Interactions Among Information Sources in Weather Scenarios: The Role of the Subjective Impulsivity

Rino Falcone and Alessandro Sapienza(✉)

ISTC-CNR, Rome, Italy
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 are relevant forecasts coming from different sources. Our idea is that agents can build their own evaluations on the future weather events integrating these different information sources also considering how much trustworthy the single source is with respect to each individual agent. These agents learn the trustworthiness of the sources in a training phase. 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 an 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 much 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 through the training phase, the different kinds of agents are able to better identify the future events.

Keywords: Trust · Social simulation · Hydrogeological risk

1 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, 8, 7], we consider here impulsivity as an attitude of making a decision just basing this decision on a partial set of evidence, also in those cases in which more evidence is easily reachable and acquirable. Sometimes this kind of behavior can produce consequences that were not taken in consideration at the moment of the decision [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 trustworthiness of the sources in a training phase. Agents are differentiated (i) on the basis of their own ability to make direct weather forecasts, (ii) on their possibility to receive bad or good forecasts from an authority, and (iii) on the

Y. Demazeau et al. (Eds.): PAAMS 2017, LNAI 10349, pp. 56–69, 2017.
DOI: 10.1007/978-3-319-59930-4_5

possibility of being influenced by the neighbors' behaviors. 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 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 results regard 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 is possible that we learned specific behaviors just basing on 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.

2 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]. Here information is represented as a probability distribution function (PDF).

In this model each information source S is represented by a trust degree called *TrustOnSource* [6], with $0 \leq TrustOnSource \leq 1$, plus a bayesian probability distribution PDF (Probability Distribution Function) 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.

2.1 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:

---

[1] The values of α and β 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.

$$newTrustOnSource = \alpha * TrustOnSource + \beta * performanceEvaluation$$

$$\alpha + \beta = 1$$

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.

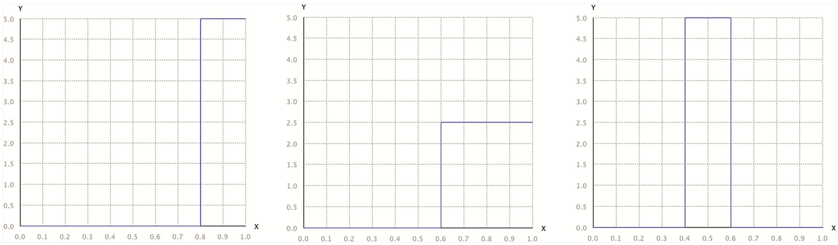

Fig. 1. (a) A source reporting a 100% probability of event 5. (b) a source reporting a 50% probability of event 5 and 50% probability of event 4. (c) a source reporting a 100% probability of event 3.

3 The Platform

Exploiting NetLogo [19], we created a very flexible platform, taking into account a lot of parameters 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 32 × 32 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.

3.1 Information Sources

To make a decision, citizens 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 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 unable 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.

3.2 Agents' Description

At the beginning, all the citizens have 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 represents complete trust and 0 complete distrust). There are two main differences between citizens. The first one relies on how much they are able to see and to read the phenomena. In fact, in the real world not all the agents have the same abilities. In order to shape this, we associated to the citizens' evaluations different values of standard deviation related to the meteorological events, dividing them 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). For better understanding which will be the future weather event 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: the have 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.

3.3 The Authority

The authority's aim is to inform citizens about what is going to happen. The best case is the one in which it to produce a correct forecast and it has the time to spread this information through all the population. However this 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 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 sill be correct in 90% of times. As a choice, 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.

3.4 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 makes 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 much they are prone to a less informed choice due to the crowd effect. This threshold is in turn 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 a sort of irrational choice 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 potential 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:

$$\boldsymbol{Ith_A = Ith_{Com} + Add_{Inf} + Add_{Exp}}$$

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.

3.5 Platform Inputs

The first thing that can be customized is the **number of citizens** and their distribution 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. One can also set the **events' probability** that is the frequency with which each event will happen. 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.

3.6 Workflow

Each 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, the 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 decided, it is not considered. If according to the social source there is one event that has a probability to happen greater than the impulsivity threshold, then they act impulsively: they will not consider the other sources. 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 decide. 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 much reliable their source are. In this phase, we will compute the accuracy of their decision (1 if correct, 0 if wrong).

4 Simulations

We investigated two main scenarios. In the first one 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. In this case impulsivity affects everyone, as even the more expert or informed can be misled by their neighbors' decisions. In the second simulation we introduce a decisional order between agents. The best informed will be the first to decide, followed by the most able to understand the events. In this second ideal world impulsivity has a much smaller influence on decision. In addition, as the most rational agents (which also possess more evidence about the events) will decide before the worst informed and able agents, there will be a positive effect on the performance of all the agents.

In order to understand and analyze each simulation, we are going to use two metrics. 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 is the **trust on the information sources**. Section 2.1 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 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 indicated in Sect. 3.2 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.

4.1 First Simulation

Here there is no decisional order, so that each citizen can influence and be influenced by everyone. 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: *number of agents* = 200; α *and* β = respectively 0.9 and 0.1; *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 time) and 0.3 for the second one (its forecasts are correct about 30% of time); *percentage of well informed citizens and percentage of expert citizens* = {10-10, 20-20, 30-30, 45-45, 60-60, 75-75}; *events' probability* (from the lightest to the most critical one) {35%, 30%, 20%, 10%, 5%}; *training phase duration* = 100 events; i*mpulsivity 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.

It is worth noting that when the impulsivity threshold ($\boldsymbol{Ith_{Com}}$) is 0.9 then well informed or expert agents are not impulsive for sure (given that for those agents $\boldsymbol{Ith_{Agent}}$ saturates the max value 1). When the impulsivity threshold ($\boldsymbol{Ith_{Com}}$) 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 in Sect. 3.4. This is clearly visible with an impulsivity threshold of 0.7, especially in Fig. 2 but also in Fig. 3: there is a big difference between 1A agents' performance and the others. In practice, in the given composition of agents showed in Figs. 2 and 3, impulsive agents are penalized. Let us explain in detail. Figure 2 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 influences negatively their neighbors through the social source, 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 Figs. 3 and 4. From Figs. 2, 3 and 4 it clearly results that the performance of 1A, 1B and 2A agents increases when we increase 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.

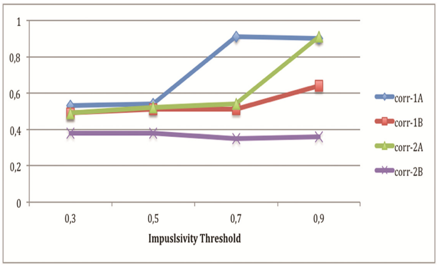

Fig. 2. Agents' correctness in the case 10-10

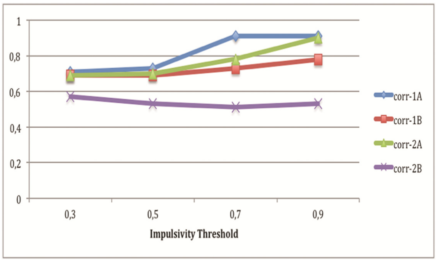

Fig. 3. Agents' correctness in the case 30-30

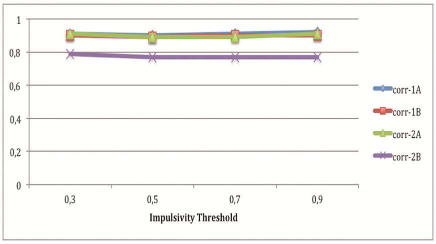

Fig. 4. Agents' correctness in the case 75-75

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 their own expertise in their final evaluations. But not being well informed or experts, there is a higher probability that they will be wrong.

4.2 Second Simulation

In this scenario we suppose that there is a decisional order, so that who is well informed (independently on its expertise) will be the first to decide; then the experts decide; finally the remaining agents decide, the ones with the lowest reliable information. Doing so there is a double effect: who has good information is not negatively influenced by who does not have it; who does not have good information is mainly much more influenced by who has it. The settings are exactly the same of the previous experiment. Let's start analyzing agents' performance.

Looking at Fig. 5, it is clear that the impulsivity has no negative effect on agents that are well informed or expert. This does not imply that they are not impulsive. Let's consider well informed agents: the first to decide will be just influenced by the authority. The others could be impulsive, following the social source. But the social source is just influenced by what the authority reports. In conclusion, the impulsivity factor has no influence in this scenario.

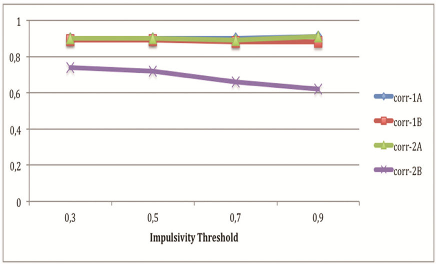

Fig. 5. Agents' correctness in the case 10-10 with decisional order.

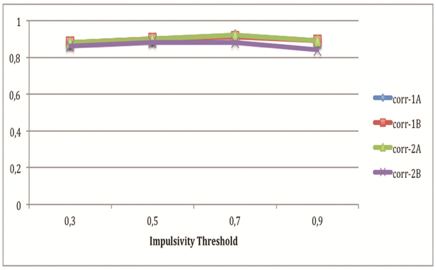

Fig. 6. Agents' correctness in the case 30-30 with decisional order

The impulsivity is still relevant for expert agents, which will be strongly influenced by the well informed ones (but their performance would be high in any case), and it is clearly fundamental for 2B agents (which are both bad evaluators and misinformed): in this case, much more than the agents 1B, impulsivity seems to help them. The proof is given by the fact that their performance decreases while their impulsivity threshold increases.

Obviously, increasing the quantity of information in the world, agents' performance improves and the 2B's curve tends to be equal to the others (Figs. 6 and 7).

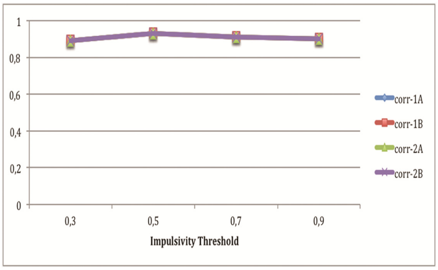

Fig. 7. Agents' correctness in the case 75-75 with decisional order

4.3 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). Figures 8, 9, 10 and 11 show these values respectively to the categories 1A, 1B, 2A and 2B. Not even the decisional order influences them, so that they are the same in both the experiments.

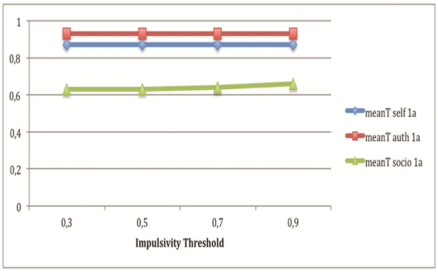

Fig. 8. Trust degrees of the agents belonging to the 1A category in the case 30-30.

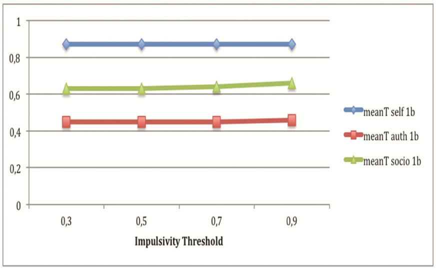

Fig. 9. Trust degrees of the agents belonging to the 1B category in the case 30-30

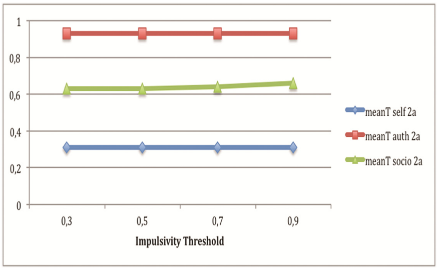

Fig. 10. Trust degrees of the agents belonging to the 2A category in the case 30-30

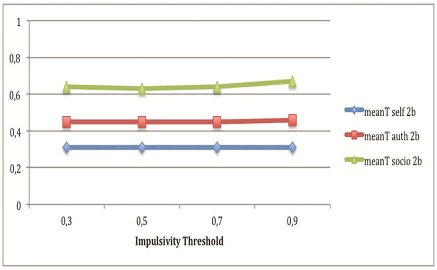

Fig. 11. Trust degrees of the agents belonging to the 2B category in the case 30-30

Of course the social trust changes. Notice that it does not depend on the agent's nature; it just depends on its neighborhood: the more performative they are, the higher the social trust will be. This is clearly visible in Fig. 12, reporting the social trust levels without decisional order. We can see how the social trust increases increasing the percentage of expert/informed citizens.

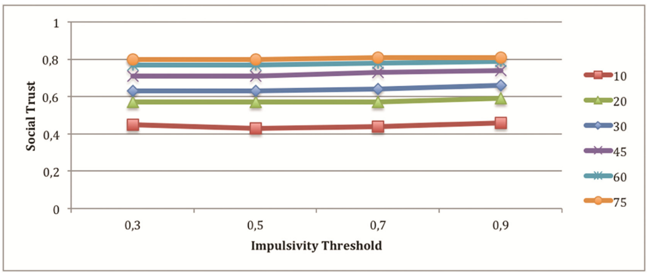

Fig. 12. Social trust of all the agents in the six cases

With the decisional order, there is a big difference for the social source: each category of agents will be differently influenced by the other categories. For instance, well informed agents (1A and 2A) are the first to decide, so that they will be influenced just by agents of the same kind. Then 1B agents decide, exploiting the decision of agents belonging to their category, but also those of well informed agents. 2B agents decide last, on the basis of everyone else's choices. Within this picture we clearly expect to have different trust levels for the different categories.

The trust values of 1A (Fig. 13) and 2A (Fig. 15) agents are almost the same[2]. Even if they have the best performance, their social trust is very low. This is reasonable as quite often they will be the first to decide: in those cases, as none of their neighbors decided yet, their social source is flat, reporting no evidence. Not being able to exploit this source, agents lower the trust in it. Then 1B agents decide, basing its decision on 1A and 2A agents. Their social trust levels (Fig. 14) are higher than those of these last. The higher trust values are the ones of 2B (Fig. 16) agents. In fact, given that all the other agents decided and their decisions were strongly influenced by well informed agents (and then indirectly by the authority), 2B agents are able to exploit a lot of correct decisions.

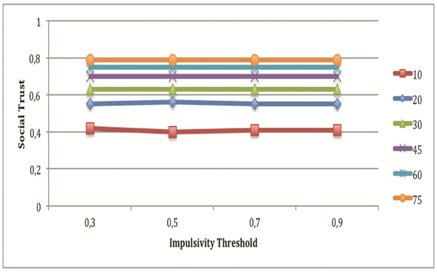

Fig. 13. Social trust of 1A agents in the six cases, with decisional order

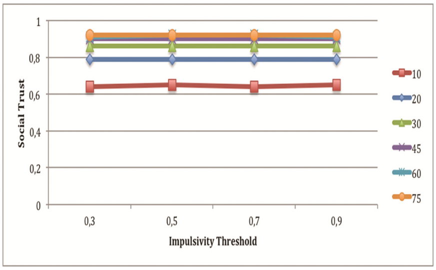

Fig. 14. Social trust of 1B agents in the six cases, with decisional order

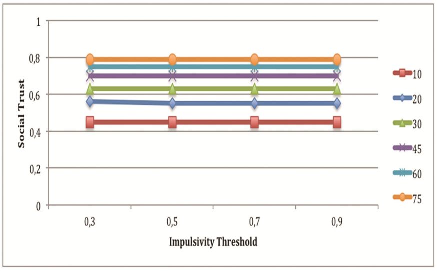

Fig. 15. Social trust of 2A agents in the six cases, with decisional order

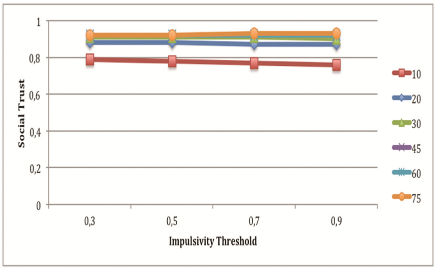

Fig. 16. Social trust of 2B agents in the six cases, with decisional order

[2] Actually there is a little difference: in the case 10-10 the social trust level is a little bit higher in the 2A case. The reason of this difference is the extremely low percentage of 1A agents in this scenario, so that it is quite unlikely that a 1A agent influences another 1A agent, but there is a higher possibility that it influences a 2A agent.

5 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 saw that impulsivity has a strongly negative impact on informed or expert agents, while on the contrary it is useful for the remaining 2B agents. In particular, 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.

Second, we saw that introducing the decisional order the agents' performance improves significantly. Doing so in fact it is possible to avoid the negative effect that impulsivity has on informed or expert agents and to increase the positive effect that it has on 2B agents, even if they represent a substantial percentage of the population.

Finally we analyzed the role played by social trust. In the first simulation, given a value for the impulsivity threshold and a percentage of informed and expert citizens, it assumes a fixed value for all the citizens, as it is independent by the agent's category. On the contrary, introducing the decisional order the social trust changes depending on the agents' category. In particular the first to decide will have the lowest trust value and it increases until the last to decide, 2B agents, just thanks to the decisional order. In this way, the agents that would be hindered by the social source use it less, while the agents that need it to decide correctly exploit it more.

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. J. Argumentation Comput. **5**(2), 191–215 (2014)
2. Barber, K.S., Kim, J.: Belief revision process based on trust: agents evaluating reputation of information sources. In: Falcone, R., Singh, M., Tan, Y.-H. (eds.) Trust in Cyber-societies. LNCS, vol. 2246, pp. 73–82. Springer, Heidelberg (2001). doi:10.1007/3-540-45547-7_5
3. Castelfranchi, C., Falcone, R.: Trust Theory: A Socio-Cognitive and Computational Model. Wiley, Hoboken (2010)
4. Cialdini, R.B.: Influence: Science and Practice, 4th edn. Allyn & Bacon, Boston (2001). ISBN 0-321-01147-3
5. Evenden, J.L.: Varieties of impulsivity. Psychopharmacol. **146**(4), 348–361 (1999). doi: 10.1007/PL00005481

6. Falcone, R., Sapienza, A., Castelfranchi, C.: The relevance of categories for trusting information sources. ACM Trans. Internet Technol. (TOIT) **15**(4), 13 (2015)
7. Genter, K., Stone, P.: 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, May 2016
8. Latané, B.: The psychology of social impact. Am. Psychol. **36**, 343–356 (1981)
9. Melaye, D., Demazeau, Y.: Bayesian dynamic trust model. In: Pěchouček, M., Petta, P., Varga, L.Z. (eds.) CEEMAS 2005. LNCS, vol. 3690, pp. 480–489. Springer, Heidelberg (2005). doi:10.1007/11559221_48
10. Melo, V.S., Panisson, A.R., Bordini, R.H.: 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, vol. 1578, paper 6. Ceur Workshop Proceedings (2016)
11. Parsons, S., Sklar, E., Singh, M.P., Levitt, K.N., Rowe, J.: An argumentation-based approach to handling trust in distributed decision making. In: AAAI Spring Symposium: Trust and Autonomous Systems, March 2013
12. Quercia, D., Hailes, S., Capra, L.: B-trust: bayesian trust framework for pervasive computing. In: Stølen, K., Winsborough, W.H., Martinelli, F., Massacci, F. (eds.) iTrust 2006. LNCS, vol. 3986, pp. 298–312. Springer, Heidelberg (2006). doi:10.1007/11755593_22
13. Sapienza, A., Falcone, R.: A bayesian computational model for trust on information sources. In: Proceedings of the Conferenze WOA 2016, vol. 1664, pp. 50–55. Ceur workshop proceedings, Catania (2016)
14. Schmidt, S., Steele, R., Dillon, T.S., Chang, E.: Fuzzy trust evaluation and credibility development in multi-agent systems. Appl. Soft Comput. **7**(2), 492—505 (2007)
15. Theodorakopoulos, G., Baras, J.S.: On trust models and trust evaluation metrics for ad hoc networks. IEEE J. Sel. Areas Commun. **24**(2), 318–328 (2006)
16. VandenBos, G.R.: APA Dictionary of Psychology. APA, Washington, DC (2007)
17. Villata, S., Boella, G., Gabbay, D.M., Torre, L.: Arguing about the trustworthiness of the information sources. In: Liu, W. (ed.) ECSQARU 2011. LNCS, vol. 6717, pp. 74–85. Springer, Heidelberg (2011). doi:10.1007/978-3-642-22152-1_7
18. Wang, Y., Vassileva, J.: Bayesian network-based trust model. In: IEEE/WIC International Conference on Web Intelligence, WI 2003 Proceedings, pp. 372–378). IEEE, October 2003
19. Wilensky, U.: NetLogo, Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL (1999). http://ccl.northwestern.edu/netlogo/