How can Subjective Impulsivity play a role among Information Sources in Weather Scenarios?

Rino Falcone and Alessandro Sapienza
Institute of Cognitive Sciences and Technologies,
ISTC – CNR, Rome, Italy
{rino.falcone, alessandro.sapienza}@istc.cnr.it"

Abstract— **The topic of critical hydrogeological phenomena, due to flooding, has a particular relevance given the risk that it implies. In this paper we simulated complex weather scenarios in which forecasts coming from different sources become relevant. Our basic idea is that agents can build their own evaluations on the future weather events integrating these different information sources also considering how trustworthy each single source is with respect to each individual agent. These agents learn the sources' trustworthiness in a training phase. Moreover, agents are differentiated on the basis of their own ability to make direct weather forecasts, on their possibility to receive bad or good forecasts from the authority, and on the possibility of being influenced by the neighbors' behaviors. Quite often in the real scenarios some irrational behaviors rise up, whereby individuals tend to impulsively follow the crowd, regardless of its reliability. To model that, we introduced an impulsivity factor that measures how agents are influenced by the neighbors' behavior, a sort of "crowd effect". The results of these simulations show that, thanks to a proper trust evaluation of their sources made in the training phase, the different kinds of agents are able to better identify the future events.**

Keywords— trust; social simulation; cognitive agents.

I. INTRODUCTION

The role of the impulsivity in human behaviors has relevant effects in the final evaluations and decisions of both individuals and groups. Although we are working in the huge domain of social influence [4][7][8] we consider here impulsivity as an attitude of taking a decision just basing on a partial set of evidence, although further evidence is easily reachable and acquirable. Sometimes this kind of behavior can produce unpredictable consequences that were not taken in consideration while deciding [16]. Impulsivity is a multifactorial concept [5], however we are interested in identifying the role that it can play in a specific set of scenarios.

In particular, in this paper we simulated complex weather scenarios in which there are relevant forecasts coming from different sources. Our basic idea is that agents can build their own evaluations on the future weather events integrating these different information sources, also considering how trustworthy the single source is with respect to each individual agent. These agents learn the sources' trustworthiness in a training phase. They are differentiated i) on the basis of their ability to make direct weather forecasts, ii) on their possibility to receive bad or good forecasts from an authority, and iii) on the possibility of being influenced by the neighbors' behaviors.

So given this picture, our simulations inquired several interactions among different kinds of agents, testing different weather scenarios with different levels of impulsivity. We also considered the role that both expertise and information play on the impulsivity factor.

The results of these simulations show that, thanks to a proper trust evaluation of their sources made through the training phase, the different kinds of agents are able to better identify the future events. Some particular and interesting result concerns the fact that impulsivity can be considered, in specific situations, as a rational and optimizing factor, in some way contradicting the nature of the concept itself. In fact, as in some human cases, it can be possible that we have learned specific behaviors based on just one information source that is enough for the more efficient behavior although we could access to other different and trustworthy sources. In that case we consider as impulsive a behavior that is in fact fully effective.

II. THE TRUST MODEL

According to the literature [1][2][10][11][17] trust is a promising way to deal with information source. In particular in this work we are going to use the computational model of [13], which is in turn based on the cognitive model of trust of Castelfranchi and Falcone [3]. It exploits the Bayesian theory, one of the most used approaches in trust evaluation [9][12][18], representing all the information as a probability distribution function (PDF).

In this model each information source S is represented by a trust degree called *TrustOnSource* [6], with 0 ≤*TrustOnSource* ≤ 1, plus a bayesian probability distribution PDF that represents the information reported by S. The *TrustOnSource* parameter is used to smooth the information referred by S: the more I trust the source, the more I consider the PDF; the less I trust it, the more the PDF is flattened. Once an agent gets the contribution from all its sources, it aggregates the information to produce the global evidence (GPDF), estimating the probability that each event is going to happen.

A. Feedback On Trust

We want to let agents adapt to the context in which they move. This means that, starting from a neutral trust level (that does not imply trust or distrust) agents will try to understand how much to rely on each single information source (*TrustOnSource)*, using direct experience for trust evaluations [14][15]. To do that, they need a way to perform feedback on trust. We propose to use weighted mean. Given the two parameters α and β[1], the new trust value is computed as:

$$newTrustOnSource = \alpha * TrustOnSource + \beta * performanceEvaluation \qquad (1)$$

$$\alpha + \beta = 1$$

TrustOnSource is the previous trust degree and *performanceEvaluation* is the objective evaluation of the source performance. This last value is obtained comparing what the source said with what actually happened. Considering the PDF reported by the source (that will be split into five parts as we have 5 possible events), we will have that the estimated probability of the event that actually occurred is completely taken into account and the estimated probability of the events immediately near to it is taken into account for just 1/3. We in fact suppose that even if the evaluation is not right, it is not, however, entirely wrong. The rest of the PDF is not considered. Let's suppose that there was the most critical event, which is event 5. A first source reported a 100% probability of event 5, a second one a 50% probability of event 5 and a 50% of event 4 and a third one asserts 100% of event 3. Their performance evaluation will be: Source1=100%; Source2=66.67% (50% + (50/3)%); Source3: 0%. Figure 1 shows the corresponding PDFs.

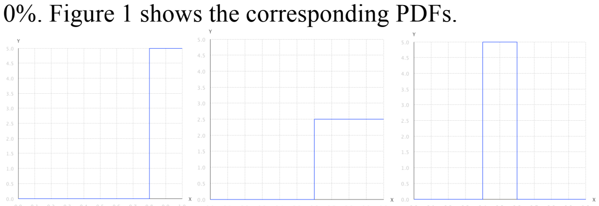

Fig. 1. (a) A source reporting a 100% probability of event 5. (b) A source reporting a 50% probability of event 5 and 50% probability of event 4. (c) A source reporting a 100% probability of event 3.

III. The Platform

Exploiting NetLogo [19], we created a very flexible platform, where a lot of parameters are taken into account to model a variety of situations.

Given a population distributed over a wide area, some weather phenomena happen in the world with a variable level of criticality.

The world is made by 32x32 patches, which wraps both horizontally and vertically where agents are distributed in a random way and is populated by a number of cognitive agents (citizens) that have to evaluate which will be the future weather event on the basis of the information sources they have and of the trustworthiness they attribute to these different sources.

We provided the framework with five possible events, going from 1 to 5, with increasing level of criticality: level 1 stands for no events, there is no risk at all for the citizens; level 5 means that there will be a tremendous event due to a very high level of rain, with possible risks for the agents sake. The other values represent intermediate events with increasing criticality.

In addition to citizens, there is another agent called authority. Its aim is to inform promptly the citizens about the weather phenomena. The problem is that, for their nature, weather forecasts improve their precision nearing to the event. Consequently, while the time passes the authority is able to produce a better forecast, but it will not be able to inform all the citizens, as there will be less time to spread information.

A. Information Sources

To make a decision, each citizen can consult a set of information sources, reporting to it some evidence about the incoming meteorological phenomenon.

We considered the presence of three kinds of information sources (whether active or passive) for citizens:

1. Their *personal judgment*, based on the direct observation of the phenomena. Although this is a direct and always true (at least in that moment) source. In general, a common citizen is not always able to understand the situation, maybe because it is not able, it does not possess any instrument or it is just not in the condition to properly evaluate a weather event. So we have introduced two kinds of agents: the expert ones and the inexpert ones.
2. *Notification from authority*: the authority distributes into the world weather forecast, trying to prepare citizens to what is going to happen. While the time pass, it is able to produce a better forecast, but it will not be able to inform everyone. In this sense we have two kinds of agents: the well-informed ones and the ill-informed ones.
3. *Others' behavior*: agents are in some way influenced by community logics, tending to partially or totally emulate their neighbors' behavior (other agents in the radius of 3 NetLogo patches). The probability of each event is directly proportional to the number of neighbors making each kind of decision. This source can have a positive influence if the neighbors behave correctly, otherwise it represents a drawback.

None of these sources is perfect. In any situation there is always the possibility that a source reports wrong information.

B. Agents' Description

At the beginning of the simulation, the world is populated by a number of citizens, having the same neutral trust value 0.5 for all their information sources. This value represents a situation in which citizens are not sure if to trust or not a given source

[1] Of course changing the values of α and β will have an impact on the trust evaluations. With high values of α/β, agents will need more time to get a precise evaluation, but a low value (below 1) will lead to an unstable evaluation, as it would depend too much on the last performance. We do not investigate these two parameters in this work, using respectively the values 0.9 and 0.1. In order to have good evaluations, we let agents make a lot of experience with their information sources.

(a value of 1 represents complete trust and 0 complete distrust).

There are two main differences between citizens. The first one relies on how able they are in seeing and reading the phenomena. In fact, in the real world not all the agents have the same abilities. For representing these different abilities we associated to the citizens' evaluations different values of standard deviation related to the meteorological events.

In order to shape this, we divided agents in two sets:

1. **Class 1:** good evaluators; they have good capabilities to read and understand what is going to happen. They will be quit always able to detect correctly the event (90% of times; standard deviation of 0.3), and then we expect them to highly trust their own opinion.
2. **Class 2:** bad evaluators; they are not so able to understand what is going on (20% of times, that is the same performance of a random output; standard deviation of 100). In order to understand which weather event is going to happen in the near future they have to consult other information sources.

The second difference is due to how easily they are reached by the authority. The idea is that the authority reaches everyone, but while the time passes it produces new updated information. There will be agents able to get update information, but not all of them will be able to do it. To model this fact, we defined two agent classes:

1. **Class A**: they possess the newest information produced by the authority; the information they receive has a 90% probability to be correct;
2. **Class B**: they are only able to get the first prevision of the authority; the information they receive has a 30% probability to be correct.

C. The authority

The authority's aim is to inform citizens about what is going to happen. The best case would be the one in which it is able to produce a correct forecast and it has the time to spread this information through all the population. However reaching everyone with correct information is as desirable as unreal. The truth is that weather forecast's precision increases while the event is approaching.

In the real world the authority does not stop making prediction and spreading it. As already said, in the simulations we modeled this dividing the population into two classes. Agents belonging to the class B will just receive the old information. This is produced with a standard deviation of 1.5, which means that this forecast will be correct in the 30% of times. Then the authority will spread updated information. Being closer to the incoming event, this forecast has a higher probability to be correct. It is produced with a standard deviation of 0.3, so that it will be correct in the 90% of times.

As a choice, we made that in the simulation it is more convenient to use as a source the authority rather than personal evaluations, except for experts that are as good as a reliable authority.

D. Citizens' Impulsivity

Sometimes impulsivity overcomes logic and rationality. This is more evident in case of critical situations, but it is still plausible in the other cases. Maybe the authority reports a *light* event, but the neighbors are escaping. In this case it is easy to be influenced by the crowd decision, to make a decision solely based on the social effect, letting "irrationality" emerge. Let us explain better this concept of "irrationality": in fact we consider that an agent follow an "irrational" behavior when it takes a decision considering just one of its own information sources although it has also other available sources to consult. In this work we consider just the social source as subjected to the impulsivity conditioning.

Impulsivity is surely a subjective factor so our citizens are endowed with an **impulsivity threshold**, which measures how prone they are to irrational choice due to the crowd effect. This threshold is affected by the other two sources, the authority and the experience, as they add rationality in the decisional process.

The threshold goes from 0 to 1, and given a value of this threshold, being well informed or an expert gives a plus 0.2 to it (it an agents is both informed and expert, it is a plus 0.4). Therefore it is important for individual to be informed, so that they are less sensible to irrationality and they are able to produce decisions based on more evidence. In our experiments we consider a common impulsivity threshold ($\boldsymbol{Ith_{Com}}$) that is the same for all the agents and two additional factors ($\boldsymbol{Add_{Inf}}$ and $\boldsymbol{Add_{Exp}}$) due to the information and the expertise each agent has that determine the individual impulsivity threshold ($\boldsymbol{Ith_{Agent}}$). In practice, given an agent A, we can say that:

$$\mathrm{Ith_A} = \mathrm{Ith_{Com}} + \mathrm{Add_{Inf}} + \mathrm{Add_{Exp}} \qquad (2)$$

The threshold is compared with the PDF reported by the social source. If there is one event that has a probability to happen (according to this source) greater than the impulsivity threshold, then the agents act impulsively.

E. Platform Input

The first thing that can be customized is the **number of citizens** in the world and how they are distributed between the **performance categories** and the **reachability categories**. Then, one can set the value of the two parameters **α and β**, used for updating the sources' trust evaluation. It is possible to change the **authority reliability** concerning each of the reachability categories. Concerning the training phase, it is possible to change its **duration**. Finally, it is possible to set the **impulsivity threshold** and how much it will be modified by each rational source.

F. Workflow

The simulation is divided into two steps. The first one is called "**training phase**" and has the aim of letting agents make experience with their information sources, so that they can determine how reliable each source is.

At the beginning of this phase, we generate a world containing an authority and a given number of citizens, with different abilities in understanding weather phenomena and different possibility to be informed by the authority. Then citizens start collecting information, in order to understand which event is going to happen. The authority gives forecast reporting its estimated level of criticality. As already explained, it produces two different forecasts. All the citizens will receive the first one, but it is less precise as it is not close enough to the event. The second one is much more precise, but being close to the event it is not possible for the authority to inform all the citizens.

In any case, being just forecasts, it is not sure that they are really going to happen. They will have a probability linked to the precision of the authority (depending on its standard deviation).

Then citizens evaluate the situation on their own and also exploit others' evaluations (by the effect of their decisions). Remember that the social source is the result of the process aggregating the agents' decisions in the neighborhood: if a neighbor has not yes decided, it is not considered. If according to the others' evaluation there is one event that has a probability to happen greater than the impulsivity threshold, then they act impulsively. This means that they are not going to consider the three sources they have, but just the social one. If this does not happen, then they consider all the information they can access and they aggregate each single contribution according to the corresponding trust value. Finally they estimate the possibility that each event happens and select the choice that minimizes the risk.

While citizens collect information they are considered as "thinking", meaning that they have not decided yet. When they reach the decisional phase, the citizens have to make a decision, which cannot be changed anymore. This information is then available for the others (neighborhood), which can in turn exploit it for their decisions. At the end of the event, citizens evaluate the performance of the source they used and adjust the corresponding trust values. This phase is repeated for 100 times (then there will be 100 events) so that agents can make enough experience to judge their sources.

After that, there is the "**testing phase**". Here we want to understand how agents perform, once they know how reliable their sources are. In this phase, we will compute the accuracy of their decision (1 if correct, 0 if wrong).

IV. Simulations

In the simulations we tested the effect of impulsivity on a population with different abilities to interpret the events and with different possibility to be informed by the authority. It is worth noting that impulsivity affects everyone, even the more expert or informed can be misled by their neighbors' decisions.

A. Simulations' Outputs

In this section we describes the metrics we used in order to understand and analyze each simulation.

The first one is **agents' performance**. Concerning a single event, the performance of an agent is considered correct (and assumes value 1) if it correctly identified the event or wrong (and assumes value 0) if it made a mistake with the events.

The second dimension we analyze is the **decisional distance**. Suppose that there will be event 5. An agent X foresees event 4, while another agent Y supposes there will be event 1. Both this decision are wrong, but the decision of agent Y is much more wrong that the one of X. Practically speaking, in case of a critical event (represented in fact by event 5) agent X could take some important measure to prevent damages to it and its properties, while agents Y just does nothing. Maybe both the agents suffer damages, but probably X manages to reduce damages or at least the probability to be damaged, while Y does not.

For a single agent its decisional distance is defined as the difference between the event that is going to happen and the agent's forecast. For instance, agent X's decisional distance is 1, while Y's is 4. We want this dimension to be the lowest; ideally in a perfect world it should be 0, meaning that the agent makes the right prediction.

A third dimension is represented by the **percentage of impulsive decision**.

The last dimension that we investigate is the **trust on the information sources**. The section "Feedback on Trust" explains how agents produce their trust evaluations, based on the source performance. They possess a trust value for each of their three sources.

We introduced these four metrics for individual agents. Actually in the results they will be presented aggregating the values of a category of agents and mediating them for the number of times that the experiment is repeated (500 times).

In particular, in order to provide a better analysis of the results, we are not going to simply consider the category of agents previously described, but their combinations: 1A = well informed and expert agents; 2A = well informed and not expert agents; 1B = less informed and expert agents; 2B = less informed and not expert agents.

B. Simulations' Scenario

In the scenarios we investigated, the percentage of well-informed citizens and the percentage of expert citizens is the same, as we are mainly interested in increasing/decreasing the quantity of good information and expertise that the population possesses. Of course, as the assignment of citizens to categories is random, it is possible an overlap between these categories: a well-informed citizen can also be an expert. Simulation settings:

1. *number of agents*: 200;
2. *α and β*: respectively 0.9 and 0.1;
3. *authority reliability*: we used a standard deviation of 1.5 to produce the first forecast reported by the authority (it is correct about 90% of the time) and 0.3 for the second one (its forecasts are correct about 30% of the time);
4. *percentage of well informed citizens and percentage of expert citizens*: {10-10,20-20,30-30,45-45,60-60,75-75}.
5. *training phase duration*: 100 events;

6. *Impulsivity threshold*: we experimented the four cases {0.3,0.5,0.7,0.9}

For sake of simplicity, as the percentage of well informed citizens and of expert citizens is the same in each experiment, we will use this value to identify the specific case. For instance, the “case 10-10” is the one with 10% of well-informed citizens and of expert citizens.

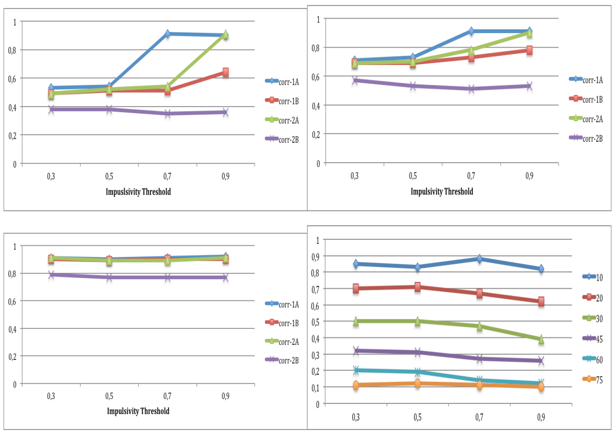

Fig. 2. (a) Agents’ correctness in the case 10-10. (b) Agents’ correctness in the case 30-30. (c) Agents’ correctness in the case 75-75. (d) Agents’ decisional distance

It is worth noting that when the impulsivity threshold (***Ith***$_{Com}$) is 0.9 then well informed or expert agents are not impulsive for sure (given that for those agents ***Ith***$_{Agent}$ saturates the max value 1). When the impulsivity threshold (***Ith***$_{Com}$) is 0.7, it is necessary to be both informed and expert to not be impulsive in any case. In the other cases agents could act impulsively, according to the modality explained above. This is clearly visible with an impulsivity threshold of 0.7, especially in Figure 2a but also in Figure 2b: there is a big difference between 1A agents’ performance and the others. In practice, in the given composition of agents showed in Figure 2a and 2b, impulsive agents are penalized. Let us explain in detail.

Figure 2a shows the case 10-10 (10% of well informed citizens and 10% of expert citizens). Here the majority of the citizens, approximately the 81%, belongs to the category 2B (not well informed and not expert) represented in violet. They are so many that their evaluation of the events when socially transmitted to their neighbors will have a negative influence on them, especially when there is a low value of common impulsivity threshold. Increasing the percentage of informed/expert citizens this effect tends to disappear, as showed by Figure 2b and 2c.

From Figure 2a, 2b and 2c it clearly results that the performance of 1A, 1B and 2A agents increases when we increases the value of the impulsivity threshold (agents are less impulsive). In fact increasing this component, these agents will not be influenced by the crowd effect and they will be able to decide on the basis of all their sources.

Differently form the others, if we focus on the 2B category (both bad evaluators and misinformed) we notice an interesting effect: in all the cases, increasing the impulsivity threshold the performance of 2B citizens decreases. This is due to the fact that, being less impulsive will have more weight on their own information and on their own expertise in their final evaluations. But not being well informed or experts, there is a higher probability that they will be wrong.

Concerning agents’ decision, it is interesting not just to see the percentage of success, but also how they differ from the correct decision. The decisional distance reports this information.

From the graphs in Figure 2d we can clearly see that increasing the quantity of information in the world (experts and informed agents) the decisional distance decreases. It also seems to decrease increasing the impulsivity threshold: in practice, the forecasts are more correct when the agents are more informed or expert and less impulsive.

C. *Trust Analysis*

Talking about trust, analyzing the four categories 1A, 1B, 2A and 2B the components of self-trust and authority trust do not change. They in fact assume a fixed value in all the cases, not being influenced by the impulsivity threshold or by the quantity of information in the world (just by its quality). Figure 3a, 3b, 3c and 3d show these values respectively to the categories 1A, 1B, 2A and 2B.

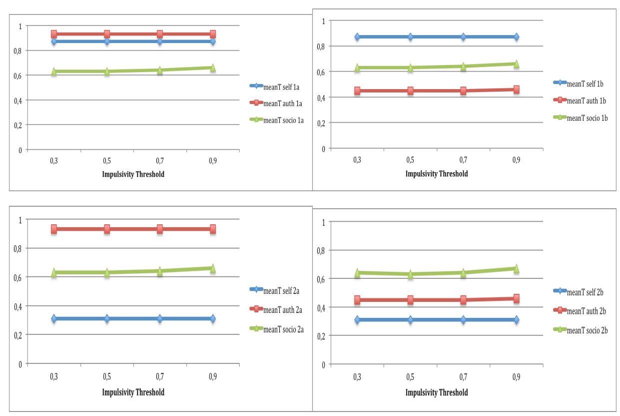

Fig. 3. (a) Trust degrees of the agents belonging to the 1A category in the case 30-30. (b) Trust degrees of the agents belonging to the 1B category in the case 30-30. (c) Trust degrees of the agents belonging to the 2A category in the case 30-30. (d) Trust degrees of the agents belonging to the 2B category in the case 30-30

What changes is of course the social trust. In fact, event if it is completely independent from the agent’s nature, it strictly depends on its neighborhood: the more performative they are, the higher the social trust will be. This is clearly visible in Figure 4. We can see how the social trust increases increasing the percentage of expert/informed citizens.

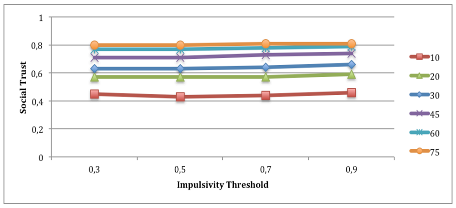

Fig. 4. Social trust of all the agents in the six cases

V. CONCLUSIONS

In this work we analyzed the effect of subjective impulsivity inside critical weather scenarios. We proposed some simulations in which a population of citizens (modeled through cognitive agents) has to face weather scenarios and needs to exploit its information sources to understand what is going to happen. In these situation agents can act "rationally" (basing their choice in the global evidence they possess) or impulsively, just emulating their neighbors due to a sort of "crowd effect".
First of all, we proved that even if impulsivity has a strongly negative impact on informed or expert agents, on the contrary it is useful for the remaining 2B agents. Further, we showed that it is not good to have a high percentage of 2B agents, as they have a negative impact also on the agents belonging to the other categories. This is a quite predictable effect, even if it is interesting appreciate the various levels of impulsivity that determine the different impacts.

Then we analyzed the role played by social trust. Given a value for the impulsivity threshold and a percentage of informed and expert citizens, we showed that it assumes a fixed value for all the citizens, as it is independent by the agent's category. However, it as a positive impact on agents with less information (2B agents) while it tends to have a negative effect on the on increasing the correctness of the information that agents own.

A last point regards the decisional distance, which provide a much more precise analysis of the decisions' correctness. We saw that it tends to decrease increasing the impulsivity threshold. This means that less impulsive agents can produce a better evaluation: even if they are wrong, their decisions are nearer to the correct decision.

ACKNOWLEDGMENTS

This work is partially supported by the project CLARA—CLoud plAtform and smart underground imaging for natural Risk Assessment, funded by the Italian Ministry of Education, University and Research (MIUR-PON).

REFERENCES

[1] Amgoud L., Demolombe R., An Argumentation-based Approach for Reasoning about Trust in Information Sources , In *Journal of Argumentation and Computation, 5(2), 2014*

[2] Barber, K. S., & Kim, J. (2001). Belief revision process based on trust: Agents evaluating reputation of information sources. In *Trust in Cyber-societies* (pp. 73-82). Springer Berlin Heidelberg.

[3] Castelfranchi C., Falcone R., Trust Theory: A Socio-Cognitive and Computational Model, John Wiley and Sons, April 2010.

[4] Cialdini, Robert B. (2001). *Influence: Science and practice (4th ed.)*. Boston: Allyn & Bacon. ISBN 0-321-01147-3

[5] Evenden, J. L. (1999). "Varieties of impulsivity". *Psychopharmacology*. **146** (4): 348–61. doi:10.1007/PL00005481

[6] Falcone, R., Sapienza, A., & Castelfranchi, C. (2015). The relevance of categories for trusting information sources. *ACM Transactions on Internet Technology (TOIT)*, *15*(4), 13.

[7] Genter, K., & Stone, P. (2016, May). Adding Influencing Agents to a Flock. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems* (pp. 615-623). International Foundation for Autonomous Agents and Multiagent Systems.

[8] Latané, B. (1981). The psychology of social impact. American Psychologist, 36, 343-356.

[9] Melaye, D., & Demazeau, Y. (2005). Bayesian dynamic trust model. In Multi-agent systems and applications IV (pp. 480-489). Springer Berlin Heidelberg.

[10] Melo, Victor S., Alison R. Panisson, and Rafael H. Bordini. "Trust on Beliefs: Source, Time and Expertise.", in Proceedings of the 18th International Workshop on Trust in Agent Societies co-located with the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2016), Singapore, May 10, 2016, Ceur Workshop Proceedings, vol 1578, paper 6.

[11] Parsons, S., Sklar, E., Singh, M. P., Levitt, K. N., & Rowe, J. (2013, March). An Argumentation-Based Approach to Handling Trust in Distributed Decision Making. In *AAAI Spring Symposium: Trust and Autonomous Systems*.

[12] Quercia, D., Hailes, S., & Capra, L. (2006). B-trust: Bayesian trust framework for pervasive computing. In Trust management (pp. 298-312). Springer Berlin Heidelberg.

[13] Sapienza, A., & Falcone, R. A Bayesian Computational Model for Trust on Information Sources, in proceedings of the conferenze WOA 2016, Catania, Ceur workshop proceedings, vol 1664, pp. 50-55.

[14] Schmidt, S., Steele, R., Dillon, T. S., & Chang, E. (2007). Fuzzy trust evaluation and credibility development in multi-agent systems. *Applied Soft Computing*, *7*(2), 492-505.

[15] Theodorakopoulos, G., & Baras, J. S. (2006). On trust models and trust evaluation metrics for ad hoc networks. *IEEE Journal on selected areas in Communications*, *24*(2), 318-328.

[16] VandenBos, G. R. (2007). *APA dictionary of psychology*. Washington, DC: APAVillata, S., Boella, G., Gabbay, D. M., & Van Der Torre, L. (2011, June). Arguing about the trustworthiness of the information sources. In *European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty* (pp. 74-85). Springer Berlin Heidelberg.

[17] Villata, S., Boella, G., Gabbay, D. M., & Van Der Torre, L. (2011, June). Arguing about the trustworthiness of the information sources. In *European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty* (pp. 74-85). Springer Berlin Heidelberg.

[18] Wang, Y., & Vassileva, J. (2003, October). Bayesian network-based trust model. In Web Intelligence, 2003. WI 2003. Proceedings. IEEE/WIC International Conference on (pp. 372-378). IEEE.

[19] Wilensky, U. (1999). NetLogo. http://ccl.northwestern.edu/netlogo/. Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL.