of the state, e.g. a regulation can impose a penalty proportional to the probability of failure. Society can provide economic incentive $z_a(\mathbf{b}) = \boldsymbol{\zeta}_a^{\mathrm{T}} \mathbf{b}$, with $\boldsymbol{\zeta}_{a,i} = Z(i, a)$, if action a is taken under belief **b**: The adjusted expected loss function l'_a is:

$$l_a'(\mathbf{b}) = l_a(\mathbf{b}) + z_a(\mathbf{b}) \tag{14}$$

Society can also provide incentives directly for promoting information collection (however, it is harder to conceive that society could put a penalty against buying almost irrelevant information).

Overall, given the results of Section 3, incentives and penalties should aim at obtaining a concave loss function for the agent, so as not to trigger undesirable effects related to information avoidance and over-evaluation.

5.3. A maintenance problem under a single reliability constraint

We consider a simple decision making problem, following [14], where the agent has the responsibility for an asset whose state is binary: it can be intact (this is state x = 1), or damaged and doomed to failure (this is state x = 2), so |X| = 2. Thus, as in the examples shown in Figs. 4-6, the agent's belief is univocally defined by failure probability $b_{\pi, 2}$, which we now indicate simply as P, for ease of notation. She has two alternative options: do-nothing (this is action a_1) or repair the asset (this is action a_2), so |A| = 2. If she does nothing and the asset is doomed to failure, she has to sustain cost of failure L_F , but the failure can be avoided by paying repair cost L_R . Thus her unconstrained optimal policy is to repair the asset only if P is above $\overline{P} = L_R/L_F$, and optimal unconstrained loss is $l^*(P) = L_R \min\{P/\overline{P}, 1\}$. We assume that society evaluates the cost of failure and repair as C_F and C_R respectively, so the policy imposed by society, consistent with its proper evaluation, is to repair the asset if $P \ge \overline{P} = C_R/C_F$, and we assume that $\overline{P} < \overline{P}$. So the constrained loss l^+ is identical to l^* for P below \overline{P} and above $\overline{\overline{P}}$, but is different from l^* and equal to L_R in the interval $\Omega_M = [\bar{P}, \overline{\bar{P}}]$. The expected loss increment is:

$$\Delta l(P) = \begin{cases} L_R(1 - P/\overline{P}) & \text{if } P \in \Omega_M \\ 0 & \text{if } P \notin \Omega_M \end{cases}$$
 (15)

So the increment is nonzero only in the range of belief between the threshold imposed by society and that identified by the unconstrained agent. If we define $P_M = \mathbb{P}\left[\mathbf{b}_{\infty} \in \Omega_M\right]$ as the probability that the posterior belief falls in Ω_M , and $\mu_M = \mathbb{E}\left[P|P \in \Omega_M\right]$ as the expected belief in interval Ω_M , then the expected posterior loss increment due to the constraint is $\Delta l_{\infty} = P_M L_R (1 - \mu_M/\overline{P})$, because Δl is an affine function in Ω_M , according to Eq. (15). As outlined above, if the prior belief is outside Ω_M , $VoI^+ \leq VoI$. Conversely, if the prior belief is inside Ω_M and the posterior belief cannot fall in Ω_M (i.e. if $P_M = 0$), then $VoI^+ \geq VoI$. Function Δl is discontinuous at \bar{P} , and the jump's value is $L_R (1 - \bar{P}/\overline{P})$. The maximum possible value of this quantity, when \bar{P} is close to zero, is L_R .

In the example of Fig. 7(a,b), $L_F = 3L_R$ so that $\overline{P} = 1/3$. The societal constraint is $\overline{P} = 0.1$, so the maximum value of the discontinuous function Δl^+ , occurring at \overline{P} , is $0.7L_R$. The available information is the

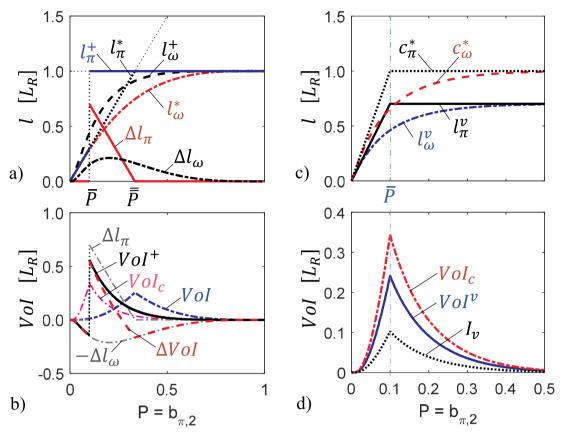


Fig. 7. Losses (a) and Vol (b) acting under external constraints, in the described example; corresponding quantities under the regulation (c-d).

same at that described in Section 2.2. and plotted in Fig. 2, with $\sigma_{\varepsilon}=1$, and the corresponding maximum value of continuous function Δl_{ω} is about $0.2L_R$. Graph (b) reports the corresponding VoI, expressed in terms of repair cost L_R . Acting under the external constraint, VoI^+ has a maximum value of $0.56L_R$ and a minimum value of $-0.14L_R$, their difference being the discontinuity of Δl^+ . As expected, variation ΔVoI is bounded by $-\Delta l_{\omega}$ and Δl_{π} . That graph also reports the unconstrained VoI, and function VoI_C , that will be introduced in the next Section. While these high values of failure probability are selected for the sake of readability of the graphs, the reader is referred to Pozzi et al. [14] for a similar example with smaller probability values, often encountered in reliability analysis of civil engineering components.

5.4. A corresponding regulation design

In the setting of previous Section, the agent would avoid collecting information when belief P is below \bar{P} and, actually, she would be willing to pay up to the 14% of the repair cost L_R , to avoid the information being collected. In the limit case of almost irrelevant information (i.e. when σ_{ε} approaches infinity) the agent can pay up to 35% of the repair cost (i.e., half of the discontinuity in Δl_{π}) to avoid collecting information or to collect information, depending on the fact that her belief P is just below or just above \bar{P} . If \bar{P} is close to zero, and under specific scenarios for almost irrelevant information, this VoI can be as large as the repair cost.

With references to the problem outlined in Section 5.3, we assume society assesses the cost of failure as $C_F > L_F$, while that of repair is consistent with that assigned by the agent: $C_R = L_R$. Thus, the societal desideratum is to fix $\bar{P} = C_R/C_F$. To align agents' preferences toward societal ones, society can enforce a penalty $(C_F - L_F)$ in case of failure. The same effect can alternatively be obtained by subsidizing the repair, proposing incentive $L_R(1-\bar{P})$, for lowering the repair cost. More generally, let ΔL_F and ΔL_R be the penalty and the incentive for failure

and repair, respectively. By imposing that the adjusted policy of the agent is consistent with that of society (i.e. $(L_R - \Delta L_R)/(L_F + \Delta L_F) = \bar{P}$), we get one linear equation, to define those quantities. Another equation can come from the condition that expected penalties should balance expected incentives, so that the regulation is self-sustainable. To do so, we need to assume a distribution of prior belief P. Let us define $P_F = \mathbb{P}\left[P < \bar{P}\right]$ the probability that the prior belief is below the threshold, and the optimal action is doing nothing, $\mu_F = \mathbb{E}\left[P|P < \bar{P}\right]$ is the expected probability of failure below the threshold, and $P_R = \mathbb{P}\left[P \ge \bar{P}\right] = 1 - P_F$ is the probability that a repair is needed. If so, incentives and penalties should be:

$$\Delta L_F = L_F \frac{\overline{\overline{P}} - \overline{P}}{r_{FR} + \overline{P}}; \quad \Delta L_R = \Delta L_F r_{FR}$$
(16)

where $r_{FR} = \mu_F P_F/P_R$. We note that the penalty can be enforced as an epistemic one, asking for payment $\Delta L_F P$ for an asset when the probability of failure is P. Again, constraints becomes redundant under this regulation. The adjusted optimal loss function for the agent is continuous and concave, so the properties of unconstrained VoI hold. The VoI assessed by the agent, that we can now call VoI^v for clarity, is proportional to that assessed by society, which we call VoI_c , and is less than that for society: $VoI^v = (1 - \Delta L_R/L_R)VoI_c$. To achieve perfect consistency in the evaluation, society should also provide to the agent additional incentive $I_v = VoI_c$ $\Delta L_R/L_R$ for collecting information. As VoI_c is a function of belief P, so is I_v .

Fig. 7(c,d) shows this correction, starting from the example of graphs (a-b). Functions c_π^* and c_ω^* represent the prior and posterior optimal expected cost, as evaluated by society, considering $C_F = 10C_R$, as $\bar{P} = 0.1$. We assume $P_F = 0.65$, $\mu_F = 0.04$, $P_R = 0.35$, so $r_{FR} = 7.43\%$, $\Delta L_F = 4.02L_R$ and $\Delta L_R = 0.30L_R$. In this condition, society imposes a penalty corresponding to about four times the repair cost, for doing nothing when P equal to \bar{P} , and offers 30% of the repair cost if the agent decides to repair. Adjusted expected loss l_π^ν and l_ω^ν are proportional to

functions c_π^* and c_ω^* , respectively. The lower graph reports VoI_c , that is maximum at \bar{P} , and $VoI^v = 0.7 \ VoI_c$. By providing incentive $I_v = 0.3 \ VoI_c$, society can make the adjusted VoI as assessed by the agent consistent with VoI_c . Properly, this regulation is not equilibrated: on one hand, the incentive for information collection is a cost for society, on the other hand, society gets a benefit from the better induced attitude towards information. This consideration could be embedded info a different (equilibrated) selection of incentive and penalty values, after some assumptions on information availability.

6. Conclusion

We have investigated the effect of societal constraints on agents' attitudes toward information. While these constraints are effective in forcing agents to take decisions consistent with society's will, they can have unwanted "second-order" effects on information collection, if this activity is controlled by these agents and unconstrained. Risk-neutral agents will collect information only if its cost is below its value. However, the VoI assessed by agents whose preferences are not aligned with society will differ from that assessed by society itself. In the illustrated example, two undesirable outcomes can occur. First, an agent forced by the constraint to repair an asset can be willing to pay too much for escaping the constraint by collecting information, up to, in the limit case, the repair cost for receiving almost irrelevant information. So, paradoxically enough, the very availability of information makes things worse for society, as its members waste resources. Second, the economic effect can be of the same magnitude (but much higher in relative terms) when the constraint is currently inactive, because the asset is judged to be safe enough. In this latter case, the agent can prefer to avoid information and, again in the limit case, pay up to the repair

To overcome these undesirable induced behaviors, society can provide economic incentives and penalties to agents, to make the loss function concave for the decision makers. The calibration of such measures to make constraints redundant, and even economically self-equilibrated, would depend on specific assumptions about agents' preferences, and we have provided an analysis of a simple problem. Society can also provide incentives for information collection, as illustrated above. Also, societal codes can require collecting data (e.g.,

[1]), and they could even prescribe to evaluate VoI according to a given formula, encoding the assessment from the societal standpoint, and force agents to buy information when its cost is below that threshold, but the implementation of such a requirement would likely be controversial.

The analysis presented in this paper can be relevant for the legal distinction between negligence and reckleness, and for the attribution of liability. Apart from proposing practical solutions for overcoming the problems of information avoidance and over-evaluation, the first goal of this paper was to clarify how mathematical properties of VoI are affected by external epistemic constraints, as this is key for many engineering applications, e.g. for the optimization of information collection.

CReditT author statement

CReditT for manuscript "Information avoidance and overvaluation under epistemic constraints: principles and implications for regulatory policies", by Matteo Pozzi, Carl Malings & Andreea Minca, submitted to RESS.

Matteo Pozzi: Conceptualization, Methodology, Writing; Carl Malings: Conceptualization, Methodology; Andreea Minca: Conceptualization, Methodology.

Declaration of Competing Interest

Conflict of Interest and Authorship Conformation Form, for manuscript "Information avoidance and overvaluation under epistemic constraints: principles and implications for regulatory policies", by Matteo Pozzi, Carl Malings & Andreea Minca, submitted to RESS.

Acknowledgements

This research is based upon work supported by the National Science Foundation under Grant no. 1638327. Also, the authors acknowledge Silvia Saccardo and MS student Ao Lao for useful insights and conversations

Appendix. Effects of almost-irrelevant information

While, trivially, even under external epistemic constraints, an independent observation with no effect on the agent's belief has no value at all, the possible value of slightly dependent information strongly depends on those constraints.

By "almost irrelevant information", we refer to a source of information that leaves the posterior belief almost identical to the prior one. Quantitatively, we can define a ball or radius ε , centered on the prior belief \mathbf{b}_{π} , so that the posterior belief \mathbf{b}_{ω} is always inside that ball, for any realized measure. We analyze the case of "almost irrelevant information" by computing the limit for ε going to zero. Alternatively, we can define the trace \mathbf{t}_{ω} of the covariance matrix $\mathbf{\Sigma}_{\omega}$ of the posterior belief \mathbf{b}_{ω} , and compute that limit for \mathbf{t}_{ω} going to zero.

Almost irrelevant information in unconstrained decision making

We formally prove that, without any external epistemic constraint, information that have almost-negligible effects on the agent's belief have also almost-negligible VoI. Let us consider the case when the optimal expected loss function l^* is a quadratic form:

$$l^*(\mathbf{b}) = \delta^{\mathrm{T}} \mathbf{Q} \delta + \mathbf{v}_{\pi}^{\mathrm{T}} \delta + q_{\pi}$$
(A.1)

where $\delta = \mathbf{b} - \mathbf{b}_{\pi}$, matrix \mathbf{Q} is negative definite, as l^* is concave, $q_{\pi} = l_{\pi}^*$ and \mathbf{v}_{π} is λ_{a_*} , where a_* is the optimal action for belief \mathbf{b}_{π} . As l^* has to be piece-wise linear, function l^* can have the form of Eq. (A.1) only for an infinite number of available actions. The posterior expected loss can be written as a function of the covariance matrix Σ_{ω} :

$$l_{\alpha}^* = \mathbb{E}_{\omega}[l^*(\mathbf{b}_{\omega})] = \mathbb{E}_{\omega}[\delta^{\mathsf{T}}\mathbf{Q}\delta + \mathbf{v}_{\pi}^{\mathsf{T}}\delta + q_{\pi}] = \mathbb{E}_{\omega}[\delta^{\mathsf{T}}\mathbf{Q}\delta] + \mathbf{v}_{\pi}^{\mathsf{T}}\mathbb{E}_{\omega}\delta + l_{\pi}^* = \operatorname{tr}(\mathbf{Q}\Sigma_{\omega}) + l_{\pi}^*$$
(A.2)

as $\mathbb{E}_{\omega}\delta = 0$ because $\mathbb{E}_{\omega}\mathbf{b}_{\omega} = \mathbf{b}_{\pi}$, and "tr" is the trace operator. So the VoI is:

$$VoI = l_{\pi}^* - l_{\omega}^* = tr(-\mathbf{Q}\Sigma_{\omega}) \tag{A.3}$$

Let us suppose we can identify a parameter of standard deviation σ_{ω} , so that the covariance matrix is proportional to the corresponding variance: $\Sigma_{\omega} = \sigma_{\omega}^2 \mathbf{M}$, and trace t_{ω} is proportional to σ_{ω}^2 . From Eq.A.3, the VoI is proportional to that variance: $VoI = \sigma_{\omega}^2 \eta$ with $\eta = \text{tr}(-\mathbf{Q} \mathbf{M})$: this shows that

function $Vol(\sigma_{\omega})$ is continuous at σ_{ω} equal to zero, where it is zero. Hence, for any monetary value, arbitrarily small, we can find a corresponding value of σ_{ω} so that the Vol is equal to that small value.

Now, generally the expected loss function l^* cannot be expressed by a quadratic form as in Eq. (A.1). However, suppose we can find a quadratic form that is never above l^* in the reachable belief's domain Ω_Y , and it is equal to l^* at prior belief \mathbf{b}_{π} . In this setting, the previous formula is an upper-bound for the VoI:

$$VoI \le tr(-\mathbf{Q} \ \Sigma_{\omega}) = \sigma_{\omega}^2 \ \eta \tag{A.4}$$

So we have proven, again, that the VoI cannot be high, when σ_{ω} is small. Actually, the quadratic bound is ineffective for some points in the beliefs domain: the kinks where the curvature is infinite. For these points, we can define a piece-wise linear lower-bound for l_{ω}^* .

Almost irrelevant information under external epistemic constraints

Under external epistemic constraints, function l^+ is not necessarily continuous. We now prove that if there is a discontinuity of value $\delta > 0$ in l^+ , then almost irrelevant information can have any value in the $[-\delta, \delta]$ interval. To do so, we focus on a simple one-dimensional case of the belief domain, without losing generality as, in a higher dimensional space, the argument can be repeated for a segment passing through the discontinuity. Let us consider this simple loss function l^* :

$$l^{+}(b) = \begin{cases} 0 & b \le 0 \\ \delta & b > 0 \end{cases} \tag{A.5}$$

where b indicates the belief. Let us consider a prior belief $b_{\pi}=0$, so that $l_{\pi}^*=0$, and a posterior one that can assume only two values:

$$b_{\omega} = \begin{cases} \varepsilon & \text{w. p. } P_{+} = 1 - \sqrt{\varepsilon} \\ \varepsilon - \sqrt{\varepsilon} & \text{w. p. } P_{-} = \sqrt{\varepsilon} \end{cases}$$
(A.6)

where ε is a small positive number. We can easily check that $\mathbb{E}[b_{\omega}] = 0 = b_{\pi}$. and that

$$\lim_{\varepsilon \to 0} l_{\omega}^{+} = \lim_{\varepsilon \to 0} \mathbb{E}_{\omega} [l^{+}(b_{\omega})] = \delta \lim_{\varepsilon \to 0} P_{+} = \delta \tag{A.7}$$

So the VoI can be as large as δ . By slightly changing the condition in A.5, so that l^+ is zero at b equal to zero, and reversing the sign of variable b_{ω} , we can also prove that the VoI can be as small as $-\delta$.

References

- (1) < number > 2008City of Los Angeles Building Code Volume 1 (Based on the 2007 CBC and the 2006 IBC) first printing ISBN-13: 978-1-58001-644-5 publication date: jan. 2008.
- [2] Bertschinger, N., Wolpert, D.H., Olbrich, E. and Jost, J., "Value of information in non co-operative games," arXiv preprint arXiv:1401.0001, 2013.
- [3] Black HC, Nolan JR, Connolly MJ, Nolan-Haley JM. Black's law dictionary: definitions of the terms and phrases of American and English jurisprudence, ancient and modern. St. Paul, MN: West Publishing Company; 1990.
- [4] Bratvold RB, Bickel JE, Lohne HP. Value of information in the oil and gas industry: past, present, and future. SPE Reservoir Eval Eng 2009;12(4):630–8.
- [5] Cover TM, Thomas JA. Elements of information theory. John Wiley & Sons; 2006.
- [6] Fischer K, Viljoen C, Köhler J, Faber MH. Optimal and acceptable reliabilities for structural design. Struct Saf 2019;76:149–61. 2019.
- [7] Golman R, Hagmann D, Loewenstein G. Information avoidance. J Econ Lit 2017;55(1):96–135. pg.
- [8] Goulet JA, Der Kiureghian A, Li B. Pre-posterior optimization of sequence of measurement and intervention actions under structural reliability constraint. Struct Saf 2015;52:1–9.
- [9] Jensen JL. Sur les fonctions convexes et les inégalités entre les valeurs moyennes. Acta Mathematica 1906;30:175–93. https://doi.org/10.1007/BF02418571. 1906.
- [10] Karlsson N, Loewenstein G, Seppi D. The ostrich effect: selective attention to information. J Risk Uncertain 2009;38(2):95–115. 2009.
- [11] Malings C, Pozzi M. Value of information for spatially distributed systems: application to sensor placement. Rel Eng Sys Safety 2016;154:219–33. 2016.
- [12] Memarzadeh M, Pozzi M. Value of information in sequential decision making: component inspection, permanent monitoring and system-level scheduling. Rel Eng Svs Safety 2016;154:137–51. 2016.
- [13] Pozzi M, Der Kiureghian A. Assessing the value of information for long-term structural health monitoring. Health Monitor Struct Biol Syst 2011 2011.

- [14] Pozzi M, Malings C, Minca AC. "Negative value of information in systems' maintenance," Proc. of the 11th International Conference on Structural Safety and Reliability (ICOSSAR2017). 2017. 6-10 August 2017.
- [15] Qin J, Thöns S, Faber MH. On the value of SHM in the context of service life integrity management. The 12th International Conference on Applications of Statistics and Probability in Civil Engineering. 2015. p. 1–8.
- [16] Raiffa H, Schlaifer R. Applied statistical decision theory. Cambridge, Massachusetts, USA: Harvard University Press; 1961. p. 1961.
- [17] Rayner S. Uncomfortable knowledge: the social construction of ignorance in science and environmental policy discourses. Econ Soc 2012;41(1):107–25.
- [18] Russell S, Norvig P. Artificial intelligence: a modern approach. Pearson Education; 2010. p. 2010.
- [19] Srinivasan R, Parlikad AK. "Value of condition monitoring in infrastructure maintenance. Comput Ind Eng 2013;66(2):233–41. 2013.
- [20] Straub D. Value of information analysis with structural reliability methods. Struct Saf 2014;49:75–85.
- [21] Sweeny K, Melnyk D, Miller W, Shepperd JA. Information avoidance: who, what, when, and why. Rev Gen Psychol 2010;14(4):340. 2010.
- [22] Thöns S. "On the value of monitoring information for the structural integrity and risk management. Comp-Aid Civil Infrastruct Eng 2018;33.1(2018):79–94.
- [23] Von Neumann J, Morgenstern O. Theory of games and economic behavior.
 Princeton University Press; 1944. p. 1944.
 [24] Zonta D, Glisic B, Adriaenssens S. "Value of information: impact of monitoring on
- [24] Zonta D, Glisic B, Adriaenssens S. "Value of information: impact of monitoring on decision-making. Struct Control Hlth 2014;21(7):1043–56.
- [25] Tonelli D, Verzobio A, Bolognani D, Cappello C, Glisic B, Zonta D, et al. The conditional value of information of SHM: what if the manager is not the owner? Health Monitor Struct Biol Syst XII 2018;10600:106002. International Society for Optics and Photonics.
- [26] Verzobio A, Bolognani D, Zonta D, Quigley J. Quantifying the Benefit of Structural Health Monitoring: Can the Value of Information be Negative? Structural Health Monitoring 2019.