



HAL
open science

An Architecture for the Evolution of Trust: Definition and Impact of the Necessary Dimensions of Opinion Making

Adrien Revault d'Allonnes

► **To cite this version:**

Adrien Revault d'Allonnes. An Architecture for the Evolution of Trust: Definition and Impact of the Necessary Dimensions of Opinion Making. Information Evaluation, Wiley, 2013, 9781848216594. hal-01559018

HAL Id: hal-01559018

<https://univ-paris8.hal.science/hal-01559018>

Submitted on 25 Oct 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Chapter 9

An architecture for the evolution of trust: definition and impact of the necessary dimensions of opinion making

9.1. Introduction

The evaluation of algorithm performance, computational techniques and methods is full of scores of trust.¹ These scores generally rate to the method rather than its output. That is, while it is acknowledged that automated processing yields uncertainty and imprecision, classic measurements usually characterise the way in which the information is constructed, rather than qualifying the degree of trust that should be invested in that piece of information. Although it is generally sufficient to know that an extraction algorithm is effective, say, eight times out of ten, or that its production is in line with expectations in the same proportions, once one starts dealing with sensitive data, at the root of potentially tragic decisions, one may wish to measure how much to trust the information item rather than a uniform doubt on its construction. Such a measurement should vary with the contents of the information itself, instead of on the basis of the tools used to produce it. Furthermore, if the elaboration method for this trust indicator is explained, legible or even adaptable, the user can then learn to grasp the system and use its indications in the rest of his intervention. To reach such a goal, it is necessary to introduce a distinction between *score*, the expressed degree of trust, and *scoring*, the method used to evaluate trust..

Chapter written by Adrien REVAULT D'ALLONNES.

¹ See Chapter 2 for a detailed analysis of the notion.

While information scoring, the way it is usually described, often focuses on evaluating the reality of the fact which it pertains to, we propose, in this chapter, to look at how much we can trust a piece of information, independently of the truth of the fact that it reports. Indeed, because it is just as possible to believe that the *Free Resistance of Dagbas* is responsible for the attacks as to believe the contrary, we propose to study the trust which can be put in the information. This way of looking at the issue raises the question of the construction of trust. Section 9.2 presents a perspective on trust, on the basis of which we build a model for both its establishment and evolution.

9.2. A perspective on trust

How does one form an opinion? What is the process by which one manages to be convinced of something, to make one's mind up? The process is simple and more or less universal: find reliable sources, cross-reference what they say and then extract the viewpoint which most closely corresponds to your convictions, knowledge and interests. Thence, new questions emerge. How is one to estimate the reliability of a source? How does one resolve eventual conflicts between reliable sources? How many confirmations does one need before being convinced? Can these certainties be revised?

With regard to information in general, many wise people tend to trust what is generally called "the media" and what, in intelligence terminology, are referred to as "open sources". On a daily basis, *compendia* of open sources are created by operators specializing in particular subjects. Journalists, representatives of the fourth power according to Tocqueville [TOC 1835], offer summaries and other press reviews, addressed to the impatient rest of the world. All of them apply the same process: first find the information and then verify it [BAU 02; SIM 05]. Making certain that it comes from a reliable and competent source is the first step. Finding other people, other points of view to support it comes later, in order to reinforce the initial trust [BOR 98]. Even scientists look for high quality references to position their own work [CHE 07].

Therefore, trust is a subject of major interest. In its broadly-accepted sense, it applies in equal measure to people ("Spontaneous or acquired belief in the moral, affective or professional value of another person, which makes it hard to imagine cheating, betrayal or incompetence on their part"), to oneself ("assuredness that one can have in one's own resources or in one's destiny"), as it does to an object or a

piece of information: “Credit, faith [...] attached to someone or something”.² Be it the question of trust between individuals or of the conviction that a piece of information is correct, all decisions integrate it, in one way or another. The consideration of sensitive subjects and the decisions which stem from it further increase the need for a reliable, comprehensible and revisable indicator of trust.

In this chapter we propose to give an unambiguous definition of the dimensions upon which trust is built and then describe the architecture of the combination of these dimensions. Finally, we show that the flexibility of the model enables it to be adapted to suit any evaluator. The aim of the proposed model is to provide its users with a bounded and interpretable view of the constitution of trust. The reader will, of course, understand that we do not mean to impose this model as the unequivocal truth behind this central point of human reasoning.

9.3. Dimensions of information scoring

Let us now present the dimensions we choose to perform information evaluation – an extension of our proposed system in [BES 08]. Our propositions originally stem from the existing criteria, presented in Tables 4.1 and 4.2. However, they break away from these existing criteria in various ways, which we shall discuss here. Among the usual criticisms made of the method for information evaluation such as it is officially defined³, the most recurring is the difficulty in interpreting and grasping the results [CHO 03; CHO 04; NIM 04]. We propose to resolve the first of these points by giving the user a unique numerical value expressing the confidence that he can have in a piece of information: its score. Giving a value on a graded scale enables comparisons between different pieces of information and to sort them by level of trust, a substantial improvement on the current approach.

Our objective of readability in terms of the score and the evaluation process is not limited to the way in which they are presented. We will, additionally, clarify the factors which we propose to use to evaluate information. In general, the quality a piece of information is measured by its conformity to a model whose quality is also guaranteed [BAT 06]. This search for quality therefore takes place in two stages. First, one produces a model capable of accurately representing all the expected information. One then checks each piece of information is detailed and not redundant. Obviously, there are many other criteria which make for quality, but the

2 N.B. The definitions proposed here are drawn from the *Trésor de la Langue Française Informatisé – ATILF* (computerized treasury of the French language) – a trustworthy source if ever there were one.

3 See Chapter 3 for more detail about the place of information evaluation in the military doctrinal corpus.

expressiveness of the model and the completeness of the information are unavoidable factors in this domain. Although they are essential for the evaluation of the knowledge contained in the model and the information, these criteria are not meant to evaluate whether the user can believe a particular piece of information or not. A model of excellent quality ensures the knowledge the information describes is complete, interpretable and exploitable. A quality recording in such a model provides a maximum of knowledge. We wish to add to these guarantees of quality an estimation of trust. In order to evaluate this trust, we propose to combine different types of dimensions, some of which bear witness to the process of creation of the information rather than its content.

The factors influencing the trust that we put in a piece of information are not limited to its content. Hence, the dimensions we use for information evaluation do not necessarily pertain to the fact it describes: although a piece of information is supposed to represent a real fact, it is not always possible to verify its realisation. The only way in which we can have faith in its veracity is by forming an opinion with the means available, such as its context and its mode of production, its conformity with our knowledge and expectations, or validation by other sources. For instance, in order to find out whether a person is tall, we not only need to measure their height using appropriate tools, but also to have an idea of the distribution of heights in the surrounding population. Also, if that information is reported to us, the first indicators we have upon which to base our belief relate to its source.

Because there is no absolute norm for evaluating how much belief we can invest in a piece of information, we propose to consider the principles governing the elaboration of trust, as mentioned above, in order to clarify the boundaries of the factors involved in information evaluation. The essential questions – although these could doubtless be enriched by others –we propose to use to evaluate the veracity of a piece of information are presented in Table 9.1.

Who is providing it?	What does he know about it?
How likely is it?	Has it been reported elsewhere?

Table 9.1. *Crucial questions influencing the estimation of trust in a piece of information*

The two questions on the upper row relate to the source of the information. They provide the first *a priori* indications about the trust which we can afford a piece of information. Even before someone gives us a piece of information, our trust in him affects the extent to which we are prepared to believe it. Note that this *a priori* indication is related to the source but does not depend on the information, as the trust one has in the source is independent of the subject, similarly to reliability in the current doctrines. After this initialisation, the second question relates to the source's qualification to express an opinion on the topic. While it still relates to the source, this capacity varies with what information is being transmitted. In an ideal universe, where the evaluation of the dimensions operates only on fact, then, unlike the previous factor, the evaluator's subjectivity should not influence the speaker's of knowledge. Indeed, the trust invested in the source varies from one listener to another, on the basis of the shared history of relations between them, whereas the speaker's abilities regarding a particular subject depend only on him. Any subjective doubts about the source's capabilities should be included in the prior evaluation of trust – at least in the model proposed here.

The remaining two questions relate to the content of the information, and are no longer connected to the source. It goes without saying that the first factor to be considered is how realistic the information looks. Note that it is not a question of verifying its reality – whether the fact described actually happened – but rather of checking its compatibility with the knowledge the evaluator has of the world. Hence, this depends solely on the studied piece of information and the frame of discernment of the user, i.e. the analyst's paradigm of interpretation. The final question, on the other hand, means to verify if the same information is known anywhere else. When dealing with a piece of information which seems unusual, i.e. if the answer to the previous question is nearly negative, and supposing that we cannot easily verify the truth of the assertion, the search for concomitant cross-references or, failing that, for other individuals who believe it, is constitutive of the establishment of faith and trust.

Note that these questions separate the dimensions relating to the source from those relating to the content of the information, in the same vein as the separation between the dimensions of absolute and relative quality of information [BER 99]. The dimensions we use are naturally divided into these categories, where reliability, which is independent of the information itself, is absolute, whereas the others are relative to the piece of information in question. To this distinction, we add another one, based on whether or not the dimension takes the context into account, i.e. if its influence is stable or varies from case to case. Indeed, the separation between questions relating to the source and to the informational content can be further clarified, if we take account of the scope of each criterion. The answers to the above four questions cross the two axes of dependency on the source or on the content of

the information, the *contextuality* or *generality* of the dimension. Thus, the four dimensions we propose are distributed in Table 9.2, where reliability deals with the first question, competence with the second, plausibility the third and credibility – a term used in the NATO doctrines [NAT 03] and discussed in Chapter 4 – with the last.

	General	Contextual
Source	Reliability	Competence
Content	Plausibility	Credibility

Table 9.2. *Object and scope of the dimensions of evaluation proposed here to clarify doctrinal factors*

While we have chosen this way of dividing the dimensions, there are other possible explanations for their independence and non-redundancy. For instance, we can cite the subjectivity of the evaluator: the reliability of the source and the plausibility of the information are two subjective dimensions. Quite evidently, the degree of trust afforded to a source depends on who is consulting that source. The plausibility of the information, envisaged as a degree of compatibility with one's existing knowledge of the world, also varies from one person to another. Note here that these restrictions on the subjectivity of the dimensions are questions of measure. Each of the dimensions presented here could be evaluated subjectively but, in order to clarify and standardise the use of the model, we propose to stipulate the methods to be used for their evaluation at the same time as describing their meaning.

The other two dimensions, on the other hand, should, ideally, not be subjective. If it is not known with absolute certainty, the real competence of the source does not depend on who is assessing it. Indeed, one can, for instance, circumvent one's doubts about a plumber's competence in electrics by refusing to pass judgment on an unknown factor, thereby avoiding favouring or penalising the final evaluation. Similarly, the way in which two pieces of information confirm one another does not depend on point of view, but rather on the evaluation method. These axes of separation between the proposed dimensions mean we can consider them not to be redundant, as they do not qualify the same objects and are therefore independent.

We represent the level of activation of each of these dimensions with a separate degree. In order to conform to the doctrinal line, we consider a scale with six levels, but this choice is not limiting. Obviously, these scales may contain more or fewer degrees depending on the intended application. There is no need to stick with the six degrees defined here, although it has been suggested that a scale with more than seven values loses in explicitness what it gains in precision [MIL 55].

A Totally reliable B Usually reliable C Fairly reliable D Rarely reliable E Unreliable F Reliability cannot be estimated	1 Expert 2 Competent 3 Partially competent 4 Insufficiently competent 5 Incompetent 6 Competence cannot be estimated
<i>Reliability</i>	<i>Competence</i>
1 Certain 2 Realistic 3 Possible 4 Not very possible 5 Impossible 6 Plausibility cannot be estimated	1 Totally confirmed 2 Partially confirmed 3 Insufficiently confirmed 4 Partially contradicted 5 Totally contradicted 6 Credibility cannot be estimated
<i>Plausibility</i>	<i>Credibility</i>

Table 9.3. *Examples of degrees for the dimensions*

In our discussion below, we detail these four dimensions for information evaluation, one by one. For each of them, we offer a definition and an interpretation of the discrete evaluation scale, all presented in Table 9.3. We also illustrate the meaning of each of them using the running example of this book.

9.4. General evaluation of the source: reliability

The quality of the source is an intrinsic quality of the information, according to the denominations given by [WAN 96] and [BER 99]. It is a crucially important criterion which usually serves to weight the interpretation of the information. Google's PageRank algorithm, for one, uses the level of trust accorded to the source (the Website) to sort the responses to a query, given an equal relevance [BRI 98].

We propose to distinguish, in this quality, that which is general from that which is contextual, and here we discuss its general aspect, which we call the *reliability of the source*. This depends neither on the question at hand nor on the already known information, and applies equally to all pieces of information supplied by one source. The reliability of the source therefore qualifies only the source, independently of the information.

It should be noted that this term is frequently used for a variety of notions: it simultaneously encompasses reputation or trust, as we conceive of it, the aptitude of the source to provide the information and any doubts that he has about the information itself [DIS 01; DEP 05]. These concepts of the criterion therefore mix its general components – relating only to the source – and its contextual components – relating to other factors. In the definition we propose, reliability depends only on the source rather than on the considered piece of information. It is, however, well known that reliability depends on the observer's position and, consequently, on who that observer and evaluator is. We shall see in the examples that while the reliability does not depend on the information, it is closely connected to the user.

We propose to express reliability on a scale with five levels, as indicated in Table 9.3, representing the degrees of reliability envisaged and their interpretations. In order to preserve a link with the original scale, we represent these degrees using letters. We have supplemented this scale with a value representing the impossibility of measuring the reliability: the degree F. Indeed, it is possible that the available knowledge about the source may not be enough to estimate his reliability. This occurs, for example, when we encounter a source who has hitherto been unknown to us. The estimated reliability of this contact can therefore be distinguished from that of an interlocutor who is known but whose reliability is medium, i.e. neither weighing in his favour or against him. Such a source will be assigned a level of reliability C. Note that this additional level appears on the scales proposed for each of the criteria considered here, as detailed below.

In addition, we must stress the dynamic nature of a source's reliability, which is bound to evolve over time as the source is consulted and acquires a reputation. Thus, we may learn more about a certain source and wish to re-evaluate his reliability. This updating poses the problem of the change in the evaluation of trust in the information. Two possibilities then arise: we can recalculate that trust or we can retain the current value as an indicator of the time when it was calculated. Of course, these issues only arise for pieces of information emitted since the reliability score changed. Consider, for instance, a source who we learn, at a date t_b , has suffered from memory problems since a date t_a (where $t_a < t_b$). All information provided

between t_a and t_b is affected by this change in reliability. However, the information emitted before t_a need not be re-evaluated.

9.4.1. Evaluation of reliability in the original scenario

The evaluation of the respective reliability of the sources in the example running through this book can be delegated to the user, in order to take into account the subjective aspect of this task and facilitate a personalised adaptation. Usbek's intelligence agencies may, *a priori* and independently of all information, assign a maximum reliability score (level A, completely reliable) to their infiltrated agent and a less high level (e.g. level D, usually not reliable) to the Ektimostanian blogger because of conflicting allegiances.⁴

9.5. Contextual evaluation of the source: competence

Having defined the general aspect of the source, his reliability, we wish to integrate into our process an axis which is often mistaken for reliability. We have seen that, such as we understand it, a source's reliability applies in exactly the same way to all pieces of information produced by that source. However, when considering the influence the source has on the trust we invest in the information that he produces, it is often tempting to include an estimation of his knowledge of the subject area. For instance, one should always take the philosophical musings of specialised scientists *cum grano salis*.

We propose to separate this dimension, which we can say is contextual, from the former concept of reliability. Indeed, the capacity of a source to provide a piece of information depends on the information itself. It is always possible that a source will provide a piece of information outside of his particular area of expertise. Yet a piece of information such as this should not be rejected out of hand for this reason alone; however, the receiver needs to take this into account when classifying the information or evaluating the trust to invest in it.

We call this factor the *competence of the source*, and propose to measure it with the scale given in Table 9.3. Here we see a level expressing the impossibility of evaluation. Indeed, it may be the case that we are unsure of the extent of a source's areas of expertise. For instance, it is not unheard of for certain IT specialists to have an unexpected cultural grounding in philosophy, and to be perfectly qualified to share their philosophical knowledge. Again, we propose to distinguish these unexpected experts from those of whose cultural background we are aware.

⁴ See the example in section 4.3 – the case of *d* and *e*.

9.5.1. Evaluation of competence in the original scenario

The competence of an infiltrated agent in the pay of Usbek about the situation of the rebel province should be evaluated equally by anyone who has all of the relevant information about the spy's mission. In reality, because he is supposed to keep track of internal manoeuvres and shifts within Ektimostan, his observation outside of his field of expertise could justifiably be estimated as "insufficiently competent", which is a score of 4 in terms of competence level. If the Ektimostan agencies, with the opposite objectives, had the same set of information, they would also attribute a competence of 4 to the enemy agent. An infiltrated agent who has not yet had any experience of an Ektimostanian blogger cannot evaluate that blogger's competence on the subject, so gives him a competence score of 6.

9.6. General content evaluation: plausibility

Having taken account of the factors about the source influencing the score, we now turn our attention to the information itself. Once more, we propose to separate the general dimension from another, more contextual. An essential criterion for belief in a piece of information is its *plausibility*. Even before considering details or confirmations, we reject any information which seems incompatible with our perception of the world. This initial evaluation is done on the basis only of the particular piece of information in question. The search for additional intelligence and confirmation takes place later.

Thus, in the intellectual process from which we draw inspiration, we would not investigate the possible reality of a piece of information which is devoid of all plausibility. This reflex for intellectual preservation is used, for example, by detractors of creationist movements, who propose a contradiction by the absurd, with their Flying Spaghetti Monster [HOR 07].

In order to evaluate the plausibility of a fact, we refer to its compatibility with our knowledge about the world. It is not a question of determining whether or not the fact is genuine, but rather of estimating the degree to which it is possible for it to happen. An example of evaluations for this criterion is presented in Table 9.3. Because it is possible that our knowledge will be insufficient to evaluate the plausibility of a piece of information, we have again chosen to propose a level 6, representing the impossibility of measurement. The major breakthroughs in scientific paradigms have certainly been produced by geniuses who have recognised situations which are not explicable by their present knowledge of the world.

9.6.1. Evaluation of plausibility in the original scenario

Thus, plausibility measures the match between the information and our knowledge of the state of the world. Suppose that the analyst dealing with the information from his spy in Ektimostan also knows that the groupuscules identified by his field agent are, in reality, groups of other agents charged with influencing public opinion and that they have received no orders regarding actions to be carried out, he will deem the participation of these factions in attacks to be “impossible” and the plausibility of the information reported will be evaluated as 5. An analyst who has no knowledge of the second plan of infiltration will, in view of the powder keg that the region represents, estimate that the existence and action of such a group are “realistic” and therefore attribute the information a plausibility of 2.

9.7. Contextual content evaluation: credibility

Once we are convinced of the source’s capabilities – first in general and then in respect of the information in question – and have verified that the information was not too incompatible with our own knowledge of the world, we generally proceed to verify the fact. When we learn something new, most of us seek to confirm it in order to ensure that it is credible. This phenomenon of confirmation seeking undoubtedly partly explains the success of search engines, constructed on the basis of the most frequent requests, and possibly also certain “buzz” phenomena. This verification is a search for correlation and corroboration of the as-yet-unknown information, preferably from reliable and competent sources. What we are looking for is the *credibility* of the information.

Because of the maximum degree of veracity on the existing information evaluation grid⁵, we choose to evaluate it as an indicator of confirmation by the pieces of information acquired. We must stress that the facts against which we compare the information being studied are themselves subjected to the same evaluation procedure. They are all drawn from more or less reliable, more or less competent sources, and their plausibility is undoubtedly comparable to that of the information being examined. With each confirmation or invalidation, the combination of the *a priori* degrees of trust (i.e. the rating resulting from the previous dimensions) weights the impact of the confirmation and finally enables us to establish whether or not to believe the information. In addition, the pieces of information against which we contrast that which we are evaluating also benefit from this update. This corresponds to the quality relating to homologous data, such as described by [BER 99]. However, credibility, as an indicator of confirmation, does not depend on the *a priori* rating. This rating relies on the impact of the

5 See section 4.4.2.1.

confirmation rather than on its measurement. The credibility will therefore be evaluated using measures of correspondence or conflict with the other pieces of information, and then integrated by combining them with their ratings.

Let us again highlight the distinction drawn between the comparisons between acquired pieces of information and their compatibility with the reasoning model. The pieces of information that we compare by evaluating the credibility are all constructed and therefore all have a changeable indicator of trust, unlike our knowledge of the world, whose level of trust is assumed to be fixed and maximal.

Table 9.3 presents the different levels on which we propose to evaluate credibility. Of particular note among these degrees of confirmation are partial confirmation and partial invalidation (levels 2 and 4). We have put forward a cumulative procedure for gaining confirmation, whereby it is not necessary to be confirmed solely by definite pieces of information in order to achieve maximum credibility [REV 07]. The accumulation of enough confirmations by less certain pieces of information yields the same result, although less quickly, as detailed in [REV 11]. However, that which is denoted by partial confirmation and invalidation relates more closely to the question of evaluation of the correlation between pieces of information. Anyone who has ever sought confirmation of a complex piece of news knows that it is rare to find complete and perfect confirmation, particularly if we insist on its coming from independent sources. It may also arise that we cannot find the link between two pieces of information although we suspect they relate to the same subject. This stems from the boundary maintained between co-occurrences, coincidences and confirmations, as anyone attempting to fix a system outside his own specialty well knows. For such cases, we again propose an inestimable, level 6, credibility.

9.7.1. Evaluation of credibility in the original scenario

Each of the pieces of information in the multi-source scenario therefore has an *a priori* rating. When looking for corroboration, we need to measure the degree to which two pieces of information confirm one another, e.g. the declarations of the Ektimostanian leaders, who confirm one another perfectly; those of Captain Ixil and the infiltrated field agent, which go in the same direction without absolutely echoing one another; or which contradict each other, as do those of Ixil and al-Adel.

9.8. Global expression of trust

On the basis of the four dimensions defined and described in the previous sections, our procedure outputs a degree of trust associated with the information in question: its rating. This rating expresses the combination of the above criteria, in a

unique and comparable score. We propose to express the rating as shown in Table 9.4.

1	Extremely likely
2	Likely
3	Possible
4	Doubtful
5	Unlikely
6	Confidence cannot be estimated

Table 9.4. *Final score: degree of trust associated with a piece of information*

It would certainly be possible to integrate other dimensions in the computation of the trust to be invested in a piece of information. For instance, we have chosen not to take account of the source's conviction. Were we be able to estimate this doubt, it goes without saying that it would be extremely important to include it in the evaluation.⁶ However, this would raise questions about the computation of the extent of the doubt, depending on the reliability and competence of a source.

In the same vein, we could envisage a factor reflecting the influence of word of mouth on the twisting of information. Besides the difficulty in machine detection and evaluation of the impact of such a factor, we think the doubts covered by this type of techniques would also be taken into account by the proposed model. If, on the other hand, it proves possible to reconstruct the original information, or to locate those responsible for the alterations [SZ 09], the contribution to the rating would be invaluable.

9.9. Architecture of information evaluation: characteristics

We have already mentioned the process which motivates our proposals for evaluation of the rating. Thus far, we have been presenting the general principle of construction of opinion, then extracting the axes along which we propose to evaluate the rating, in the previous section. The underlying process is crucial to our propositions regarding information evaluation. It structures not only the dimensions we use to rate information, but also the way in which we combine those dimensions, from the order of integration to its resulting sequentiality. Now we give a detailed

⁶ See Chapter 7 for an example of a semi-automated information evaluation method which makes use of this dimension.

description of this model which underlies our work and its implications, and then discuss some of its consequences, in section 9.11.

We consider the rating of the information as an indicator of the trust that can be invested in it. In order to establish a graspable and interpretable process of rating, we adopt a model for evaluating that trust. It is that model which we summarised above as the four questions applied to a piece of information:

- 1) Who is providing it?
- 2) What does he know about it?
- 3) Is it likely?
- 4) Has it been reported elsewhere?

Remember that we are not seeking to evaluate the certainty of the information: we are not interested in knowing whether the information is true, but rather in determining the extent to which we are convinced by it – be it precise, vague or uncertain. In order to achieve this result, we consider the process by which we are able to form an opinion about a hitherto unknown piece of information. Having established the definition of the factors which we integrate to the evaluation of the rating, we now turn our attention to the order in which they are taken into account, which derives from our model. In order to illustrate this process, consider what happens when we learn something new.

9.9.1. Order of integration of the dimensions

Whatever the information in which we are interested, the first *a priori* factor which influences our opinion is our trust in the organ presenting that information to us. In fact, when we look for a piece of information, we lend preference to sources which are certain, as indeed search engines do for us. Similarly, our opinion is already moulded by the person giving us the information – in this case, we shall suppose, it is a highly trusted friend.

It is only afterwards that we look at the competence of the source on the topic of the information, aided by our knowledge of his specialties, or the intersection of our tastes. If the information deals with emerging research in number theory, we can consider that our friend – an artistic photographer with not a trace of a scientific mind – is not the best source, but this should not put into question his usual reliability.

Whether it is reported as certain or – more likely here – in a vague and imprecise manner, we then consider whether the information does or does not contradict our own knowledge. For instance, the discovery of Fermat's own original proof of his last theorem – while it may be surprising – is not inconceivable.

Finally, to consolidate our opinion, we check if the information is given elsewhere. The search for confirmation follows a similar path, and the correlations which stem from this enable us to decide whether or not we believe the information.

Here let us insist on the process of integrating the dimensions into the rating. Our trust in our friend is independent, as already pointed out above, from what he tells us – the effect of which is projected onto everything he tells us. His competence on the topic in question, or at least our perception of that dimension, intervenes only later in the forming of our opinion. Also, the plausibility of the information transmitted and then its corroboration intervene in turn later on. From these observations about the forming of an opinion, notably borne out by the existing body of literature on intelligence [DIS 01] and on journalism [BOR 98], we deduce the ordering of the chain of information evaluation represented in Figure 9.1. It is from here as well that the notion of projection of the dimensions onto the information stems, because those dimensions are more part of our own convictions than of the information *per se*.




Figure 9.1. *Direction of the influence of the dimensions on trust: integration of the reliability of the source, his competence, the plausibility of the information and its credibility*

9.9.2. *Sequentiality of the information evaluation chain*

The information evaluation chain presented in Figure 9.1, which the rest of this chapter attempts to illuminate, thus formalises the elaboration of trust. Apart from the order in which the dimensions are integrated, the chain model also presents information evaluation as a sequential process. The acceptance of our model lends itself to this conception, and this choice offers greater interpretability for the rating. The illustrative examples given below highlight the advantage to a comprehensible process of opinion construction. In addition, the constant availability of a rating reflecting the advancement in the process helps to track, understand and validate the obtained result. Sequential calculations are no more costly, and therefore do not hinder the implementation of the model, and the evolution of the rating in this presentation of the process gives clear expression to the influence of the projection of each of the dimensions.

9.10. Architecture of information evaluation: a description

The architecture we propose for the information evaluation process is therefore a formalisation of the establishment of trust which we use it to represent. The information evaluation chain shown in Figure 9.1 gives an overview of its operation. This schema simultaneously represents the evaluation of the criteria and their impact on the evolution of the score.

The sequential process of information evaluation is divided into four stages. The running score, after integration of those dimensions which have already been evaluated, is represented by its current level: a number between 1 and 6 in a circle in different levels of grey for each stage. The complete scale is therefore present four times in Figure 9.1. The sensor represents the source, whose reliability may be activated at different levels. The arrows emanating from the sensor symbolise the impact of their level of activation on the initial estimation of the score, to which they directly lead. Thereafter, the arrows represent the levels of activation of the particular dimension in question and indicate its influence on the score. The intensity of activation of the criterion is symbolised by a variation in level of gray: the darker the arrow, the more intense the evaluation of the dimension. Light-coloured arrows, on the other hand, represent a minimum intensity.

9.10.1. Reminders about the evaluation of the dimensions

As we saw earlier on, the evaluation of each criterion, like the evaluation of the score, operates on a different scale but one whose construction is the same – i.e. the degrees of activation of the criteria range from 1 (“high”) to 5 (“low”), to which we add a degree of indecision 6 (“unquantifiable”). However, each evaluation has its own interpretation, as specified by Table 9.3.

This reminder is intended to draw the reader’s attention to two points: the particular behaviour when we encounter an “unquantifiable” score, on the one hand, and the multiple arrows representing the impact of taking account of the dimensions corresponding to those levels of activation, on the other. As a consequence of the constraints imposed on the evaluation, certain arrows become indistinguishable, as explained below.

We shall now detail the usage of the process of information evaluation and present the expected constraints on the influence of the integration of each dimension. In section 9.11, we shall come back to the possibilities opened up by playing with these influences.

9.10.2. Reliability of the source

In the model of trust evaluation and in the information evaluation chain, we propose that the first criterion to come into play should be the reliability of the source reporting a piece of information. We therefore use the reliability of the source as an initialisation of the information evaluation process.

Suppose we receive a new piece of information, which therefore has not yet been evaluated. Because this information is unknown, it is believed to the extent that we trust its source. Thus, the first step in the process attributes an unknown piece of information the level of trust (Table 9.3, greyscale in Figure 9.1) equivalent to the reliability of its source.

9.10.3. Competence and plausibility

The stage of integration of competence into our process has a negative impact on the level of trust. Indeed, as the initialisation of the score takes place at the level of the trust invested in the source, we would undoubtedly have supposed that the said source “knows what he is talking about”, or at least that we trust him to the level of that supposition. It is therefore natural that the only possible impact of competence is a negative one. If a completely trusted source is providing a piece of information outside of his domain of expertise, we are led to take that information with circumspection, whereas a source known for lies and manipulation will have difficulty in convincing us, even if he is talking about his specialty.

Similarly, when we take account of the plausibility of the information, it tends to lower the overall evaluation. If we believe a piece of information because of its source and if the source is competent on the subject in question, the fact that the information is possible should not reinforce our confidence, because we have begun by supposing it to be at least possible. If, on the other hand, all other things being equal, it seems improbable, we would begin to doubt it.

The arrows representing the impact of competence and plausibility therefore direct the score downwards. The greater the level of activation of the criterion – i.e. the more competent the source or the more plausible the information – the less the score drops by. All activation levels are considered, although there is no arrow for each one. This phenomenon is explained firstly by the implication of a total lack of knowledge and, secondly, by the fact that the constraints that we imposed are on the tendencies, rather than on the extent of the impact. We shall see, in section 9.11, that this flexibility means it is possible to adapt to each individual user.

In order to explain this phenomenon, note that we impose special treatment on an unquantifiable dimension. Because a lack of information does not enrich our knowledge, the score does not suffer from the impossibility of evaluation one or other of the dimensions. If, for instance, the plausibility of the information “cannot be evaluated”, the score, after its integration, does not change, regardless of the current score.

In addition, we can state that if it has been possible to evaluate the score at any one stage, it will be possible to evaluate it at all following stages. Thus, no arrow points to an unquantifiable score. Therefore, the lowest accessible level is level 5: “Improbable”. Therefore, because competence and plausibility have a negative impact, a piece of information scored as 5 before one or other of these stages will not change, in spite of the activation of the dimension. The arrows for all these levels are therefore identical to the horizontal. Similarly, with the exception of a piece of information that is “extremely probable” (scored as 1), certain levels of activation become indistinguishable.

We still need to explain the particular approach for a piece of information scored as “level of trust cannot be evaluated”. Such a score is reached when none of the previous dimensions have been evaluated. In the absence of any knowledge at all, the first criterion which can be evaluated takes on the role of initialisation. This explains why all trust levels are accessible for a piece of information which has not yet been evaluated, on the basis of the activation of the initialisation.

9.10.4. *Credibility*

Unlike the previous steps, the last phase of corroboration of the information offers a possibility for the score to increase. Indeed, the evolution of trust on the basis of confirmations and invalidations is a natural process. However, given that we envisage contradiction of the information, this criterion may just as easily lower the level of trust as raise it. This final step in Figure 9.1 therefore alters the score in a unique manner.

Naturally, a piece of information scored 1, even abundantly confirmed, cannot be believed any more than it already is. On the other hand, the lowering of its level of trust is more limited than in the previous stages. Indeed, the five degrees testifying to a correlation between the pieces of information under consideration range between a positive (confirmation) and a negative influence (invalidation). Supposing that the credibility activated as “insufficient confirmation” is neutral, the negative impact is reduced to the minimum two levels. On the other hand, as it is impossible for invalidation to increase the score, we only take into account the confirmations of

a piece of information that was previously unqualifiable or whose current score is minimal.

Finally, let us briefly return to the practical integration of corroboration into the information evaluation process. Remember that, in this stage, we compare already evaluated pieces of information. While the trust invested in the correlated pieces of information does not enter into the evaluation of the activation of the credibility, it is sensible to take it into account in the evolution of the score. In the same way as we lend greater credence to a column written by Andrew Wiles than to the affirmations of our friend when it comes to Fermat's theorem, the integration of invalidation or confirmation will be weighted by the *a priori* score of the correlated pieces of information, as proposed by [BES 07].

9.11. Personalisation of information evaluation: modelling levels of gullibility

The architecture we propose here has a certain subtle level of flexibility, which is precious to the user. Indeed, the arrows in Figure 9.1 represent the tendencies which we prescribe for the influence of each estimation of a criterion. However, we do not impose the strength of these influences. Given these degrees of freedom, the impact of the factors may be different depending on the importance the user attaches to each one. In addition, two users may have different configurations. These degrees of freedom therefore mean it is possible to adapt the method to the user and to represent different attitudes to levels of trust, as we describe in [REV 09].

Let us introduce the following definition:

STRATEGY.— *A strategy of usage of the information evaluation chain is a set of intensities for the influence of criteria. For each dimension, the strategy gives the result of the projection of its level of activation on the current score.*

This definition leads us to what these strategies enable us to represent:

CREDULITY STANCE.— *A credulity stance is an a priori choice made by the user about his expectations regarding the system's behaviour. It represents the user's sensitivity to the various dimensions and his gullibility. The choice of a strategy determines the user's preferred stance.*

It goes without saying that, if we are in the middle of the evaluation process, we cannot change our strategy.

In the rest of this section, we shall use examples to illustrate the notion of a stance of credulity for each of the stages, considering two extreme strategies and one neutral strategy.

9.11.1. Reliability of the source

Hitherto, we have presented the initialisation of the chain of information evaluation as being a balance between the reliability of the source and the *a priori* opinion regarding the level of trust. However, the important point for us was less equivalence between the levels than the importance of the initial opinion. Indeed, we can imagine that a person with extreme scepticism will refuse to believe even someone in whom he has total trust, merely on say-so.

The columns in Table 9.5 show three different strategies regarding the reliability of the source. It provides the initial score for each level of activation of the reliability of the source (from A to F, as shown by the rows). Strategy S_1 is the default strategy from Figure 9.1, with which we associate a neutral credulity stance. This associates A with 1, ... and F with 6. Note that apart from the need for equivalence between the “unquantifiable” degrees – meaning $F \equiv 6$ – any other acceptable strategy– i.e. where $A \geq B \geq C \geq D \geq E$ – can be used.

We can see that strategy S_2 does not lend any more credit to a completely reliable source than it does to a source whose reliability is slightly lower. Similarly, with the lowest level of reliability (levels D and E), it retracts all trust (score of 5). This extreme strategy represents a posture of scepticism on the part of the user. Conversely, the strategy S_3 remains open to a source whose reliability is shaky. That user begins to doubt if the source is absolutely not credible, but is prepared to believe any person who is relatively trustworthy. This strategy, even less common than the previous one, corresponds to a gullible user.

	S_1	S_2	S_3
A	1	2	1
B	2	2	1
C	3	3	2
D	4	5	3
E	5	5	4
F	6	6	6

Table 9.5. Three strategies for integration of reliability

Consider three users (U_1 , U_2 and U_3), each manifesting one of the three stances of credulity. U_1 is neutral, U_2 sceptical and U_3 gullible. Suppose that, in spite of the differences in their credulity, the three users have the same opinion about our photographer, whom they consider to be perfectly reliable. After initialisation, each of them has an *a priori* stance of $c_1=1$ (for U_1), $c_2=2$ (for U_2) and $c_3=1$ (for U_3).

9.11.2. Competence and plausibility

Competence and plausibility have similar impacts on the evolution of the score. For this reason and to make this document easier to read, Table 9.6 presents the same stances of credulity for both dimensions. Yet this is not a constraint of the model. A user may, for instance, prefer to trust in the source's competence rather than in his personal knowledge of the subject.

This table proposes a different formulation of the strategies, as we suppose that the evaluation process has been initialised. Therefore, the influence of the evaluation of the dimension (between 1 and 6, indexing the rows on a dark background) is calculated on the basis of the current level of the score, c , and the combination is bounded by the extent of the evaluation scale, meaning that the evaluated score must remain below 5.

For clarity's sake, the case where $c=6$ is not shown in the table. In that case, we would apply the initialisation strategies shown Table 9.5, with the activation levels for reliability being replaced by their equivalent for the dimension in question.

	S_1	S_2	S_3
1	c	c	c
2	$\min(c + 1, 5)$	$\min(c + 2, 5)$	c
3	$\min(c + 2, 5)$	$\min(c + 3, 5)$	$\min(c + 1, 5)$
4	$\min(c + 3, 5)$	$\min(c + 4, 5)$	$\min(c + 2, 5)$
5	$\min(c + 4, 5)$	$\min(c + 4, 5)$	$\min(c + 3, 5)$
6	c	c	c

Table 9.6. Three strategies for integrating competence and plausibility

Suppose that our three users are unaware that the mentioned photographer is not well versed in mathematics and therefore suppose him to be competent, i.e. level 2. The integration of competence into their evaluation of his statements will therefore give: $c_1 = \min(c + 1, 5) = 2$, $c_2 = \min(c + 2, 5) = 4$ and $c_3 = c = 1$. We can already

see a divergence of points of view depending on the users' respective stances of credulity.

All three users, in perfect harmony, also estimate that although it is surprising, the discovery of Fermat's proof is "realistic". The new update of the score is: $c_1 = \min(c + 1, 5) = 3$, $c_2 = \min(c + 2, 5) = 5$ and $c_3 = c = 1$. The opinions range between the gullible, who believes the information, and the sceptic, who denies it. U_1 , who is neutral, assumes the role of the agnostic.

Note that we can also represent a strategy by its association table for each dimension. The default strategy S_1 would, for competence and plausibility, have the association table represented in Table 9.7. We can see that as long as we do not try to represent more than one strategy, it is easy to include all combinations, particularly when $c = 6$. In addition, the development of the first column from Table 9.6 is again to be found here.

		Activation of the dimension					
		1	2	3	4	5	6
Current score	1	1	2	3	4	5	1
	2	2	3	4	5	5	2
	3	3	4	5	5	5	3
	4	4	5	5	5	5	4
	5	5	5	5	5	5	5
	6	1	2	3	4	5	6

Table 9.7. *Default strategy: updating of trust, on the basis of the evaluation of the competence or plausibility*

9.11.3. Credibility

As previously stated, the integration of the credibility in the evaluation is a tricky business. On the one hand, this is the only dimension which can increase the trust score. On the other, when it is projected onto the information, the *a priori* score (the score without the credibility) for the correlated information is taken into account. Table 9.8, once again, shows three different strategies for credibility. Again, the level of activation of the dimension indexes the rows in the table. We can see a variation in the contrast of the users, where the mistrusting stance favours invalidations over confirmations. Here again, we see the expression of the influence on the basis of the current trust score. For the sake of legibility, we suppose that

these expressions include the *a priori* score for the correlated information. The modifications are, again, bounded so as to remain within the authorised range of levels of trust.

	S ₁	S ₂	S ₃
1	$\max(c-2,1)$	$\max(c-1,1)$	$\max(c-2,1)$
2	$\max(c-1,1)$	c	$\max(c-2,1)$
3	c	$\min(c+1,5)$	C
4	$\min(c+1,5)$	$\min(c+2,5)$	C
5	$\min(c+2,5)$	$\min(c+3,5)$	$\min(c+1,5)$
6	c	c	C

Table 9.8. *Three strategies for integration of credibility*

Having learnt that Fermat's elegant proof has finally been discovered, the three users launch themselves – doubtless with different goals – in search of confirmations. As soon as they begin their search, they find Andrew Wiles' column, which admits the elegance of the new proof. A reliable source ($d_1 = 1, d_2 = 2, d_3 = 1$) and extremely competent ($d_1 = 1, d_2 = 2, d_3 = 1$) totally confirms a realistic fact ($d_1 = 2, d_2 = 4, d_3 = 1$), where d_i represents the score of the new piece of information for the user U_i .

U_1 therefore is convinced ($c_1 = \max(c-2, 1) = 1$), U_2 remains doubtful ($c_2 = \max(c-1, 1) = 4$) and U_3 , who has not changed at all, remains in agreement with himself ($c_3 = \max(c-2, 1) = 1$).

9.11.4. Discussion

This section briefly reviews the points of the proposed model whose constraints may seem significant. We also take the opportunity to suggest an avenue for potential refinement of the evaluation of the sources.

9.11.4.1. Order and sequentiality

The model proposed here presents information evaluation as an ordered and sequential process. This conception of the way in which trust is built may appear restrictive in the eyes of some people. Let us begin by reminding readers that there is no universally recognised model of this phenomenon in existence. In addition, the presentation given here of the process of information evaluation is more rigid than we actually view that process as being. For us, information evaluation is not a finite

process. As the dimensions can be altered during the process, the score can also be re-evaluated. Also, although to us it seems natural and justified, the order of integration of the dimensions which we recommend is not set in stone. The presence for all the evaluations of an unquantifiable level means, in particular, that the inclusion of a certain dimension may be delayed or omitted.

9.11.4.2. Updating

After credibility in the evaluation process, the information in question was evaluated on the basis of the reliability of its source, his competence in the domain of definition of the information, the plausibility of that information in view of the evaluator's perception of the world and, finally, an indicator of confirmation between uncertain pieces of information. At this point, we can consider updating the reliability score of the source, in the same vein as that proposed by [DEL 07], who updates the reliability of a source, when fusing different pieces of information, on the basis of its cumulative contradiction with other sources. This feedback of trust in the information regarding the reliability of its source is not part of the chain, but does have a very important effect. If a trusted source systematically provides unlikely pieces of information, it may be wise to review our judgment of his reliability, as shown by Figure 9.2. The rules governing the updating depend heavily on the source, and on the particular application. We see certain procedures for studying the reputation of sources of the literature as promising avenues to learn to qualify them in an automated fashion. This constitutes an important issue when handling large volumes of data, such as open sources.




Figure 9.2. *Updating of the source's reliability score following the information evaluation process*

9.12. Conclusion

This entire volume is dedicated to an essential task in the processing of sensitive data: information evaluation. Generally, both here and in previous literature and doctrines, information evaluation is considered to indicate the extent to which the events described by the information are accurate. Knowing whether the fact is real is, of course, of primordial importance for informed decision-making, but it may be difficult to evaluate. In such cases, most of us fall back on our own evaluation of the believability of the information, i.e. how much trust can be invested in it. Note that

this evaluation is not contradictory with the previous one: before we determine whether a fact is real or not, we all have an opinion about how much trust to lend to the piece of information which alerted us to that fact.

This chapter looks at the formalisation of such a process of constitution of trust, and then its evolution. Based on a simplified view of its establishment, a certain set of dimensions participating in its development are identified and described. It almost goes without saying that the selection, and definition, of the axes whereby trust evolves is debatable. However, as the objective is to propose a legible and interpretable model, each dimension needs to be interpreted in the same way by all users. Similarly, the next stage in the model, describing the influence of the dimensions on the trust, facilitates the sharing of evaluations between operators of differing levels of sensitivity.

This work opens up interesting perspectives for the study of trust. On the one hand, can humans be described by credulity stances? Are those stances constant or variable from one subject to another? How many different forms of credulity are there? The importance of the impossibility of evaluation, frequently mentioned here, also offers a whole area for future research. Must we, as we too frequently do, allow “I cannot say” to be equivalent to “whatever”? To what extent does this difference play a part in general human reasoning? Are we prepared to believe anything at all if we know nothing? The study of the evolution of trust itself brings out a variety of questions. Does the order in which pieces of information are discovered have an influence on the level of trust, if our gullibility remains the same? Finally, at the end of the chapter, we touch on the question of the source. When and how, for instance, do we change our minds about a source?

Many of the doors opened by the work discussed in this chapter correspond to essential themes in other research on information evaluation. Whatever way we conceive of it, information evaluation offers avenues for collaborations and other questions.

9.13. Bibliography

- [BAT 06] BATINI C., SCANNAPIECO M., *Data Quality*, Springer, Berlin, 2006.
- [BAU 02] BAUD J., *Encyclopédie du renseignement et des services secrets*, Charles Lavauzelle, Paris, 2002.
- [BER 99] BERTI L., “Quality and Recommendation of Multi-source Data for Assisting Technological Intelligence Applications”, *Lecture Notes in Computer Science*, p. 282-291, 1999.

- [BES 07] BESOMBES J., CHOLVY L., "Information fusion: using an ontology to information evaluation", *International Colloquium on Information Fusion*, p. 416-422, Xi'an, China, 2007.
- [BES 08] BESOMBES J., REVAULT D'ALLONNES A., "An Extension of STANAG2022 for Information Scoring", *11th International Conference on Information Fusion*, p. 1 635-1 641, Cologne, Germany, 2008.
- [BOR 98] BORDEN D.L., HARVEY K., *The Electronic Grapevine: Rumor, Reputation, and Reporting in the New On-Line Environment*, Lawrence Erlbaum Associates, Mahwah, United States, 1998.
- [BRI 98] BRIN S., PAGE, L., "The anatomy of a large-scale hypertextual Web search engine", *Computer networks and ISDN systems*, vol. 30, n° 1-7, p. 107-117, 1998.
- [CHE 07] CHEN P., XIE H., MASLOV S., REDNER S., "Finding scientific gems with Google's PageRank algorithm", *Journal of Informetrics*, vol. 1, n° 1, p. 8-15, 2007.
- [CHO 03] CHOLVY L., NIMIER L., "Information Evaluation: discussion about STANAG 2022 recommendations", *NATO-IST Symposium on military data and information fusion*, Prague, Czech Republic, 2003.
- [CHO 04] CHOLVY L. "Information Evaluation in fusion: a case study", *International Conference on Processing and Management of Uncertainty in Knowledge-based Systems (IPMU 2004)*, Perugia, Italy, 2004.
- [DEL 07] DELMOTTE F., "Detection of defective sources in the setting of possibility theory", *Fuzzy Sets and Systems*, vol. 158, n° 5, p. 555-571, 2007.
- [DEP 05] DEPARTMENT OF THE ARMY HEADQUARTERS, Human Intelligence Collector Operations, Field Manual, FM 2-22.3, 2005.
- [DIS 01] DISS & DISC DEFENCE INTELLIGENCE AND SECURITY SCHOOL & DEFENCE INTELLIGENCE AND SECURITY CENTRE, Intelligence Wing Student Précis, 2001.
- [HOR 07] HORN G. Van, JOHNSTON L., "Evolutionary Controversy and a Side of Pasta: The Flying Spaghetti Monster and the Subversive Function of Religious Parody", *GOLEM: Journal of Religion and Monsters*, p. 1-32, 2007.
- [MIL 55] MILLER G.A., "The Magical Number Seven, Plus or Minus Two: Some Limits on our Capacity for Processing Information", *Psychological Review*, vol. 101, n° 2, p. 343-352, 1955.
- [NIM 04] NIMIER V., "Information Evaluation: a Formalisation of Operational Recommendations", *Seventh International Conference on Information Fusion*, p. 1166-1171, Stockholm, Sweden, 2004.
- [REV 07] REVAULT D'ALLONNES A., AKDAG H., POIREL O., "Trust-moderated information-likelihood. A multi-valued logics approach", *Computation and Logic in the Real World*, Third Conference on Computability in Europe, p. 1-6, Sienna, Italy, 2007.
- [REV 09] REVAULT D'ALLONNES A., BESOMBES J., "Critères d'évaluation contextuelle pour le traitement automatique", *Qualité des données et des connaissances (QDC'09). Atelier des 9^e Journées francophones extraction et gestion des connaissances (EGC)*, p. A6 13-20, Strasbourg, 2009.

- [REV 11] REVAULT D'ALLONNES A., Evaluation sémantique d'informations symboliques : la cotation, Doctoral thesis, Université Pierre et Marie Curie, Paris 6, 2011.
- [SHA 09] SHAH D., ZAMAN T.R., "Rumors in a Network: Who's the Culprit?", *Neural Information Processing Systems (NIPS)*, Vancouver, Canada, 2009.
- [SIM 05] SIMON F., *Journaliste : dans les pas d'Hubert Beuve-Méry*, Arléa, Paris, 2005.
- [TOC 1835] TOCQUEVILLE A. de, *De la démocratie en Amérique*, Folio Histoire, Paris, 1835.
- [WAN 96] WANG R.Y., STRONG D.M., "Beyond Accuracy: What Data Quality Means to Data Consumers", *Journal of Management Information Systems*, vol. 12, n° 4, p. 5-33, 1996.