Studying Citizens' Trust to Monitor Measures Acceptance during COVID-19 Pandemic

Alessandro Sapienza and Rino Falcone
Institute of Cognitive Sciences and Technologies,
ISTC-CNR
Rome, Italy, 00185
alessandro.sapienza,rino.falcone@istc.cnr.it

Abstract

Over the last year, the COVID-19 pandemic has strongly affected everyone's lives. An unprecedented situation, which required enormous sacrifices and very stringent limitations. Within this context, trust has played a crucial role: people decided to trust their institutions to tackle the pandemic and that trust made the strong restrictive measures effective. This work aims to study the response of the Italian population to the early stages of the pandemic. Making use of a survey addressed to 4260 Italian citizens, we realized an agent-based simulation to model and analyze citizen trust starting from its main cognitive sub-components, with particular reference to the dimensions of competence and willingness. The results of this work can be of great interest, both for understanding what happened in the past, but also for designing effective strategies in the future.

Keywords: COVID-19, Trust, Multi-agent system, Social Simulation

1 Introduction

The first case of COVID-19 was reported in Wuhan, China in December 2019 [Li20] which then spread around the world. In the last months, global health authorities approved a few vaccines and vaccination has already started worldwide. Unfortunately, it is not yet clear how long it will take to make these vaccines available to the vast part of the world population. Since, this problem seems far from being solved[Bon20], there is no other choice but to prepare to live with the virus for a while longer.

Italy was affected by COVID-19 in late February, being the first European country. After a a few weeks, the pandemic spread throughout the rest of the country [Tui20]. The only measures available to deal with the emergency involved highly restrictive social distancing measures [Bak20, Pre20, Jef08], as pandemics spread exploiting the social and community tendency of individuals [Zng20]. However, establishing costly and drastic measures would not have been enough, as societal safety was strongly depending on the commitment of individual citizens[Cai20, Fll20]. In fact, the efficacy of the containment measures relies almost entirely on the citizens' willingness to cooperate with institutions and such measures would have been vain if a significant percentage of the population had considered them inadequate or inappropriate, not respecting them.

In: R. Falcone, J. Zhang, and D. Wang (eds.): Proceedings of the 22nd International Workshop on Trust in Agent Societies, London, UK on May 3-7, 2021, published at http://ceur-ws.org

Such considerations suggest the need to investigate what happened during the first wave, in order to identify what led to this high acceptance rate. This analysis could help the authorities realize better intervention policies and communication strategies in the future [Pag20].

As it is well known, citizens' trust in public institutions has a key role during pandemics[Sie14], influencing people's willingness to adopt recommended behaviour [Lew20]. Indeed, in the first pandemic wave, the Italian population decided to trust its institutions to face COVID-19[Fal20a] and this has actually led to an exceptionally positive result.

As [Fel20] stated, understanding how trust is formed and when it breaks down is a key aspect in understanding large social behavioural trends during this time There is a clear need for a comprehensive theory and a model of the citizens' decision-making process. We need a detailed analysis of trust and its cognitive components, to identify what elements stimulated trust and how trust actually led citizens to compliance. Given the limited resources, in fact, public authorities need to choose where to focus their efforts. Knowing which cognitive components should be stimulated to influence at best citizens' opinion - and, consequently, their behaviour - would be of great help.

The tool of simulations has been widely used to study this pandemic. It proved to be particularly useful for reducing the impact of COVID-19 [Cur20], supporting different decisions that arise during a public health emergency such as this one. Indeed, over the past months, many contributions studied the pandemic using SIR models [Ana20, Bar20, Fan20] or agent-based simulation [Man20, Cre20, Hun20, Roc20, Sil20] focusing on the spread of the contagion. If, on the one hand, these issues have widely been addressed, on the other hand, much less attention has been given to simulating citizens' acceptance of the restrictive measures. The measures necessary to reduce the infections and damages of the pandemic have been clearly detected, yet a detailed analysis of the conditions leading to compliance is still lacking. At this very moment, we believe that such analysis is fundamental to be better prepared for the future, exploiting the huge amount of data we already gathered in the past months.

Presenting this study, we mean to contribute to the vast body of knowledge on the COVID-19 pandemic, using agent-based simulation to model citizens' decision-making process and, then, analysing the reaction of the whole community. Specifically, on the basis of a survey addressed to 4260 Italian citizens, we applied linear regression as an instrument to reproduce citizens' trust and decision-making model from their main cognitive sub-components.

In particular, we aim to contribute to the assessment of public policies to control the incidence of COVID-19 in two ways. First of all, we produced a detailed analysis of what happened during the first pandemic wave, investigating the behaviour of the Italian population. Secondly, based on this detailed analysis, we are able to make predictions about possible future behaviours, by hypothesizing scenarios of new pandemic waves.

2 Survey and Sample

The survey [Fal20a] was conducted on Italian citizens (N = 4260) between 9 March and 14 March 2020, using a snowball sampling method to determine the respondents. The main characteristics of the sample are synthesized in Table 1. It should be noted that the mean educational level of participants is very high: almost three quarters of respondents have a degree (38%) or post-graduate specialization (34%).

The questionnaire was aimed to investigate the participant's overall trust towards public authorities, along with the factors beyond this trust. The questionnaire was based on the socio-cognitive model of trust developed by [Cas10] and explored participants' opinions on five main dimensions, in relation to the current COVID-19 crisis:

1. Evaluation of the *competence* of public institutions;
2. Evaluation of the *willingness* of public institutions;
3. *Purposes* and *effectiveness* of the public institutions' intervention;
4. Trust and *information sources*: the most used sources of information and their perceived trustworthiness;
5. Expectations about the *future scenarios* that will arise, once the COVID-19 crisis is over.

Within this work, we focus on the first three points, considering all the cognitive items related to them. The questionnaire fully complied with ethical guidelines for human subject research and participation was conditional

Table 1: Sample characteristics

	Regions most affected % (30%)	**Regions less affected % (70%)**	**Total %**
Gender			
Men	45	42	43
Women	55	58	57
Total	100	100	100
Age (Mean=46)			
18-29	19	11	13
30-39	23	18	19
40-49	23	24	24
50-59	21	28	26
60-69	11	15	14
>70	3	4	4
Total	100	100	100
Educational Level			
Middle school	3	2	2
High school	24	27	26
University degree	41	36	38
Post-graduate specialization	32	35	34
Total	100	100	100
Geographical provenance			
Northern Italy	96	7	33
Central Italy	4	53	39
Southern Italy/islands	0	40	28
Total	100	100	100

on the preliminary approval of an informed consent by each subject; the compilation took an average time of 10 minutes.

3 Trust Model

3.1 Theoretical Formulation

We aim to identify the complex relations between trust and its components, with respect to the context of the COVID-19 pandemic. The model of trust we realized was inspired by the socio-cognitive model developed by [Cas10], the same exploited in the original survey [Fal20a]. Looking at the main predictors of trust in order to perform regression analysis, we chose this approach because it provides a rich and nuanced description of the various cognitive components of trust. Therefore, we realized an instance of this theoretical model, related to the specific context of interest.

We consider a population composed by n citizens $\mathcal{P} = \{c_1, \ldots, c_n\}$ and a public authority PA. Upon the occurrence of a pandemic C, the citizens need to find someone reliable to carry out a specific task τ, which implies the identification of the necessary rules and restrictions to face the pandemic threat, ensuring the well-being of citizens.

$$\tau = (\alpha, g) \tag{1}$$

In 1, α represents the action. In this specific case, α implies "to establish the necessary measures". Actually, this is not a simple action, but rather a whole class of actions. Concerning g, it is the goal of the task, which in this case is "ensuring the well-being of the whole community, thus including c_i".
Given the nature of the problem and of the domain we are considering, there is only one potential trustee able to carry out τ, precisely PA. Such role is well defined within the society, and it is established that who is holding it must also perform τ. Thus, citizens evaluate PA and consider how appropriate it is for τ, within the COVID-19

pandemic context C. Referring to the socio-cognitive model, trust is formulated as in equation 2.

$$\text{Trust}(c_i, PA, \tau, C), \forall 1 < i < n \tag{2}$$

The citizen c_i will actually decide to trust and to rely on the authority PA to carry out the task τ (which, in turn, implies the execution of the action α to realize the goal g) just if the trust evaluation on PA is good enough, i.e. above a certain threshold of acceptance σ_i.

Putting the theoretical framework into practice, it is necessary to investigate the specific features of trust. In fact, as it is well known in the literature [Cas00, Ram04], trust is made of different cognitive sub-components. Among the many, two assume a particular importance: *competence* and *willingness*. Competence represents the set of qualities making the trustee good for the task τ: skills, know how, expertise, knowledge, self-esteem, self-confidence, and so on. Willingness, the second fundamental dimension, is a prediction of the trustee's behaviour: the fact that it is reliable, predictable, that we can count on it. Willingness evaluates weather the trustee is willing (it really has the intention to do α in order to realize the goal g) and persistent [Cas98, Fal01] (it is not prone to give up, but it will insist on achieving g). Many other beliefs, depending on different kinds of delegation and different kinds of agents, enrich and support these two components of trust as an attitude towards the trustee. Nevertheless, competence and willingness represent the real cognitive kernel of trust. With respect to the context we are interested in:

1. the trustor c_i has to believe that the trustee PA has the ability to realize $\tau => Bel_{c_i}(CanDo_{PA}(\tau))$
2. the trustor c_i has to believe that PA is really willing to realize $\tau => Bel_{c_i}(WillDo_{PA}(\tau))$

According to Castelfranchi and Falcone, competence and willingness are the necessary components of trust and trustworthiness. This does not imply that, in order to believe that PA is trustworthy (and possibly to decide to trust it), c_i should necessarily have a good evaluation of its competence and willingness. Trust, as an evaluation, is not a yes/no status; c_i's trust in PA (evaluation of trustworthiness) must be just sufficient enough to take a risk on it. For instance, it may happens that the competence has a negative evaluation, but the overall evaluation of the trustee is positive.

Of course, internal change in the trustee (new goals, the need to redirect funds, etc.) or external inputs (rumors about what is going to happen, dissemination of fake news, etc.) may affect these two components and, in turn, the evaluation of citizens[Bun20]. We are not interested in modeling the process of knowledge acquisition and revision of citizens and how this affects trust [Fal20b]. Instead, within this framework, we consider a higher level representation. Therefore, we focus only on the actual variation of competence and willingness, without deepening the reasons for this change.

Notice that, in this specific scenario, if c_i actually decides to rely on PA for τ, as a consequence, it chooses to accept the rules and the restrictions that PA established. In other words, in order to successfully pursuit τ, in turn, PA will assign a further task τ_1 (equation 3) to citizens:

$$\tau_1 = (\alpha_1, g_1) \tag{3}$$

In equation 3, the action α_1 is "to follow the restrictions that PA established", while the goal g_1 is a sub-goal necessary for achieving "the well-being of the whole community" (respect for spacing, cleanliness of hands and wearing the mask will result in partial acquisition of the overall goal of defending public health). In fact, only maximizing the number of citizens adopting g_1 it will be actually possible to realize the overall goal g, i.e. ensuring the well-being of the whole community. While the authority PA explicitly aims to protect the whole community (g as result of the sum of all g_1), the primary purpose of the citizen c_i is, usually, to protect itself (indeed, it may also understand that it can contribute to achieve g, which is the only way to really save itself). Maybe c_i does not believe g_1 or g either. For the sake of completeness, we care to underline that other different reasons may induce citizens to adopt the goal g_1, even if they do not possess g. As an example, compliance may aim at achieving other goals, such as avoiding fines. But there may also be more elaborate motivations behind c_i's choice, such as compliance to the norms governing the society (the pact at the basis of common coexistence in an organized society). The citizens c_i may not agree with g_1, but it could still decide to adopt g_1 (and as a consequence g) in order to respect social norms.

Besides competence and willingness, many other cognitive variables may come into play during the trust assessment process. For instance, the subjective importance of the task τ_1, i.e. how important τ_1 is for the

Table 2: Linear regression coefficients for trust on PA

	Coefficient
Competence	0.3385
Willingness	0.2558
PA's goal: to contain the contagion	0.2149
Importance of own contribution	0.1892
Expectation on others' behaviour: full compliance	0.1169
k	-0.8646

trustor. Notice that τ_1 does not necessary imply a direct utility, since following the rules could not be directly useful to the specific citizen, but rather to the whole community. Carrying out τ_1 is a necessary precondition to realize g. Nevertheless, the actual achievement of this goal is related to the execution of τ_1 by a significant part of the population. The intertwined nature of the relationship between g and g_1 (respectively τ and τ_1) introduces many other cognitive variables, refining the decision-making phase: if the trustor perceives its contribution as useful for the community and/or for its relatives; if it believes that a significant part of the population will actually carry out τ_1 and then it will be possible to realize g; if it believes that an insufficient number of people will follow the rules, thus it is not worth accepting these restrictions. Although all these variables will be used to shape trust, within the simulation we will mainly focus on two specific dimensions: competence and willingness.

3.2 Linear Regression on trust

Linear regression is a linear approach to modelling the relationship between a dependent variable and one or more explanatory variables. The advantage of such approach is that, unlike in correlation analysis, it does not simply investigate the relationship existing between two variables, telling us how much one affects the other. On the contrary, this instrument takes into consideration the entire picture, assessing the interaction of all variables to determine their overall effect on trust. This tool fits perfectly with the considered problem since, by the means of appropriate weights, it relates trust to its different cognitive sub-components. Specifically, for each citizens c_i, we are interested in a function modelling its trust evaluation as in equation 4:

$$\text{Trust}(c_i, PA, \tau, C) = k_i + \sum_{h=1}^{q} w_h * \text{x}_\text{h}(c_i, PA, \tau, C) \tag{4}$$

While we already have the values of the different variables x_1, x_2, ... , x_q, extrapolating them from the cited dataset, linear regression provides us with weights w_1, w_2, ... , w_q and constant k_i.

We applied linear regression to the whole sample, estimating a unique trust function for all the population. Data analysis was performed using Weka (version 3.8.3) statistical software [Wit02].

Table 2 reports all the weights used in linear regression. These data confirms a clear and strong link between trust and the dimensions of competence and willingness. In addition, the difference in weight between competence and willingness appears clear. Overall, citizens consider competence as more important.

The correlation coefficient is very high (0.7405). Concerning the MAE, it hovers around 0.5 (0.4788). This value appears more than reasonable, given that trust functions range in [1, 5]. Moreover, while the subjects could express their evaluation in natural number $\mathbb{N}$, thus excluding the intermediate values, trust functions work in $\mathbb{R}$. Such values suggest that the defined functions represent a quite well approximation of citizens' trust assessment. Nevertheless, our goal is not estimating trust of specific individuals, but that of the whole community. At community level, MAE is significantly reduced: the estimation of the average community trust is almost exactly equal to the actual average trust of the sample. In other words, trust estimation at community level is much more precise ($MAE < 0.01$) than that on individual subjects.

3.3 Linear Regression on the acceptance threshold

In their theoretical model, Castelfranchi and Falcone state that trust as an evaluation can be defined as a degree. However, producing an evaluation is only the first step. We are also interested in studying the following cognitive

phases. We would like to understand when and how this evaluation becomes a decision and then an actual action, the act itself of trusting and relying on the trustee. As highlighted in [Cas10], trust as a decision or action cannot be defined as a degree, but it is a boolean variable: true/false, to trust/not to trust.

If the trustor perceives the trust evaluation as good enough (above a certain internal threshold), then it decides to trust. Otherwise, it excludes this possibility. Referring to the context of our interest, after the evaluation, the citizen c_i compares $\text{Trust}(c_i, PA, \tau, C)$ with an internal threshold σ_i. Such a comparison will determine the final decision of c_i. As σ_i is an internal property of the citizen c_i, it is not given in advance. Nevertheless, it is possible to estimate it. First of all, we care underlining that the trust assessment process we introduced does not simply evaluated an a-priori trust, an abstract evaluation of the trustee PA. On the contrary, this evaluation takes into account all the specific cognitive dimensions related to the task and the context (perceived importance of the task, expectation on others' behaviour, additional reasons for trusting PA, and so on). Since all these dimensions have already been weighed, with a first acceptable approximation, we may suppose that σ_i is around 0.5 (the half of the interval in which the trust value is included). In other words, if the trust evaluation is greater than half, c_i decides to trust PA. Such statement represents a reasonable choice, both qualitatively and quantitatively, since a trust value over 0.5 actually represents a situation in which the reasons for trusting are more than the reasons for not doing so.
As a second approach, which we actually implemented in this work, it is possible to estimate the actual value of σ_i, on the basis of the answers of the subjects. Practically speaking, we are interested in identifying a trust threshold such that above this value we may reasonably assume c_i effectively trusts PA. In other words, c_i considers useful to follow the rules determined by PA. Then, making again use of linear regression, we linked trust in PA for COVID-19 management with the evaluation of the measures to face this emergency. Then, we exploited these functions to determine the acceptance threshold. Basically, we estimated we have estimated how much trust the citizens need to perceive the measures at least as enough useful (namely, greater than 3 in Likert scale, or 50%). Remarkably, results from linear regression analysis show that σ values hover around 50% (51.83%, 3.0733 in Likert scale). This is actually very interesting, as there is an actual correspondence between the threshold we deduced and the one we computed.

4 Simulations

We developed an agent-based model for a fine-grained computational simulation of the institutional trust of Italian people during COVID-19 pandemic. Such framework can be used to explain how and to what extent the individual sub-components have affected trust and, above all, to understand how trust evolves when these components change. The model has been developed in the environment of NetLogo 5.2 [Wil99].

Within the simulations, we use as input the already cited corpus. Given that we can exploit 4260 subjects, we considered as many agents in the simulation. We used the items corresponding to the subjects' answers as variables to describe and characterize the different agents.
We proceed estimating individuals' trust on PA, using the framework previously described. Thus, $\text{Trust}(\mathcal{P}, PA, \tau, C)$ is computed as the mean of individuals' trust evaluation (equation 5). Given that the population $\mathcal{P}$ is made of n citizens

$$\text{Trust}(\mathcal{P}, PA, \tau, C) = \frac{\sum_{j=1}^{n} \text{Trust}(c_j, PA, \tau, C)}{n} \tag{5}$$

The framework we defined can be used to verify the evolution of trust when its different cognitive components change. In particular, in this study we focused on the two dimensions of competence and willingness, keeping fixes the values of the other parameters. Given that the subjects were asked to answer in a 5-point Likert scale, we also considered the variation of competence and willingness with respect to the same scale. Instead, trust values will be reported both in percentage and in Likert scale.

Of course, within the framework we defined, it would be theoretically possible to consider different dimensions, which would have a more or less significant impact on trust. Nevertheless, our choice fell precisely on competence and willingness. The reasons behind this choice are basically two. First of all, Castelfranchi and Falcone themselves identify these component as the most important for trust, being part of what they defined "core trust". Indeed, the analysis conducted on the subjects confirms their theory, as competence and willingness are two of the components that most correlate with trust: respectively ($R = 0.587$, $p < 0.0001$) and ($R = 0.551$, $p < 0.0001$). In turn, this implies that their variation has also a higher impact on the final trust value.

Secondly, it has to been considered that PA can directly act on these variables, in order to modify its actual trustworthiness and citizens' perception of its competence and willingness. This situation does not hold for the other variables. Certainly, the authority PA may still decide to work on other dimensions affecting citizens' trust evaluation (for instance importance of own contribution, expectation on others' behaviour, etc.). However, this would be a much more complex way to (possibly) achieve the same result. Even managing to indirectly manipulate such variables, this would still have a lower impact, as these secondary variables have a weaker correlation with trust.

5 Results

Below, we start analysing citizens' reaction in the first pandemic phase. After that, we take into consideration a series of possible hypothetical scenarios, in which the values of competence and willingness change.

5.1 The first pandemic episode

In the first instance, the realized platform can be used to analyse citizens' reaction during the first pandemic episode. As far as it concerns the average trust of the whole population on PA, we estimated Trust$(\mathcal{P}, PA, \tau, C)$ = 72.35%. This is definitely a high value, especially if we think of the historical mistrust in the public institutions in Italy. Such a result confirms the actual decision of the population to rely on public institutions, in order to face COVID-19.

Again, we stress that the positive outcome of the measures determined by PA is conditional on the acceptance of these restrictions by a significant percentage of the population. Thus, a generic attempt to raise the average trust, besides being very challenging, may not be the best strategy to increase citizens' adoption. As we will see later, it is undoubtedly true that increasing average trust results in a higher acceptance, but different acceptance rates may correspond to the same trust value. Identifying the most critical dimensions and working on them may be a better strategy to optimize the global result, minimizing the effort to increase trust values.

5.2 Analysis of community Trust

In the previous section we provided a picture of what happened in the past months, during the first pandemic wave. This kind of analysis provides useful hints to understand what worked and what were the weaknesses. In this section, instead, we consider some hypothetical scenarios, evaluating what might have happened (or what may happen in future) upon the occurrence of specific conditions. In particular, we verify trust evolution in 81 different simulation scenarios, on the basis of the following experimental setting:

1. number of agents: 4260
2. competence: 1, 1.5, 2, 2.5, 3, 3.5, 4, 4.5, 5;
3. willingness: 1, 1.5, 2, 2.5, 3, 3.5, 4, 4.5, 5;

Figure 1 shows community trust evolution when the components of competence and willingness change. Specifically, it varies from a minimum of 22.78% to a maximum of 82.14%. It is worth noting that the curve has a substantial asymmetry. In the presence of a maximal competence and a minimal willingness, trust is estimated at 56.56%. On the contrary, with a minimal competence and a maximal willingness, we get 48.29%. Such an effect is the consequence of a different influence these two cognitive components have on trust. Indeed, subjects assign a higher weight to competence.

In figure 1, the dark gray area of the curve highlights a situation of substantial average distrust. It describes a critical situation, which PA needs to avoid at any cost. In fact, in such cases the average trust of the population is below 40%. In turn, this may result in a strong negative impact on the citizens' acceptance of the measures. In contrast, the lighter gray area describes a positive situation, in which the average trust is greater than 60%. While such general consideration already provides some useful indications, the observation of the whole community may not be enough. An average value does not describe how trust is actually distributed among the population.

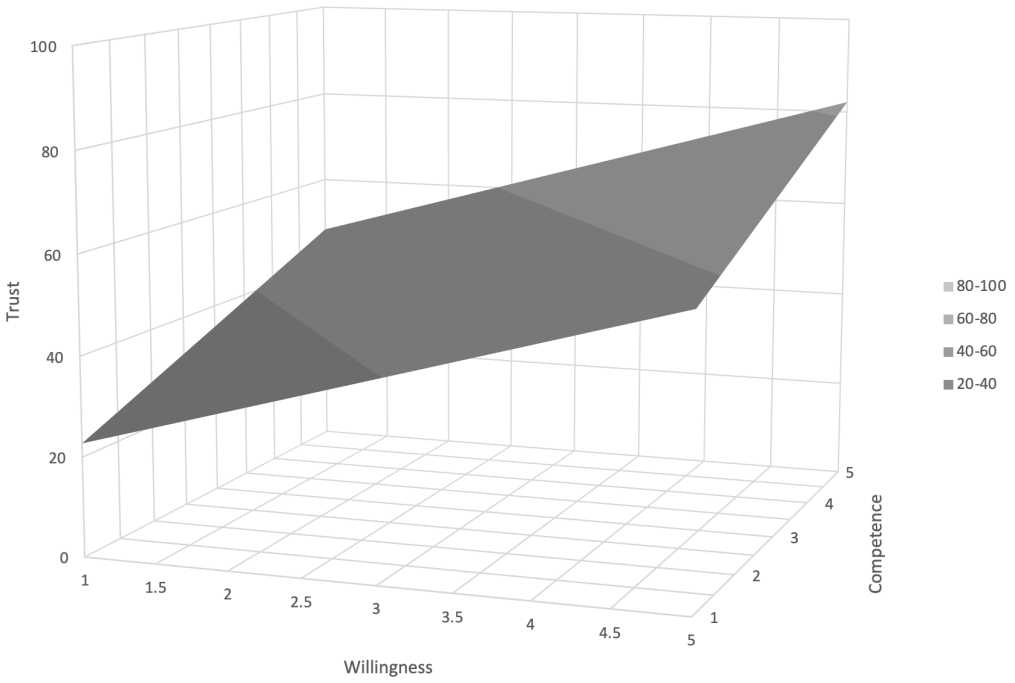

Figure 1: Evolution of community trust, when the evaluation of competence and willingness changes

5.3 Citizens' acceptance

In this section, we investigate the actual acceptance rate for the measures determined by PA. Such dimension is subject to a sufficient trust evaluation, i.e. over a given intrinsic acceptance threshold σ. In order to investigate this topic, we made use of the σ value reported in section 3.3.

Figure 2 shows the results of the simulation. It is easy to notice that this curve has a completely different behaviour, with respect to trust. It is worth underlining that there is a strong difference between trust and acceptance, although these two dimensions are strictly tied. The first represents a means the authority PA has to convince the citizens accepting the restriction, while the second tells us the actual percentage of citizens accepting them. Such a difference explains why we found limit values (maximal and minimal) for trust evaluation, while this does not hold for acceptance rate, which varies from 0% to 99.6%, covering all the available interval.

Concerning the analysis of the corpus, subjects reported 89.72% acceptance rate. This high value actually corresponds to what happened in the early stages of the pandemic.

6 Conclusions

Within this study, we proposed investigated the reaction of the Italian population to the first pandemic episode of COVID-19, taking into consideration the strong link between this dimension and trust in Public Authorities. This kind of analysis is particularly interesting, since it allows us to understand what happened in the past months and, identifying what worked and what did not, it also offers useful suggestions on how to manage further institutional interventions in the future. Most notably, the key conceptual advantage of this research is the attention to the different trust dimensions, since trust is a complex and layered construct, with its own internal dynamics. The core idea, widely highlighted in literature [Sie14, Lew20], is that trust in public authorities is pivotal, both before and during the pandemic, to positively influencing people's willingness to adopt recommended behaviour.

The first result of this study confirms a definitely high trust in public institutions, equal to 72.35%. Considering the historical distrust in the Italian institutions, this is a surprising result. Although such value is very different from what we would have expected, many other studies confirm such strong change. For instance, a survey on a representative sample of Italian citizens (N = 1028, 16–17 March 2020), conducted in the same period

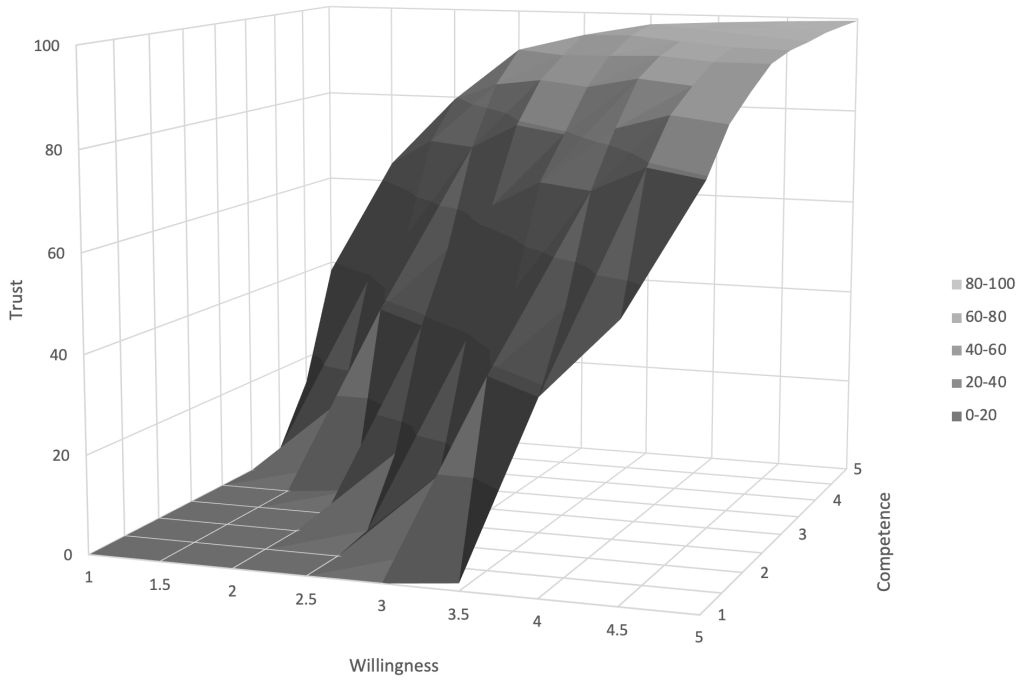

Figure 2: Acceptance rate within the population

by the independent research centre Demos and Pi (`http://www.demos.it/a01705.php`), reported that 71% of citizens trust both the Italian government and the current prime minister, compared to the 44% of the previous month, and 94% approval of the adopted measures. Moreover, according to the analysis conducted by Statista (`https://www.statista.com/statistics/977223/support-for-prime-minister-conte-in-italy/`), trust in the Italian prime minister in February was around 39%. The same report shows a sudden rise of trust up to 54% immediately after the emergency arrival.

Therefore, we considered the evolution of trust when its components change, investigating what could have happened or could still happen in the future in the presence of different conditions. Specifically, we are interested in two specific dimensions: competence and willingness. Results revealed that, depending on these two features, the maximum trust value was 82.14%. Such value represents a useful term of comparison, giving us a clearer view of how high trust was during the pandemic. Conversely, the minimum trust value is estimated at 22.78%.

Remarkably, we identified a considerable difference in weight between competence and willingness. This kind of phenomenon can be traced back to different factors. For example, it is possible that willingness is taken for granted: given the specific role PA holds within the society, citizens may assume that it is certainly willing to protect the population or that it is its responsibility to take care of such problems. Moreover, this lower weight may be the consequence of the assumption that there are no secondary goals behind the actions of the authority (creating alarmism, speculating on the crisis). Once malevolent reasons are excluded, the only possibility is that the authority is actually determined to protect the population. Besides that, concerns about competence need to be verified: weather PA is able to determine adequate and effective rules; if it made appropriate use of experts contribution (e.g. epidemiologists, virologists, paediatricians, psychologists, etc.) and so forth. In such a view, competence automatically plays a critical role.

The most noteworthy outcome of this work concerns the actual acceptance of the restrictive measures by the population. We care to deepen this concept and to provide some clarification, as this is a key point in our framework. It is important to underline that the final goal of the public authority is not just to increase trust, but to increase the acceptance rate of the restrictive measures. For instance, the authority may estimate that an acceptance rate of 80% would ensure an adequate response to the pandemic. From this perspective, a simple

trust assessment would not be a useful index, as different acceptance rates may correspond to a single trust value.

In other words, every trust value enough high to exceed the acceptance threshold σ is fine. Although there is a strong interconnection, trying to maximize trust does not necessarily coincide with maximizing the acceptance rate too. Referring to the first pandemic episode, we estimate an average trust value of 72.35%, while the acceptance rate is definitely higher (89.72%). Although this is an excellent result, we estimate that, under specific circumstances, even a 60% trust value may be enough to convince more than 90% of the population. High trust values may be interpreted as a positive signal, but it is not necessary to reach such values and they would not be a guarantee of success.

In conclusion, this work, besides providing a detailed analysis and interpretation of what happened during the first pandemic episode, could be of great help to understand and explain the decision-making processes of the citizens even in future scenarios, guiding health authorities in providing informed interventions and clear communications.

As future work, we would like to apply the category concept within this framework. We believe that the findings of this work are already interesting. Nevertheless, understanding pandemics requires fine-grained data representing specific local conditions and the social reactions of individuals, in order to cover the full behavioural and social complexity of societies under pandemic crisis. Considering different social categories may provide a more detailed perspective on the problem.

References

[Ana20] Anastassopoulou, C., Russo, L., Tsakris, A., & Siettos, C. (2020). Data-based analysis, modelling and forecasting of the COVID-19 outbreak. PloS one, 15(3), e0230405.

[Bak20] Bakker, M., Berke, A., Groh, M., Pentland, A. S., & Moro, E. (2020). *Effect of social distancing measures in the New York City metropolitan area*. Boston: Massachusetts Institute of Technology.

[Bar20] Barlow, N. S., & Weinstein, S. J. (2020). Accurate closed-form solution of the SIR epidemic model. Physica D: Nonlinear Phenomena, 408, 132540.

[Bon20] Bontempi, E. (2020). *The Europe second wave of COVID-19 infection and the Italy "strange" situation.* Environmental Research, 110476.

[Bun20] Bunker, D. (2020). Who do you trust? The digital destruction of shared situational awareness and the COVID-19 infodemic. International Journal of Information Management, 55, 102201.

[Cai20] Cairney, P., & Wellstead, A. (2020). *COVID-19: effective policymaking depends on trust in experts, politicians, and the public.* Policy Design and Practice, 1-14.

[Cas98] Castelfranchi, C., & Falcone, R. (1998, July). Principles of trust for MAS: Cognitive anatomy, social importance, and quantification. In Proceedings International Conference on Multi Agent Systems (Cat. No. 98EX160) (pp. 72-79). IEEE.

[Cas00] Castelfranchi, C., & Falcone, R. (2000, January). Trust is much more than subjective probability: Mental components and sources of trust. In Proceedings of the 33rd annual Hawaii international conference on system sciences (pp. 10-pp). IEEE.

[Cas10] Castelfranchi, C., & Falcone, R. (2010). Trust theory: A socio-cognitive and computational model (Vol. 18). John Wiley & Sons.

[Cre20] Cremonini, M., & Maghool, S. (2020). The unknown of the pandemic: an agent-based model of final phase risks. Available at SSRN 3584368.

[Cur20] Currie, C. S., Fowler, J. W., Kotiadis, K., Monks, T., Onggo, B. S., Robertson, D. A., & Tako, A. A. (2020). How simulation modelling can help reduce the impact of COVID-19. Journal of Simulation, 14(2), 83-97.

[Fal01] Falcone, R., & Castelfranchi, C. (2001). Social trust: A cognitive approach. In Trust and deception in virtual societies (pp. 55-90). Springer, Dordrecht.

[Fal20a] Falcone, R., Colì, E., Felletti, S., Sapienza, A., Castelfranchi, C., & Paglieri, F. (2020). *All we need is trust: How the COVID-19 outbreak reconfigured trust in Italian public institutions.* Frontiers in Psychology, 11.

[Fal20b] Falcone, R., & Sapienza, A. (2020). How COVID-19 Changed the Information Needs of Italian Citizens. International Journal of Environmental Research and Public Health, 17(19), 6988.

[Fan20] Fanelli, D., & Piazza, F. (2020). Analysis and forecast of COVID-19 spreading in China, Italy and France. Chaos, Solitons & Fractals, 134, 109761.

[Fel20] Fell, L. (2020). Trust and COVID-19: Implications for Interpersonal, Workplace, Institutional, and Information-Based Trust. Digital Government: Research and Practice, 2(1), 1-5.

[Fll20] Felletti, S. (2020). *"Trust me, I'm your neighbour" How to improve epidemic risk containment through community trust.* Mind & Society, 1-4.

[Hun20] Hunter, E., Mac Namee, B., & Kelleher, J. (2020). A Hybrid Agent-Based and Equation Based Model for the Spread of Infectious Diseases. Journal of Artificial Societies and Social Simulation, 23(4), 1-14.

[Jef08] Jefferson, T., Foxlee, R., Del Mar, C., Dooley, L., Ferroni, E., Hewak, B., ... & Rivetti, A. (2008). *Physical interventions to interrupt or reduce the spread of respiratory viruses: systematic review.* Bmj, 336(7635), 77-80.

[Lew20] Lewnard, J. A., & Lo, N. C. (2020). Scientific and ethical basis for social-distancing interventions against COVID-19. The Lancet Infectious Diseases, 20(6), 631-633.

[Li20] Li, Q., Guan, X., Wu, P., Wang, X., Zhou, L., Tong, Y., ... & Feng, Z. (2020). Early transmission dynamics in Wuhan, China, of novel coronavirus–infected pneumonia. New England journal of medicine.

[Man20] Manzo, G., & van de Rijt, A. (2020). Halting SARS-CoV-2 by Targeting High-Contact Individuals. Journal of Artificial Societies and Social Simulation, 23(4), 1-10.

[Pag20] Pagnini, F., Bonanomi, A., Tagliabue, S., Balconi, M., Bertolotti, M., Confalonieri, E., ... & Villani, D. (2020). *Knowledge, concerns, and behaviours of individuals during the first week of the coronavirus disease 2019 pandemic in Italy.* JAMA network open, 3(7), e2015821-e2015821.

[Pre20] Prem, K., Liu, Y., Russell, T. W., Kucharski, A. J., Eggo, R. M., Davies, N., ... & Klepac, P. (2020). *The effect of control strategies to reduce social mixing on outcomes of the COVID-19 epidemic in Wuhan, China: a modelling study.* The Lancet Public Health, 5(5), e261-e270.

[Ram04] Ramchurn, S. D., Huynh, T. D., & Jennings, N. R. (2004). Trust in multiagent systems. The knowledge engineering review, 19(1), 1-25.

[Roc20] Rockett, R. J., Arnott, A., Lam, C., Sadsad, R., Timms, V., Gray, K. A., ... & Sintchenko, V. (2020). Revealing COVID-19 transmission in Australia by SARS-CoV-2 genome sequencing and agent-based modeling. Nature medicine, 26(9), 1398-1404.

[Sie14] Siegrist, M., & Zingg, A. (2014). The role of public trust during pandemics: Implications for crisis communication. European Psychologist, 19(1), 23.

[Sil20] Silva, P. C., Batista, P. V., Lima, H. S., Alves, M. A., Guimarães, F. G., & Silva, R. C. (2020). COVID-ABS: An agent-based model of COVID-19 epidemic to simulate health and economic effects of social distancing interventions. Chaos, Solitons & Fractals, 139, 110088.

[Tui20] Tuite, A. R., Ng, V., Rees, E., & Fisman, D. (2020). *Estimation of COVID-19 outbreak size in Italy.* The Lancet. Infectious Diseases, 20(5), 537.

[Wil99] Wilensky, U. (1999). NetLogo. Evanston, IL: Center for connected learning and computer-based modeling, Northwestern University.

[Wit02] Witten, I. H., & Frank, E. (2002). Data mining: practical machine learning tools and techniques with Java implementations. Acm Sigmod Record, 31(1), 76-77.

[Zng20] Zhang, J., Litvinova, M., Liang, Y., Wang, Y., Wang, W., Zhao, S., ... & Yu, H. (2020). Changes in contact patterns shape the dynamics of the COVID-19 outbreak in China. Science, 368(6498), 1481-1486.