Interpersonal trust to enhance cyber crisis management

Florent Bollon, Anne-Lise Marchand, Nicolas Maille, Colin Blättler, Laurent Chaudron, Jean-Marc Salotti

HAL Id: hal-02901604
https://hal.science/hal-02901604

Submitted on 17 Jul 2020

Interpersonal trust to enhance cyber crisis management

Florent Bollon[1,2], Anne-Lise Marchand[3], Nicolas Maille[1], Colin Blättler[3],
Laurent Chaudron[4], & Jean-Marc Salotti[2]
[1]ONERA (French Aerospace Lab), [2]IMS UMR CNRS 5218,
[3]CREA (French Air Force Academy Research Centre),
[4]Theorik-Lab
France

Abstract

In the field of cyber-security, software performance optimization is a major focus of research to better prevent cyber threats. However, once threats are detected, they have to be managed by a human operator or more often by human operators' joint actions. The purpose of this study is to show that in these collaborative situations, the interpersonal trust level between these actors shapes their handling of the threat. Forty-five participants performed, with twenty-eight different fictive teammates, a collaborative counting task that included aleatory phases of jamming. Each fictive teammate was described through two adjectives selected to induce a predefined level of interpersonal trust (low or high). The subject and his collaborator worked on different systems with different objects to count and different jamming phases. Nevertheless, each participant had the possibility of supervising his teammate's work by checking out his task and modifying his answers (number of targets and jamming events reported) if required. The subject was responsible for validating the team's final result. The experimental data show that, in this type of collaborative task, the interpersonal trust level has indeed an influence on the supervision strategy used and the team performance.

Introduction

In order to prevent the increase in the number of cyber-attacks, States are setting up cyber operations centers (C2Cyb). The operators of these C2Cybs, who monitor the state of systems and the information flows, are collectively responsible for detecting, correlating and analyzing the various indicators that can *make sense* of a cyber crisis (Boin, Busuioc, & Groenleer, 2014). These indicators, which are difficult to perceive but that predict perturbations in the system, are called *weak signals* (Saritas & Smith, 2011) and are discrete, ephemeral, distributed and difficult to interpret.

In a complex and highly interconnected cyberspace, the collection, detection, analysis and comprehension of *weak signals* requires aggregating information from various actors, both human and material, engaged in monitoring the global system. The amount and complexity of the information available in cyberspace makes it impossible for a single operator to compile all the information in a limited amount of time. The heterogeneous nature of the signals also increases the uncertainty of operators about how to interpret them. As a result, decisions made by the C2Cyb team

leader are based on information that is usually unverifiable and transmitted by his/her teammate. This information can sometimes contradict the leader's information. A question therefore arises: how does the team leader in C2Cyb consider this contradictory information when making decisions in a situation of uncertainty?

The decision-making strategies studied in psychology and economics are sometimes based on theories that adopt probabilistic visions. In particular, the dual-process theory presupposes the existence of two distinct rationality processes (De Neys, 2006; Evans, 2003; Evans, 2011; Kahneman & Frederick, 2007) used in optimizing decision-making. According to this theory, two systems, called system 1 and system 2, coexist. System 1 is a fast, intuitive system that does not require the use of working memory (Evans, 2011). System 2 is used for tasks requiring thoughtful decision-making, and, by extension, a calculation of the probabilities of possible futures generated by the decision. System 2 is slower than system 1 and requires greater cognitive resources and task-specific access to working memory (Evans, 2011). Thus, when a person uses system 2, s/he performs a conditional probability calculation in order to make the best decision.

In the work underlying this theory, the probability distributions of the different options are usually clearly identifiable by the participant, assisting decision-making (Kahneman & Tversky, 1979). However, due to the abundance of information in cyberspace, no probability distribution seems to be applicable by the operator to analyze the veracity and the impact of *weak signals*. In fact, when a team leader has to make a decision, he can only do it based on his own information (the *weak signals* directly perceived) and the information transmitted by his teammates without being able to check it or to compare it with a probability distribution. In these cases, other mechanisms that facilitate decision-making should therefore come in play. Among these mechanisms, trust is often described as a uncertainty reducer (Meyerson, Weick, & Kramer, 1996) that facilitates decision-making (Bell, 1982). This article proposes to study in environments with high uncertainty, what the role of trust is in the leader's decision-making when he cannot verify the data transmitted by his teammate and when these data are different from his own.

Posten and Mussweiler (2019) established a trust predictability function, i.e. trust would allow us to anticipate the possibilities by calculating their probabilities of occurrence. This is what Gambetta indicated (1988: p. 217) when he defined trust as "a particular level of the subjective probability with which an agent assesses that another agent or group of agents will perform a particular action, both before he can monitor such action (or independently of his capacity ever to monitor it) and in a context in which it affects his own action". Gambetta's definition and, more generally, the research conducted in economics (Williamson, 1993) and sociology (Coleman, 1990) link the phenomenon of trust to the notion of probabilistic evaluation and are thus in accordance with the dual-process theory approach. Trust can be considered as the calculation of the perceived cost-benefit (Williamson, 1993) of a relationship. In this calculative approach, "Trust emerges when the trustor perceives that the trustee intends to perform an action that is beneficial" (Rousseau, Sitkin, Burt & Camerer, 1998, p.399). Indeed, trust can only occur in relationships that bring rewards to both

parties (Lewicki & Bunker, 1995) and can be summarized, from an economic perspective, by a probability calculation (Williamson, 1993).

This notion of probability calculation is the link between the literature on trust and the literature on decision-making. In theory, the decision corresponds to "a choice or a set of choices drawn from the available alternatives" (Bellman & Zadeh, 1970). Like trust, decision-making is the choice of the alternative that subjectively presents the best cost/benefit ratio. In this approach, decision-making is no more than the result of a probabilistic assessment of the consequences of different choices (Lowenstein, 2003). In the decision-making process, the trust mechanism could therefore be seen as a readjustment of the probabilities perceived by an operator of the possible futures generated by different options, the option chosen by the operator being the option with the best cost/benefit ratio. This interpretation is consistent with Lewis and Weigert's (1985, p.969) definition of trust when they describe it as "to trust is to live as if certain rationally possible futures will not occur". In teams operating in uncertain environments such as cyberspace where operators cannot assign probabilities about future events generated by a decision (Duncan, 1972), trust may therefore facilitate decision-making. In cases where the leader cannot verify *in situ* the information transmitted by his teammate, and therefore assign a probability as to the accuracy of this information, the level of trust could be a determining factor in decision-making, in particular by facilitating acceptance by the leader of the information transmitted by his teammate. When the level of trust between a leader and his teammate is high, the information provided by the teammate should be perceived by the leader as probably more accurate than when the level of trust is low.

Hypothesis 1: For a team leader, a high level of trust in his teammate leads to a greater acceptance of the unreliable information that the teammate transmits.

In C2Cyb, *weak signals* reported by a teammate are often unverifiable by the leader. This impossibility of verifying the information means that it is impossible for the leader to assign an effective probability to these *weak signals*. When the leader cannot rely on actual probabilities, he has to assign a subjective probability (Kahneman & Tversky, 1972) to these *weak signals*. To do this, he can only rely on his own information, particularly the evaluation of the *weak signals* that he has himself received. He can therefore compare the *weak signals* he has perceived directly with those communicated to him; if all these *weak signals* correspond, they will be considered *consistent*. In this case, the leader should perceive the information transmitted by his teammate as probably more reliable than in the case of *non-consistent signals*.

According to the dual-process theory, in the case of *weak consistent signals*, decision-making is fast and intuitive (system 1). In the case of *non-consistent* signals, because of the necessary probability calculation, the response is slower (system 2) (Hypothesis 2). In this case, when the level of trust between team members is low, if the leader has not perceived any evidence of an attack "directly", he may judge as unlikely the elements in favour of an attack that are provided by the teammate. In other words, a leader will be more inclined to accept the contradiction if he trusts his teammate (hypothesis 3).

Hypothesis 2: *Consistent* signals are processed more quickly by the leader than *non-consistent* signals

Hypothesis 3: The level of trust has an indirect effect on decision-making by modulating the *consistency* consideration

Material and procedure

Method

To test these hypotheses, it is necessary to create an experimental context similar to that faced by cyber leaders. This environment must offer the participant (here, a team leader) a main task and a supervision task on which can be grafted one or more weak signals directly perceived by the leader or transmitted by a teammate. Despite the "*weak*" character, these signals must be sufficiently detectable. The leader has to make a decision based on these *weak signals* that he cannot verify in situations where he has a variable level of trust in his teammate and where these signals are not always *consistent*.

The chosen task fulfils these conditions: it offers the participant a main task of counting aircraft on a photograph with the possibility of checking a similar task with a teammate. The teammate is fictional and only presented by a predefined and controlled level of trust (Bollon, Maille, Marchand, & Blättler, 2019). During this task, "jamming" (see Figure 1) constituting the *weak signals* may occur. The participant has to indicate the number of jamming events without being able to check the number indicated by his teammate. This consideration of the teammate's data corresponds to a "blind" decision. It is this decision that is analyzed in this study and not the decisions related to the main task that can be checked on the teammate's side.
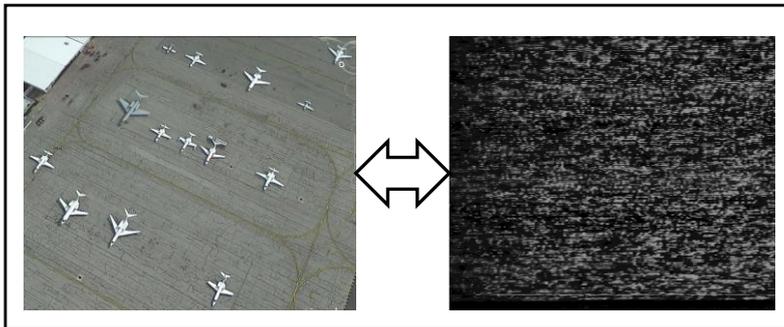

Figure 1. The picture on the left is an example of a photograph used in the experiment; the picture on the right is the jamming that can occur at any time. In the event of jamming, the image on the right appears for one second before disappearing.

In order to test the impact of trust on acceptance of the information transmitted (hypothesis 1), it is necessary to induce different levels of trust in the participant, to check this induction and to test, for each level, the percentage of information transmitted by the teammate and accepted by the participant. The trust-level induction is an independent variable (IV) with two controlled levels (low and high) that will be

called "trust-levels" in the following section of this article. The percentage of information transmitted by the teammate and accepted by the leader (in %) is a dependent variable (DV) collected during the experiment that will be called "decision" in the following section of this article.

In order to test the impact of consistency on the choice of decision system (system 1 or system 2) (hypothesis 2), it is necessary to induce *consistent* and *non-consistent* signals and to compare the time taken by participants to validate a decision according to these signals. The *consistent* or *non-consistent* nature of the signals is an IV which will be called "consistency" in the following section of this article. The *consistency* distribution is controlled by the occurrence of the *weak signals* transmitted. The time taken by participants to validate a decision (in ms) is a DV, called " time ", collected during the experiment.

In order to determine the impact of trust on decision-making during consistent and/or non-consistent events (hypothesis 3), the two IVs explained above as well as the DV "decision" are used.

Participants

45 people (46.6% women and 53.3% men) with an average age of 22.7 years (SD: 1.09%) participated in this study. All participants were second-year engineering students. No participants had any health problems; all had normal or corrected vision.

Protocol

Before the start of the experiment a briefing was carried out, and all participants completed an informed consent sheet. Following this, the participants carried out a 5-minute training session before starting the experiment. The experiment was divided into 28 trials, each with 4 phases. For each trial the participant worked with a different teammate (computer simulated behaviour). 14 trials were performed with a trustworthy teammate (high trust) and 14 trials with a non-trustworthy teammate (low trust). In order to avoid an order effect, teammates' profiles were randomly drawn. All participants therefore worked with all teammate profiles but in a different order.

Participants performed the experiment in groups in computer rooms that did not allow them to see what was happening on the other participants' screens. At each trial, the participant thought s/he was doing the task in collaboration with one of the other participants in the room, although in reality all the teammates were fictitious. Each participant performed the task on an ordinary desktop computer using the keyboard and mouse. The screens of all participants were similar in terms of resolution and brightness.

The task was carried out in 4 phases. The first phase of each trial was designed to introduce to the participant the characteristics of his new teammate who was more or less trustworthy (IV "trust level"). The first display showed a pair of words characterizing this teammate (see Figure 2 "1"). This pair of words allowed the participant to induce a level of trust in his teammate, either low (thanks to rather negative elements: unreliable, disloyal, etc.) or high (thanks to rather rewarding

elements: professional, organized, etc.) (Bollon et al. 2019). These word pairs were obtained by following the protocol described by Bollon et al. (2019) which uses social psychology methods to identify social representations of trust in given social groups. In order to ensure that the participant had taken the teammate's characteristics into account, he was asked, on a second display, to find these two characteristics among 8 distractors (Bollon et al. 2019).
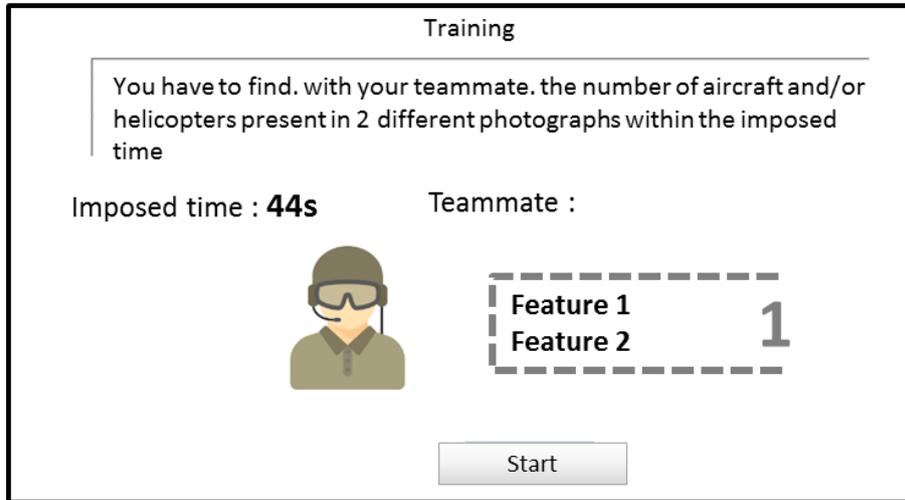

Figure 2 . First display of phase 1. On this display the participant was informed of the instructions (similar throughout the experiment), the time allocated to the task (the time differed depending on the photograph) and the characteristics of his teammate (noted "1" on the image above). These characteristics induced a low or high level of trust in the participant.

The second phase corresponded to the completion of the aircraft counting and jamming counting tasks. The participant had a control display that allowed him to see the countdown of the remaining time as well as the sum of the aircraft counted in the two photographs. This screen contained 4 buttons that allow the participant to (see Figure 3):

1. Display the image on which s/he had to count the aircraft and jamming events
2. Display his teammate's image in order to check the count made by his teammate if necessary
3. Modify the total score, if the participant considered that the number of aircraft counted in the two photographs was not correct
4. Complete this task and move on to the next phase. This button was only active after the participant had validated the number of aircraft present in his photograph.
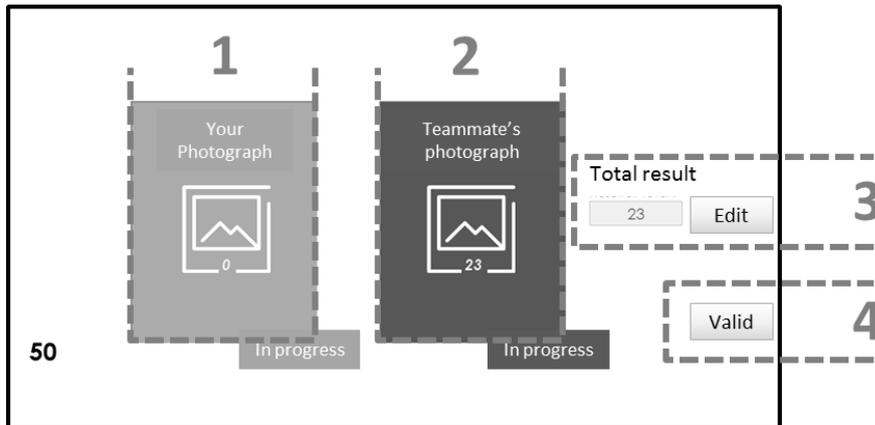
Figure 3. Main display of phase 2. On this display the participant can see the time remaining as well as the number of aircraft counted by his teammate. With the help of different buttons, the participant can access his own image (button "1"), access the image of his teammate (button "2"), modify the total score (button "3") or complete phase 2 (interlocutor "4")

On the display allowing him to perform his own counting task, the participant found his photograph, the remaining time (see Figure 4 "1") as well as 5 buttons that allowed him to:

- Increment or decrement the count by the number of aircraft (see Figure 4 "2"),
- Increase the number of jamming events detected (see Figure 4 "3"),
- Validate his count of the number of aircraft (see Figure 4 "4")
- Return to the control display (see Figure 4 "5").

The teammate's display was exactly the same as the participant's one. However, on the teammate's screen the buttons were not clickable (except for the button used to come back to the control screen). On the teammate's display, the photograph was different from the one presented on the participant's screen and s/he had to do the aircraft and jamming counting tasks on this other photograph. Moreover, on the teammate's screen, it was impossible for the participant to see the jamming (jamming events were never displayed on the teammate's screen). In this experiment, the participant was not aware that it was impossible for him to see the jamming events occurring on the teammate's screen.

The participant had to count the aircraft in his photograph and validate the team's total result before the end of the time limit. If this was not the case, the trial was failed and an additional trial with a teammate of the same level of trust was added at the end of the session. The validation of the total score allowed the participants to move on to the next phase.
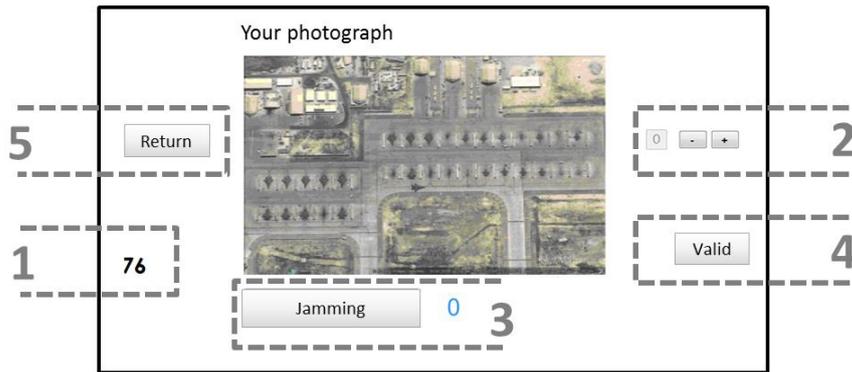
Figure 4. Display used by the participant in phase 2 to perform his counting task. On this display the participant can see the remaining time ("1"). Using different buttons, the participant can count the aircraft (button "2"), increment the interference counter (button "3"), validate his aircraft count (button "4") or return to the main display (button "5") (see Figure 3).

The third phase was devoted to validation by the participant of the jamming events detected on the two photographs. The display showed the number of jamming events detected by the teammate (see Figure 5 "1") and the number of jamming events detected by the participant (see Figure 5 "2"). Because no real jamming was displayed on the teammate's screen, the participant could not see these jamming events and therefore could not assess the validity of the information transmitted by his teammate. Next to each of these numbers, there were 3 buttons to validate or invalidate the jamming (none, 1 or more). The participant had to make a decision on the number of jamming events to validate on the teammate's photograph (see Figure 5 "3") as well as the number of jamming events to validate on his own photograph. Once this was done, the participant could move on to phase 4.

The different DVs used to test the 3 hypotheses were collected in phase 3. The "decision" DV used to test hypotheses 1 and 3 corresponds to the percentage of jamming events transmitted by the teammate and not validated by the participant (in %). The "time" DV used to test hypothesis 2 corresponds to the time taken by the participant to validate this third phase (in milliseconds).

In order to control the IV "consistency", in this experiment, the jamming events presented to the participant were linked to the jamming events transmitted by the fictitious teammate in order to obtain the following 4 cases:

- No jamming was presented to the participant and the number of jamming events detected by the teammate was 0 (25% of cases)
- 1 or 2 jamming events were presented to the participant and the number of jamming events detected by the teammate was 1 or 2 (25% of cases)
- No jamming was presented to the participant but the number of jamming events detected by the teammate was 1 or 2 (25% of cases)
- 1 or 2 jamming events were presented to the participant but the number of jamming events detected by the teammate was 0 (25% of cases)

Figure 5. Display used by the participant in phase 3. On this display the participant can see the number of jamming events detected by his teammate ("1") and the number of jamming events he had himself indicated ("2"). The participant had to make a decision on the number of jamming events to be validated for the participant (button "3") and for himself before he could complete phase 3 by clicking on the validation button (button "4").

The first two cases were the so-called *consistent* cases and the other two *non-consistent* cases.

Finally, Phase 4 was the subjective assessment of the participant's level of trust in the results (number of aircraft) reported by his teammate. The purpose of this evaluation on non-segmented scales was to verify that the experimental trust induction equipment was working well and that the participant was working with teammates whom he perceived as trustworthy and others as less trustworthy (Bollon et al., 2019). As a high level trust induction should lead to a higher subjective evaluation by the participant of his teammate's performance than a low level trust induction (Dirks & Ferrin, 2001), the smooth operation of the experimental protocol should therefore lead the participant to assign a high evaluation to teammates in whom he had high trust and a lower one to teammates in whom he had less trust.

Results

Data from the 45 participants were included in the analysis. Before analysing the results required for hypothesis testing, the verification of the induction of trust in the experimental protocol was performed. The subjective evaluation data of the results transmitted by the teammate (recovered in Phase 4) show that when the trust level was high ($M = 5.68$, $SD = 2.33$) the subjective evaluation of the teammate's performance seem to be higher than when the trust level was low ($M = 5.34$, $SD = 2.44$). In order to validate these results, a one-way repeated measure ANOVA, with the IV "trust level" as a factor, has been carried out. The significant results ($F(1,44) = 4.11$, $p = .04$) validated the presence of two levels of trust (high and low).

In order to test the hypothesis that a high level of trust between team members leads to greater acceptance by the leader of the information transmitted by his teammate (hypothesis 1), the DV "decision" and the IV "trust level" were used. For each participant, the data obtained were averaged, for each level of trust. The data indicate (see Figure 6) that between the low trust level ($M = 23.4\%$, $SD = 35.7\%$) and the high trust level ($M = 21\%$, $SD = 33\%$) the performances are relatively similar. A one-way repeated measure ANOVA, with the IV "trust level" as a factor, has been carried out. The insignificant results ($F(1,44) = 1.1$, $p = .30$) do not support hypothesis 1. In other words, trust does not seem to have a direct effect on the validation of the results reported by the teammate.
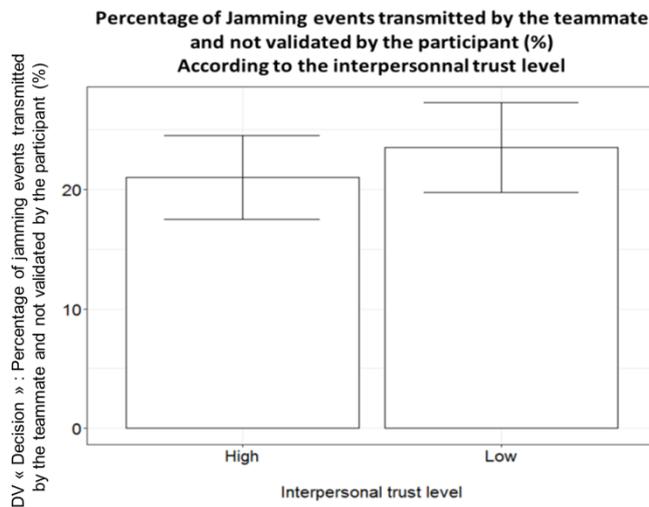

Figure 6 . Percentages of jamming events transmitted by the teammate and not validated by the participant according to the trust level

In order to test the hypothesis that consistent signals are processed faster by the leader than non-consistent signals (hypothesis 2), the "time" DV and the "consistency" IV were used. For each participant, the data obtained were averaged, for each level of consistency. The results show that when the jamming events were *consistent* ($M = 3750.5$ ms, $SD = 1192.4$ ms) the participants seem to validate phase 3 more quickly than when the jamming events were *non-consistent* ($M = 4272.7$ ms, $SD = 1515.5$ ms). In order to validate these results, a one-way repeated measure ANOVA, with the IV "consistency" as a factor, has been carried out. The significant results ($F(1,44) = 13.37$, $p < .001$) validated hypothesis 2. It would seem that the participants had a different perception of the consistency of the signals.

In order to test the hypothesis that the trust level has an indirect effect on decision making through the modulation of the *consistency* consideration (hypothesis 3), the DV "decision", the IV "trust level" and the IV "consistency" were used. For each participant, the data obtained were averaged, for each trust level, according to their *consistency*. The data show (see Figure 7) that when the information transmitted by the teammate was in line with the event perceived as the most likely (consistent case)

the leader seems to validate the information transmitted by his teammate, irrespective of whether the teammate was associated with a high ($M = 21.3\%$, $SD = 30.8\%$) or low ($M = 18.5\%$, SD = 32%) trust level. However, when the information transmitted by the teammate supported an event perceived as unlikely (non-consistent cases), when the trust level was high ($M = 20.6\%$, $SD = 35.4\%$), the leader seem to validate the information transmitted by his teammate more easily than when the trust level was low ($M = 28.4\%$, $SD = 38.7\%$). In order to validate these results, a two-way repeated measure ANOVA, with the IV "trust level" and the IV "consistency" as a factor, has been carried out. The results of the ANOVA showed an interaction effect ($F(1,132) = 4.86$, $p = .02$ (eta-squared =.068))). A post hoc analysis performed with a Tukey HSD test indicated a significant difference in trust levels for *non-consistent* trials ($p = .02$) and no difference for *consistent* trials ($p = .40$).
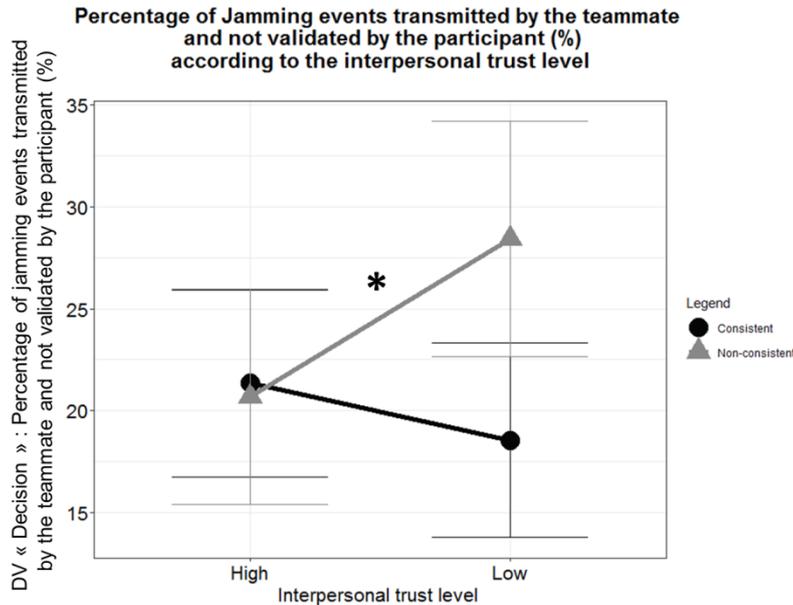

Figure 7. Percentages of interference transmitted by the teammate and not validated by the participant according to the trust level and the consistent or non-consistent character of the tests

Discussion

This study has investigated the relationship between interpersonal trust and decision-making in uncertain environments. On the basis of the dual-process theory (De Neys, 2006; Evans, 2003; Evans, 2011; Kahneman & Frederick, 2007), it is expected that decision-making can be supported either by a rapid and intuitive mechanism (system 1) that requires few resources or by a slower mechanism (system 2) (Evans 2011) involving an assessment of probabilities in relation to the possible situations, risks and benefits of certain alternatives. Applied in a micro-world resulting from cyber crisis management, the experiment aimed to better understand the impact of trust between operators and the consistency of the information exchanged on the decision-making

mechanism (through the time taken to complete the task), but also on the decision itself (validation of the partner's response).

The results show that trust does not directly impact decision-making when it is made on unverifiable elements (hypothesis 1). This result seems to contradict existing models that link trust and decision-making (Kim, Ferrin, & Rao, 2008). However, the current literature studies trust in collaborative tasks where participants can at least access the teammate's work to assess it (Bollon et al., 2019; Dirks, 1999), while the protocol presented here proposes a "blind" decision. It seems necessary to further study this type of situation and its impact on trust. On the other hand, the *consistency* of the information exchanged directly modifies the time taken to take the decision (hypothesis 2). In other words, the *consistency* of the elements exchanged between operators appears to be the primary criterion that determines the mechanism underlying the decision-making process. Once system 1 or 2 has been chosen, trust comes into the decision to the extent that the system 2 leader agrees more with the teammate's result when he or she has trust even if the information given is *non-consistent*. Once system 1 or 2 has been chosen, trust becomes an important factor in the decision-making. In fact, the leader in system 2 accepts the teammate's result to a greater extent when he trusts him even if the information given is *non-consistent*. (Hypothesis 3) (*see* Figure 8).
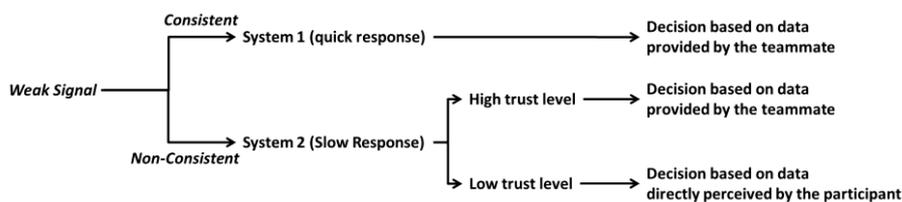

Figure 8 . When the weak signals directly received by the leader and the weak signals transmitted by the teammate are perceived as consistent, decision-making is fast and intuitive (system 1) and is independent of the trust level. However, in the case of weak signals perceived as non-consistent, decision-making is slower (system 2) and involves the trust level. When the trust level is high, the leader's decision-making is in line with the information provided by the teammate and when the trust level is low, the decision making is in line with the information he has himself perceived.

Thus, the study shows that the result of the decision, in terms of the acceptance or non-acceptance of the teammate's information, is linked both to the consistency of the information transmitted and to the level of trust between the operators. When the information received is consistent with the teammate's observations, decision-making is intuitive and not linked to the level of trust between operators and all information is accepted by the leader. On the other hand, when the information is non-consistent and the leader uses system 2 to make his decision, then the level of trust in the teammate who gave him the information can change the decision; the more trust the leader has in his teammate, the more inclined he is to accept his information, whether the latter confirms or invalidates his observations. The level of trust appears therefore to have a significant impact on the probability that the leader will associate with the information received, which the literature has suggested since Gambetta's (1988) work.

The direct impact of consistency is significant in the implementation of C2Cyb. Indeed, it is important in these safety-critical operations to better understand what can impact the way decisions are made. This can make it possible to adapt operator training by making them aware of the effect of consistency on their decision-making (rapid decision versus rational decision). These results can also contribute to a better understanding of how information is presented on the interfaces in order to help in better decision-making.

In terms of trust, the experiment shows that in the context of a decision made by assessing the risks or costs associated with each choice, trust in the source of the information changes the decision. This result is also important from an applicative point of view because it shows that some weak signals sent back to the decision-maker could be taken into account differently in the decision depending on the relationship between the people. Trust between people therefore changes the trust placed in the data itself. It will therefore be important for socio-technical systems such as C2s to take this dimension into account to optimize its effect on the functioning of the system.

One of the methodological contributions of this study is that we have confirmed experimentally the implementation of different decision-making mechanisms according to consistency, in accordance with the dual process theory. In other words, this *micro-world* may affect decision-making in either system 1 or system 2. However, the protocol used does not make it possible to check whether the time delay observed as a function of consistency corresponds to a probability calculation. A future study should make it possible to test this probability calculation by detailing how the decision-making process is carried out. It could use this micro-world to better understand the cognitive mechanisms really at work in each strategy.

This study considered two factors, consistency and trust, which combine to modulate the decision-making mechanism and decision content in collaborative activities. It would now be appropriate to investigate how these results are related to the interaction between human operators or whether they are more general. Are the mechanism and decision similar if the operator acts in cooperation with an automated system or artificial intelligence?

References

Bell, D. (1982). Regret in Decision Making Under Uncertainty. *Operations Research, 30*, 961–81.

Bellman, R.E., & Zadeh, L.A. (1970). Decision-Making in a Fuzzy Environment. *Management Science, 17*(4), B141–B164.

Boin, A., Busuioc, M., & Groenleer, M. (2014). Building European Union capacity to manage transboundary crises: Network or lead-agency model? *Regulation & Governance, 8*(4).

Bollon, F., Maille, N., Marchand, A., & Blättler, C. (2019). Cyber-attaques : Organiser la confiance. *Dixième colloque de Psychologie Ergonomique*, (pp. 243–250). Arpege Science Publishing.

Coleman, J.S. (1990). *Foundations of Social Theory*. Harvard University Press.

De Neys, W. (2006). Dual Processing in Reasoning: Two Systems but One Reasoner. *Psychological Science, 17*, 428–433.

Dirks, K.T. (1999). The effects of interpersonal trust on work group performance. *The Journal of Applied Psychology, 84*, 445–455.

Dirks, K.T., & Ferrin, D.L. (2001). The Role of Trust in Organizational Settings. *Organization Science, 12*, 450–467.

Duncan, R. (1972). The characteristics of organizational environments and perceived environmental uncertainty. *Administrative Science Quarterly, 17*, 313–327.

Evans, J.S.B. (2003). In two minds: dual-process accounts of reasoning. *Trends in Cognitive Sciences, 7*, 454–459.

Evans, J.S.B. (2011). Dual-Process Theories of Reasoning: Contemporary Issues and Developmental Applications. *Developmental Review, 31*, 86–102.

Gambetta, D.G. (1988). Can we trust trust? In D.G. Gambetta (Ed.), *Trust* (pp. 213-237). New York: Basil Blackwell.

Kahneman, D., & Frederick, S. (2007). Frames and brains: elicitation and control of response tendencies. *Trends in Cognitive Sciences, 11*, 45–46.

Kahneman, D., & Tversky, A. (1972). Subjective probability: judgment of representativeness. *Cognitive Psychology, 3*, 430–454.

Kahneman, D., & Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica, 47*, 263.

Kim, D.J., Ferrin, D.L., & Rao, H.R. (2008). A trust-based consumer decision-making model in electronic commerce: The role of trust, perceived risk, and their antecedents. *Decision Support Systems, 44*, 544–564.

Lewis, J.D., & Weigert, B. (1985). "Trust as a Social Reality." *Social Forces, 63,* 967–985.

Lewicki, R.J., & Bunker, B.B. (1995). Trust in relationships. *Administrative Science Quarterly*, *5*, 583-601

Lowenstein, G. (2003). *The Role of Affect in Decision Making*. Handbook of Affective Sciences, (pp. 619–642).

Meyerson, D., Weick, K.E., & Kramer, R.M. (1996). Swift trust and temporary groups. In *Trust in organizations: Frontiers of theory and research* (pp. 166–195).

Posten, A.C., & T Mussweiler. T. (2019) Egocentric Foundations of Trust. *Journal of Experimental Social Psychology, 84*, Article 103820. https://doi.org/10.1016/j.jesp.2019.103820

Rousseau, D.M., Sitkin, S.B., Burt, R. S., & Camerer, C. (1998). Introduction to Special Topic Forum: Not so Different after All: A Cross-Discipline View of Trust. *The Academy of Management Review, 23*, 393–404.

Saritas, O., & Smith, J. (2011). The big picture-trends, drivers, wild cards, discontinuities and weak signals. *Futures, 43*, 292–312.

Williamson, O.E. (1993). Calculativeness, Trust, and Economic Organization. The Journal of Law and Economics, 36 (Part 1, Part 2), pp. 453–486.