Safety, Reliability, Risk, Resilience and Sustainability of Structures and Infrastructure
12th Int. Conf. on Structural Safety and Reliability, Vienna, Austria, 6–10 August 2017
Christian Bucher, Bruce R. Ellingwood, Dan M. Frangopol (Editors)
©2017 TU-Verlag Vienna, ISBN 978-3-903024-28-1

Negative value of information in systems' maintenance

Matteo Pozzi[a], Carl Malings[a] and Andreea C. Minca[b]

[a]Department of Civil and Environmental Engineering, Carnegie Mellon University
[b]Department of Operations Research and Information Engineering, Cornell University

Abstract: The value of information (VoI) provides a rational metric to assess the impact of data in decision processes, including maintenance of engineering systems. According to the principle that "information never hurts", VoI is guaranteed to be non-negative when a single agent aims at minimizing an expected cost. However, 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. A similar result may occur for managers of infrastructure assets following the prescriptions of codes and regulations. Modern codes require the probability of some failure events be below a threshold, so managers are forced to retrofit assets if that probability is too high. If the economic incentive of those agents disagrees with the code requirements, the VoI associated with tests or inspections may be negative. In this paper, we investigate under what circumstance this happens, and how severe the effects of this issue can be.

1 Introduction

Design, operation and maintenance of structures and infrastructure components can be formulated as a decision making process, under uncertainty on hazard, demands, capacity and long-term evolution. Managers and stakeholders (whom we will hereafter refer to as "agents") can take these decisions with the aim of optimizing their own revenues, 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 policies to affect or even control them. Specifically, agents may be prone to accept risks higher than that which the society can tolerate, possibly because they do not include all costs relevant for society in their analysis. To prevent agents making excessively risky decisions, 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 policy that it considers optimal, balancing costs for construction, maintenance and renovation with risks related to failures and malfunctioning.

However, agents also take decisions about the collection of information, e.g. using sensors and inspectors, and they allocate economic resources to these activities. In this paper, we investigate the effect that the constraints on decisions have to the information collection. We assume that society does not impose any direct constraint on information collection, and agents are free to select information by evaluating its cost and benefit. However, as information is useful as far as it guides the management process, that evaluation is influenced by society's constraint. Specifically, the agent may find it convenient to avoid information, even when it is free, in order to escape a constraint that society imposes. In Section 2, we illustrate why the Value of Information (VoI) assessed by agents can be negative, and in Section 3 we describe in detail a setting in which this happens.

2 Positive and negative Value of Information

VoI analysis provides a rational metric for assessing the impact of information [1-3]. While this analysis can be complicated in some problems [4-5], here we refer to the simple setting of rational agents taking a one-stage decision under uncertainty. An agent has to select action a in domain A, while the world the agent is interacting with is in state x in domain X. The loss the agent received is quantified by function L.

2.1 Why the VoI is (usually) non-negative

In this Section, we provide a short intuitive proof of the principle that "information can't hurt" [6], that is, that VoI is non-negative. Figure 1 illustrates this proof, and follows the notation of probabilistic graphical models and decision graphs. Let us start with the case of perfect information, as in diagram (a). Loss function $L(a,x)$ depends on decision variable a and random variable x, described by distribution p_X. Without observing x, the agent selects action a^* to get expected loss $L^*(\emptyset) = \min_A \mathbb{E}_X L(a,x) = \mathbb{E}_X L(a^*,x)$, where $\mathbb{E}_Z[f(z)]$ indicates the expectation of function f according to distribution p_Z of z. Observing the world's state x in advance, the agent can calibrate the action depending on the observation, getting $L^*(X) = \mathbb{E}_X \min_A L(a,x)$. The expected value of perfect information is $\text{EVPI}(X) = L^*(\emptyset) - L^*(X)$ and defines the expected reduction of loss due to that observation. We can define the regret taking action a^* as $R(x) = L(a^*,x) - \min_A L(a,x) \geq 0$. As the regret is always non-negative, so is the EVPI: $EVPI(X) = \mathbb{E}_X R(x) \geq 0$.

The expected loss when indirect measure y of the world's state, defined by conditional probability $p_{Y|X}$, is observed in advance is $L^*(Y) = \mathbb{E}_Y \min_A \mathbb{E}_{X|y} L(a,x) = \mathbb{E}_Y \min_A L'(a,y)$, where we define $L'(a,y)$ as $\mathbb{E}_{X|y} L(a,x)$. We can re-define the prior loss as $L^*(\emptyset) = \mathbb{E}_X L(a^*,x) = \mathbb{E}_{XY} L(a^*,x) = \mathbb{E}_Y \mathbb{E}_{X|y} L(a^*,x) = \mathbb{E}_Y L'(a^*,y)$. So the VoI of observing y is $VoI(Y) = L^*(\emptyset) - L^*(Y) = \mathbb{E}_Y R(y)$, where regret $R(y) = L'(a^*,y) - \min_A L'(a,y) \geq 0$ is non-negative, and so is the VoI. Diagrams (b-d) show how to transform an indirect observation into a direct one, using Bayes' rule (b-c) and eliminating intermediate variable x (c-d). These steps are encoded in the definition of $L^*(Y)$.

We conclude that the agent should always accept free information, as it "can't hurt", even when it is indirectly (and even loosely) related to the world.

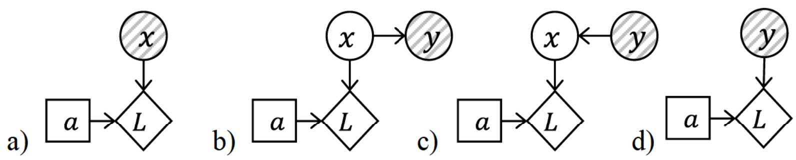

Figure 1: probabilistic graphical models for the VoI: perfect information (a), imperfect information (b), transformed to equivalent perfect information (c-d).

2.2 and why, sometimes, it is actually negative

However, experience suggests that sometimes it is better to neglect or refuse information. First, we can that argue that "too much information can harm". When variable Y is independent of X (so that $\mathbb{E}_{X|y}[\cdot] = \mathbb{E}_X[\cdot]$), the VoI of Y is nil. Now, if the observation is free, the agents should have nothing against collecting it. However, no information is actually free, when considering all costs related for collecting and processing it, and so they should neglect irrelevant information.

Second, an agent can take a decision, then revise it based on noisy measures, while in reality the prior decision was actually correct and the information has misled her. However, we point out that non-negativity of VoI is a result holding in the expected sense while, in a given realization (say of X and Y), it may not hold true for the corresponding realized quantities. We can also argue that if the processing model is incorrect, then the impact of information can be detrimental. For example, consider an agent over-confident in the precision of a sensor, or unaware of its systematic bias, and so assuming an incorrect conditional probability $p_{Y|X}$. 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 perceived uncertainty in the relation between the world and measures.

Our investigation refers to a third class of problems when the VoI can be negative. VoI analysis assesses the impact of information; now, who is the information revealed to? To the agent herself, first of all, who adapts her policy depending on the outcome. But, in many circumstances, when the information is revealed, it is also made available to other agents. Those agents can modify the environment where the decision-maker acts, and influence her loss. When including this indirect effect, the overall impact of information can be negative. As an example, suppose an entrepreneur shares a market with a competitor: a piece of information may be irrelevant for her, but key for her competitor, who can improve his strategy and reduce her share of the market [7]. In that case, the impact of information is clearly negative, when assessed by the agent.

2.3 Value of information when acting under an external constraint

In the analysis of Section 2.1 we assumed that the environment surrounding the agent, defined by the available actions and by the loss function, is not affected by the measures collected. To model the cases suggested at the end of previous Section, we now remove that assumption. Our motivating case is that of a manager of an infrastructure asset following a regulation, e.g., a building code. The code embodies a public policy limiting risk to an acceptable level. Depending on the knowledge of world state X, the code poses a constraint, and allows only for some actions to be taken. To model this, we define $\mathcal{A} \subseteq A$ as the subset of actions available, as a function of the probability distribution of X: either prior probability p_X or posterior probability $p_{X|y}$. Prior optimal loss is now $L_S^*(\emptyset) = \min_{a\in\mathcal{A}[p_X]} \mathbb{E}_X L(a,x)$, and posterior expected loss is $L_S^*(Y) = \mathbb{E}_Y \min_{a\in\mathcal{A}[p_{X|y}]} \mathbb{E}_{X|y} L(a,x)$, so that VoI is now $VoI_S(Y) = L_S^*(\emptyset) - L_S^*(Y)$. We have no guarantee that this quantity is non-negative. In the next Section we investigate how this value behaves in an example.

3 VoI in systems' maintenance under an external constraint

3.1 How external constraints can make the value of information negative

We consider an agent managing an asset under uncertainty about its state. The world's state X is binary, 0 indicating a healthy asset and 1 a damaged one. p_X can be completely defined by the probability P_D of the damaged condition. Action set A is also binary, with 0 indicating doing nothing and 1 retrofitting. Cost for retrofitting is summarized in loss value L_R, while the cost for failure, occurring when a damaged asset is not retrofitted, is given by loss value L_F. The loss function can be expressed as $L(x, a) = a\, L_R + x(1-a)L_F$. Such a function indicates that failure can be avoided by retrofitting.

In this setting, an agent not constrained by the code would retrofit if and only if P_D surpasses $P_L = L_R/L_F$. However, let the agent be constrained by regulation prescribing to retrofit if P_D exceeds threshold value P_T as, above that value, society considers the asset too risky to be in operation: hence while set $\mathcal{A}$ usually includes both actions, it is restricted to $\mathcal{A} = \{1\}$ if $P_D \geq P_T$. So the constraint is inactive if $P_L \leq P_T$ as, in that case, the agent's policy is more conservative than the code. If, on the other hand, $P_L > P_T$, the regulation can force the agent to retrofit when, if left free, she would have preferred to accept the risk of doing nothing.

We consider an available "inspection" of the asset, modeled by noisy binary observation Y, where 0 indicates "silence" and 1 "alarm", whose value may differ from X with probability ε. Hence, for inaccuracy ε equal to zero the agent can observe X directly, while for ε equal to 50% the observation is independent of the actual asset state. The relationship between P_D and the VoI (normalized by L_R) is plotted in Figure 2(a) for L_F equal to a thousand times L_R, and 6 values of ε: 0, 10, 20, 30, 40 and 50%. Without constraints, VoI is a piecewise linear function, made by increasing linear function $\alpha = (\Delta L\, \zeta + L_R\varepsilon)P_D - L_R\varepsilon$ and decreasing linear function $\beta = -(\Delta L\, \varepsilon + L_R\zeta)P_D + L_R\zeta$, where measure accuracy ζ is $(1-\varepsilon)$ and $\Delta L = L_F - L_R$. The lines meet at P_L. For P_D ranging from 0 to 1, the VoI is nil up to when α becomes positive, it follows α up to P_L, then it follows β down to zero, and then it stays at zero. Piecewise linearity is masked, in graphs (a-b), by the log-scale for the horizontal axis.

Under the constraint imposed by the code, the VoI looks as in Figure 2(b), where we select a specific value for the inspection's inaccuracy among those plotted in (a), i.e., $\varepsilon = 30\%$, and threshold $P_T = 0.03\%$. VoI jumps to function α at a value approximately equal to $P_T\varepsilon/\zeta$, then to function β at P_T and to zero approximately at $P_T\zeta/\varepsilon$. Hence, the constraint highly influences the VoI and it can be negative, for P_D smaller than P_T.

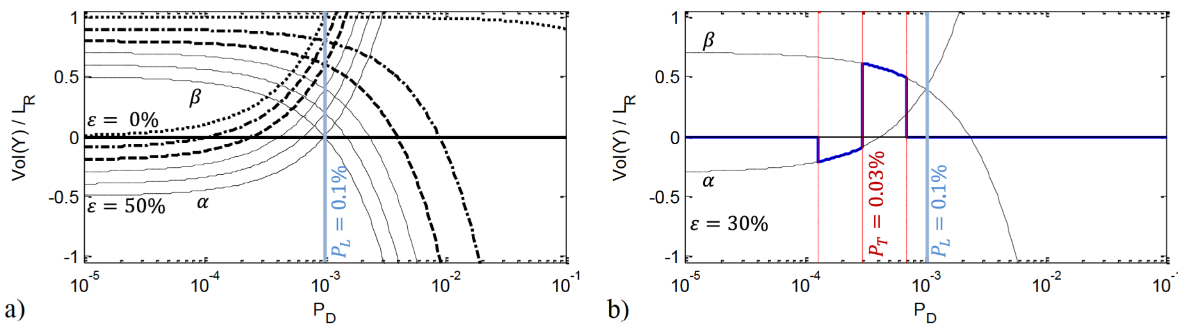

Figure 2: Construction lines for the VoI as a function of the probability of damage, for different values of inaccuracy (a), specific VoI for a set of parameter values (b).

3.2 Agent behaviour and societal expected cost depending on the constraint

In the prior condition, the agent will do nothing only if P_D is below both P_T and P_L; otherwise she has to retrofit, either because the perceived risk is too high (when P_D is above P_L) or because the code imposes it (when P_D is above P_T). P_A, P_S, $P_{D|A}$ and $P_{D|S}$ indicates the marginal probability of the alarm observation, of the silence one, and the conditional probability of damage given these outcomes, respectively.

The agent assesses the VoI, as shown in previous Section, and compares it with observation cost L_I, so the binary decision variable I is $I[L_I \leq VoI(Y)]$, where $I[\cdot]$ is the indicator function: I is 1 when the agent decides to inspect and 0 otherwise.

We assume that society models the process using the same framework the agent uses, only with different values. Societal costs for failure, retrofit and inspection are C_F, C_R and C_I respectively. In the prior condition, clearly the optimal value for threshold P_T is C_R/C_F, as that forces agent's action to be consistent with societal utility. Including the possibility of inspection, the expected societal cost is $\mathbb{E}[C|P_T] = C(\emptyset)(1-I) + [C(Y) + C_I]I$, where $\mathrm{C}(\emptyset)$ and $C(Y)$ are defined as $L(\emptyset)$ and $L(\mathrm{Y})$, but using societal costs. We also assume that the agent and society agree on all probabilities and costs except that modeling the effect of failure, so C_R and C_I are identical to L_R and L_I respectively.

We normalized the cost for society to that of retrofitting, defining $r_{C/R} = C_F/C_R = 100$, $r_{I/R} = C_I/C_R = 50\%$, $r_{L/C} = L_F/C_F = 10\%$, and a probability of damage as $P_D = 2\%$. We start considering inaccuracy $\varepsilon = 5\%$, so that the updated probability given alarm and silence are $P_{D|A} = 27.94\%$ and $P_{D|S} = 0.107\%$ respectively, while the marginal probability of alarm is $P_A = 6.8\%$. For the agent, the maximum risk before retrofitting is $P_L = (r_{L/C} r_{C/R})^{-1} = 10\%$ so, after an alarm, she will find retrofitting convenient. As shown in Figure 3(a), when P_T is above P_D, the constraint is inactive: the agent does nothing in the prior condition, and the VoI is rather low, as shown in graph (b). When P_T is below $P_{D|S}$ the agent is forced to retrofit no matter what, and the VoI is nil. For P_T between $P_{D|S}$ and P_D, the agent is forced to retrofit, but she can hope to receive silence and avoid retrofitting, and the VoI is much higher than that without the constraint. As the cost C_I is half of C_R, the agent will inspect only in that range of thresholds, as shown in graph (c). In this problem, the VoI assessed by society is also above C_I. Graph (d) reports the expected societal cost as a function of P_T. By assigning a threshold P_T between $P_{D|S}$ and P_D, agent's behavior corresponds to the optimal one for society: to inspect, and then to retrofit only following an alarm.

Results are different if we consider a less accurate inspector, with $\varepsilon = 20\%$, as we show in graphs (e-h). Now $P_{D|A}$ and $P_{D|S}$ are 7.55% and 0.51% respectively, while P_A is 21.2%. As $P_{D|S}$ is less than P_L, without constraint the agent would never retrofit, and VoI would be nil. For P_T between P_D and $P_{D|A}$ the agent can avoid to retrofit in the prior setting, but an alarm would force her to retrofit, while she would prefer not to, hence the VoI is negative. As in the previous case, the VoI is relatively high for P_T between $P_{D|S}$ and P_D, but in this case the VoI assessed by society is actually lower than C_I. In this case, the optimal constraint would be below $P_{D|S}$, as the best action for society is to retrofit without inspecting.

This analysis shows that no value of $\mathrm{P_T}$ can guarantee the optimal cost for society: as society cannot control the decision of inspection directly, the agent's behavior, which is optimal for minimizing agent's loss L, is not optimal for minimizing societal cost C.

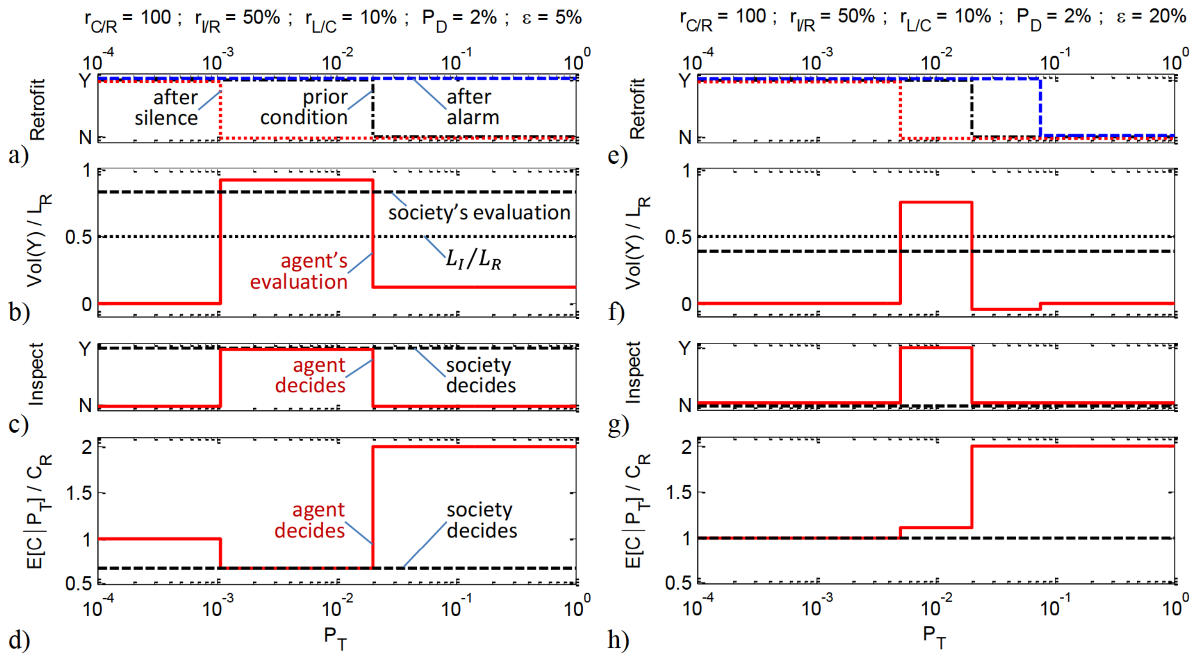

Figure 3: Decisions on retrofit (a), VoI analysis (b), decisions on inspection (c) and expected societal cost (d) vs threshold on acceptable probability of damage when inaccuracy is 5% and corresponding graphs when inaccuracy is 20% (e-h).

3.3 Bounds for suboptimal behaviour in information collection

In this Section, we investigate bounds to the inefficiency in information collection in a management process. We start considering the case when P_D is above P_T: society finds retrofitting convenient and imposes this prior decision to the agent. Of course, information has a value according to society: if this is above its cost, society would like to recommend that it be collected. In this case, the agent will always agree with society when the latter suggests collecting information, as the VoI assessed by an agent is greater than that assessed by society, for P_T/P_L less than one. However, exactly for this reason, the agent may be willing to acquire a piece of information that is too expensive from the societal standpoint.

We focus on limit cases. Let us consider an agent for whom L_F is identical to C_R, so that the agent finds retrofitting convenient only when damage is certain (i.e., $r_{L/C} = P_T$). We assume P_D is slightly above P_T, so the agent is forced to retrofit. However, an observation is available, with inaccuracy ε slightly below 50%. Clearly, this information is almost irrelevant for society, and the corresponding VoI is almost zero. However, the agent sees the collection of that information as a 50% chance of escaping the constraint, hoping in the silence outcome that would take the posterior probability of damage below P_T. The VoI, assessed by the agent, is almost $C_R(1-P_T)/2 \cong C_R/2$ and she is willing to pay up to such a high cost, if needed. Hence, the availability of information makes the overall cost grow by up to 50% in this setting. Figure 4(a) reports the costs for these limit cases, depending on P_D, for $P_T = 0.1\%$. When P_D is higher than P_T, the extreme case occurs when $P_{D|S}$ is equal to (actually, slightly less than) P_T, so that the agent can intended a silence outcome as a way for escaping the constraint. As shown in graph (c), the corresponding limit case inaccuracy, ε^*, decays from 50% down to zero. The corresponding probability of silence, P_S, grows, and so do value VoI^* and the corresponding expected cost $\mathbb{E}[C^*]$ in the limit case. However, VoI^* has to approach zero when P_D goes to one, and so the function (as well as P_S) decreases along the way.

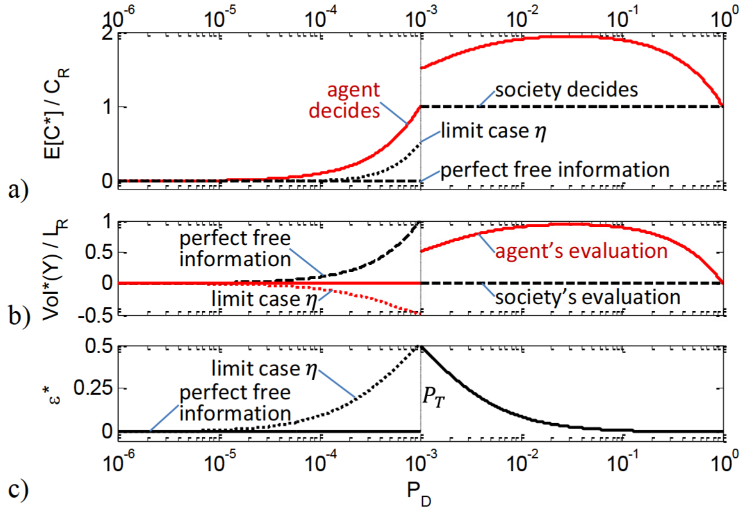

Figure 4: Expected societal cost when agent decides about the inspection and when society does (a), VoI analysis (b) and observation inaccuracy in the limit cases (c).

At its maximum, VoI^* cannot be higher than L_R, because it makes no sense to invest in information more than in an action that would prevent any additional required cost. Consequently, $\mathbb{E}[C^*]$ cannot be higher than the double of C_R.

Let us now focus on the case when P_D is below P_T and, consequently, the prior action is to do nothing. One limit case is given by free perfect information. When L_F is C_R, the agent will assign no value to it, while society will assign $VoI^* = (C_F - C_R)P_D$, as in graph (b). $\mathbb{E}[C^*]$ is $C_F P_D$ when the agent is in control and doesn't collect information, and it would be $C_R P_D$ if society could control the decision on inspection, as in graph (a). Another limit case (that we call η) is related to negative VoI assessed by the agent. For example, when P_D is slightly less than P_T, the agent interprets an observation with ε slightly below 50% as an almost 50% chance of being forced to retrofit, and she will be willing to pay $C_R(1 - P_T)/2 \cong C_R/2$ to avoid the observation. Similarly, we define that limit case by imposing $P_{D|A}$ equal to P_T. The corresponding inaccuracy decays to zero when P_D decreases to zero, as plotted in graph (c). In the meanwhile, the limit-case VoI, as assessed by the agent, stays negative, but its magnitude also decays, as $P_{D|A}$ does, as in graph (b). The corresponding VoI assessed by society is zero. For evaluating $\mathbb{E}[C^*]$ in that case, we assume that such information is actually available at a negative cost equal to the VoI assessed by the agent, so that society can decide to take the revenues related to the information. Graph (a) also reports the cost related to that case, which is always higher than that when free perfect information is available.

4 Conclusions

We have investigated the effect of society imposing a constraint to agents taking decisions on risk management. While this constraint is effective in forcing agents to take decisions consistent with society's will, it can have unwanted second-order effects on information collection, if this activity is controlled by these agents and unconstrained. Risk-neutral agents will assess the value of available of information and collect only that whose cost is below that value. However, the VoI assessed by agents whose preferences are not aligned with society will differ from that assessed by society itself. We have restricted our analysis to a simple case, where agents have to decide to retrofit an asset or not, and their loss function differs from that of society only because the latter assigns a higher cost to the asset failure. In that setting, two negative outcomes can derive. An agent forced by the constraint to retrofit can be willing to pay too much for escaping the constraint by collecting information, leading to a total expected cost that is bounded from above by the double of the retrofit cost. So, paradoxically enough, the very availability of information makes things worse for society. 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 even when it is free, or even if she can receive an economic reward for collecting data, despite the fact that it can reduce the expected societal management cost by several orders of magnitude.

Given the current structure of regulations, it is hard to overcome this inefficiently in collecting information. On the one hand, codes can require to collect data (e.g., [8]), 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. The implementation of such a requirement would likely be controversial. On the other hand, society could remove the constraint, and instead introduce incentives for aligning agents' preferences with societal ones. The scope of this paper, however, is not to propose solutions, but to highlight and understand the issue, illustrating why the value of information can be negative in some applications.

Acknowledgements

This research is based upon work supported by the National Science Foundation under Grant No. 1638327.

References

[1] H. Raiffa and R. Schlaifer, Applied Statistical Decision Theory. Cambridge, Massachusetts, USA: Harvard University Press, 1961.
[2] M. Pozzi and A. Der Kiureghian, "Computation of lifetime value of information for monitoring systems," in Proceedings of the 8th International Workshop on Structural Health Monitoring, Stanford University, California, USA, 2011.
[3] J.A. Goulet, A. Der Kiureghian, and B. Li, "Pre-posterior optimization of sequence of measurement and intervention actions under structural reliability constraint," Structural Safety, 52, pp.1-9, 2015.
[4] C. Malings and M. Pozzi, "Value of Information for Spatially Distributed Systems: application to sensor placement," Rel. Eng. & Sys. Safety, vol. 154, pp. 219–233, 2016.
[5] M. Memarzadeh and M. Pozzi, "Value of information in sequential decision making: Component inspection, permanent monitoring and system-level scheduling," Rel. Eng. & Sys. Safety, vol. 154, pp. 137–151, 2016.
[6] T.M. Cover and J.A. Thomas, Elements of information theory. John Wiley & Sons, 2012.
[7] N. Bertschinger, D.H. Wolpert, E. Olbrich and J. Jost, "Value of information in noncooperative games," arXiv preprint arXiv:1401.0001, 2013.
[8] 2008 City 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