

Master thesis

Trust and the structure of human cooperation

Impact of trust on the clustering of social networks

Roxane Van den Bossche

under the supervision of Jean-Baptiste André and Coralie Chevallier

at *Institut Jean-Nicod*

2023-2024

Abstract. This study examines the impact of social trust on the structure of social networks. We model reputation-based cooperation in social networks, varying the level of trust to assess its effect on both the network's clustering coefficient and size of social groups. We present two models and the results of simulations. Under a wide range of parameter values, we show that low social trust induces more clustered social networks. We then test these findings by examining the correlation between clustering and trust using empirical data from two previous studies conducted in the United States. The observed negative correlation between social trust and clustering supports our theoretical predictions.

Originality statement

The present study is original in that we investigate the consequences of social trust on the shape of social networks. As described in more detail in the introduction, such work has not been done before, especially the precise aspect that we study: social clustering and the size of social groups.

Previous studies have investigated the effect of social clustering on cooperation. They are briefly described in the Introduction of this report. Many studies have looked at social trust and its origins. One empirical article discussed in introduction goes in the same direction as our work, but they used surveys and experiments, while we present a model of the phenomenon.

Declaration of contribution

All the simulations and results presented in this report were formalised and coded by Roxane Van den Bossche. The ideas for the models were found during discussions with Jean-Baptiste André, Coralie Chevalier and Roxane Van den Bossche. The theme of interest was introduced by Jean-Baptiste André, Coralie Chevalier and Nicolas Baumard, in line with the global nature of their research. The empirical data were collected by the studies in the Materials and Methods section, but were normalised to the state level and statistically compared by Roxane Van den Bossche. Juliette Lelièvre also participated in this project as an L2 student in an independent part of the project, which is not presented in this report. This report was written by Roxane Van den Bossche, with the contribution of Coralie Chevalier and Jean-Baptiste André for corrections, in particular to improve the introductory section.

Acknowledgements

I would like to thank Jean-Baptiste André and Coralie Chevalier for supervising this internship. I would like to thank (again) Jean-Baptiste André, Nicolas Baumard and Pat Barclay for their fantastic lectures on evolutionary anthropology, which introduced me to this field of research. I would also like to thank all the IJN staff for the great atmosphere, discussions and lab meetings during this internship. Finally, I would like to thank my family and friends for listening to me when I wanted to talk about my work at all costs, especially Rémy Défossez, who helped me look for a patch for the first model and found the article proving that no patch would be possible.

Contents

1	Introduction	3
2	Material and methods	7
2.1	Reputational Flow	7
2.2	Model one	9
2.3	Model two	11
2.4	Measures on social networks	14
2.5	Empirical data	16
3	Results	18
3.1	Model one	18
3.2	Model two	20
3.3	Empirical data	22
4	Discussion	24
4.1	The big picture	24
4.2	Limits	25
4.3	Further work	26
A	Robustness tests	27
B	Empirical data	33
C	ChatGPT	34

1 Introduction

Human societies across time and space exhibit remarkable diversity in their structure, characterized by variations in both the size of social units and the degree of separation between them. Social units can range from small, intimate groups such as nuclear families or hunting bands, often found in hunter-gatherer communities, to larger, more complex aggregations such as tribes or chiefdoms, to vast social networks that span entire nations and connect individuals who may never meet in person [Turchin et al., 2022]. In addition, the degree to which these units are isolated from one another varies considerably [Glowacki and Lew-Levy, 2022]. In some cases, social networks are highly clustered, with small groups cooperating only internally and remaining largely separate or hostile to one another. In other cases, social networks are continuous, with no clear boundaries between groups, allowing for widespread and cohesive cooperation.

The origin of variability in social structure has been the subject of important work, in particular through the prism of group size and political organization. Three main families of explanations have been proposed to account for the variability in social structure.

Cultural explanations propose that cultural norms, and in particular the cooperative quality of these norms, determine the scale of cooperation, thereby accounting for the influence of culture on social organization [Henrich et al., 2010]. However, these theories do not explain where these norms come from and why some societies develop effective large-scale cooperative norms, while others do not, leaving the exogenous factors underlying variations in social structure unaddressed.

Political explanations, on the other hand, propose that the quality of a society’s political institutions explain the size of its social units [Robinson and Acemoglu, 2012]. These explanations are grounded in the observation that variations in the size of social units consistently correlate with changes in their political organization [Turchin et al., 2022]. However, similar to cultural theories, political theories do not adequately address why some societies develop effective institutions for large-scale cooperation while others do not.

Finally, ecological explanations suggest that a society’s means of production and its ability to extract wealth from the environment shape its structure through an effect on inequality and population density. For example, this family of explanations accounts for how the Neolithic revolution facilitated the emergence of larger-scale political organizations [Powers and Lehmann, 2014]. Its main limitation, however, is the emphasis placed on demography as a key driving force. Empirical evidence suggests that demographic expansion often does not increase cooperation, but rather leads to increased conflict, especially when combined with a decline in per capita wealth [Acemoglu et al., 2020]. This suggests that in addition to demographic factors, individual factors significantly influence the structure of human cooperation.

The purpose of this master’s thesis is to explore an alternative explanation that is grounded in psychology. Our theory posits that the structure of human cooperation is influenced by the level of social *trust* individuals experience in a given environment. Specifically, this thesis suggests that a human preference for interacting in clustered social networks acts as a compensatory response to low levels of social trust. This preference facilitates the formation of small and isolated groups that, in turn, create environments where trust between *specific* individuals can be effectively established and reinforced.

Social trust is defined as the *a priori* belief held by individuals in a given society about the likelihood of cooperation from others whom they do not know personally (it is also sometimes referred to as “generalized trust”, [Nannestad, 2008]; in this report we will sometimes also simply call it “trust”).

Many empirical studies support the fact that social trust varies significantly across places and time [Nannestad, 2008]. [The social capital project, 2017] measured trust across different states in the United States and found notable differences. For example, 28.6% of people in Tennessee reported that “most people can be trusted most or all of the time,” compared to 60.3% in Oregon. [Roth, 2024] analyzed cross-national data from 1980 to 2020 and showed that trust also fluctuates over time within countries.

Well-established theories, supported by numerous empirical observations, suggest that this variability in trust is causally related to variability in individual wealth and access to resources [Mell et al., 2021, Lie-Panis and André, 2022, de Courson et al., 2024, Kesternich et al., 2020, Mell et al., 2022, Reiter et al., 2023]. Beyond

understanding the *origins* of trust variation, the focus of the present study is to examine its *consequences*, in particular how these variations affect the structure of social networks.

At first glance, one might assume that a decline in trust would reduce the level of cooperation among all members of a society, in a uniform way. Our hypothesis challenges this assumption. Instead, we posit that the primary effect of declining trust is that societies become fragmented into smaller, and more clustered, interaction groups. This counterintuitive effect is complex and beyond the scope of mere verbal reasoning. More formal methods are needed to explore this phenomenon. Thus, the goal of this thesis was to develop a simulation model to investigate whether and how variations in social trust influence the structure of social networks. This model is based on two premises.

Premise 1: Mutual connections enhance trust.

People choose the partners they want to cooperate with and those they want to avoid according to several criteria. Our first premise is that at least one of these criteria is the number of mutual connections they have with them, as a greater number of mutual connections often correlates with both better reputational information and stronger reputational incentives for cooperation.

This premise is justified by both theoretical considerations and empirical observations. From a theoretical point of view, evolutionary theories are very consistent in asserting that human cooperation with nonkin is based on reputation [Roberts et al., 2021, Raihani and Barclay, 2016, Barclay and Willer, 2007]. Cooperating with partners who are not genetically related provides no benefit in terms of inclusive fitness, at least not in the general case [West et al., 2011]. The fitness benefit that underlies cooperation with non kin is therefore reputational. This has two important implications. First, in the absence of direct information about a given potential partner, people will prefer to cooperate with that partner if they have indirect information on this partner, through their interactions with mutual friends. Second, other things being equal, people will prefer to cooperate with someone they know will face high reputational costs if they cheat because this would be revealed to mutual friends, leading to severe social consequences. As a result of these two effects, the more mutual friends there are between two individuals, the more likely they are to trust and cooperate with each other.

Empirically, numerous experimental studies show that people choose their cooperation partners based on their reputations (reviewed in [Roberts et al., 2021]), including indirect information about their behavior toward third parties [Barclay and Willer, 2007]. More specifically, a recent survey of first-year students at a Japanese university showed that students do prefer to interact with partners they already have common friends with whenever possible [Igarashi and Hirashima, 2021].

Premise 2 : Mutual connections are more critical when trust is low

The second premise is that choosing partners with whom one has a high degree of mutual connection is a useful compensatory mechanism when trust is low, but it becomes less important when trust is high. As a result, people’s preference for mutual connections is stronger in low trust social settings.

Theoretically, this second premise is a direct consequence of evolutionary theories of reputation-based cooperation. The higher the risk of encountering cheaters, the more important it is to choose partners (i) about whom one has a lot of redundant information through common friends, and (ii) who will not be tempted to cheat because they fear that their dishonesty could lead to widespread disapproval among common acquaintances. On the contrary, if the general level of trust in the population is uniformly high, we can expect other social preferences to become more important, such as the preference for partners with high levels of skill or with interests that match one’s own, etc.

Empirical studies support this idea. [Igarashi and Hirashima, 2021] conducted a survey of first-year students at a Japanese university, measuring both social trust and individual networks. Their results showed that students with higher levels of trust were more likely to delete and form new connections and were less likely to be part of triangular relationships, i.e. less likely to share mutual friends. Similarly, [Horita and Yamazaki, 2022] observed that social trust correlated with a higher proportion of outgoing connections in the social networks of high-trust individuals, indicating a tendency to initiate non-reciprocal relationships. Con-

versely, individuals with low trust had a greater number of incoming connections, representing non-reciprocal relationships initiated by others. While the latter study did not specifically analyze the existence of mutual connections, it nevertheless demonstrates the significant influence of trust on individuals' freedom to form new relationships.

The goal of this Master's thesis was to develop numerical simulations to study the consequences of these two premises on the dynamics of a social network, and to test whether they can indeed explain how low social trust can lead to increased fragmentation and more clustered social networks.

Numerous previous studies examine the relationship between social network structure and cooperation (reviewed in [Takács et al., 2021]).

Several theoretical papers show the consequences of clustering for cooperation in social networks. They show that higher clustering promotes cooperation through gossip and reputation-based cooperation. This was shown by [David-Barrett, 2023], who modelled cooperation in networks with different clusterings and calculated how many iterations cooperation needed to become an optimal strategy. He showed that higher clustering allows cooperation to become optimal in fewer time steps. This is made possible by the gossip system, in which an individual will only cooperate with another if the majority of his friends consider that person to be reliable (i.e. they have not yet been cheated by him). [Coleman, 1988] had already mentioned in 1988 that the closure of social networks is a facilitator of social capital, with closure being an idea very close to clustering. [Kuperman and Risau-Gusman, 2012] worked on larger regular graphs (with thousands of nodes) and showed that higher clustering facilitates the evolutionary invasion of cooperators. They suggest that clusters allow the formation of cooperators that are protected by the external nodes, which in turn are kept alive by the resources of the internal group.

Studies modelling cooperation on graphs are numerous and address different problems. Some study only regular graphs (graphs in which each vertex has as many links as the others) [David-Barrett, 2023, Kuperman and Risau-Gusman, 2012], some use fixed graphs ([David-Barrett, 2023, Kuperman and Risau-Gusman, 2012]), while others use dynamic links ([Pacheco et al., 2006a, Pacheco et al., 2006b, Pacheco et al., 2008]). Most of them rely on the Prisoner's Dilemma from game theory and run a small game on each link of the graph, but some variants exist, such as [Hauert et al., 2002], which uses public good games instead of the Prisoner's Dilemma.

The shape of social networks has many other important socio-economic consequences. Relational mobility has many consequences for psychology and inequality, [Thomson et al., 2018, Yuki and Schug, 2020] and the shape of social networks could have a significant effect on it. This relationship has not been studied much and could be a continuation of this work (see Subsection 2.4.3).

Although many articles study the effects of the structure of social networks, very few study the origins of this structure and there is no theoretical study that proposes a model of this origin. The present work aims to fill this gap.

To do so, we will simulate the dynamics of social networks in which connections are created and removed based on individuals' preferences, in particular their preference for mutual connections.

We will then measure the degree of clustering observed in the networks generated at equilibrium, which will provide an assessment of the extent to which the network is fragmented into internally connected but externally isolated groups. In the time allotted for this master's project, we were unable to obtain a satisfactory measure of group size per se, which is a complex and difficult concept to quantify in graphs. However, we will provide visual representations of the generated networks, allowing at least a visual assessment of group sizes.

Ultimately, we will show that individual preferences and micro-decisions lead to the emergence of macroscopic patterns in social structure through the natural dynamics of social network construction. Specifically, we will show that less trusting individuals who seek mutual connections tend to form highly clustered networks, resulting in small, fragmented groups.

In the final stage of the project, we aim to test our theoretical predictions using empirical data from two

independent sources. We use data on the degree of clustering in the Facebook social network across different US states, as documented by [Chetty et al., 2022], and correlate these with independent data measuring trust levels in the same US states from [The social capital project, 2017]. This analysis shows that US states with lower levels of social trust tend to have more clustered social networks on Facebook, suggesting a potential causal link between social trust and network clustering, as predicted by our model.

2 Material and methods

We model reputation-based cooperation in a networked population of 100 individuals. We introduce the notion of reputation flow to model the impact of reputation on cooperation. Reputational flow quantifies how reputational information travels through the network. Individuals then take reputation into account when deciding whether or not to cooperate.

We hypothesise that when social trust is low, people pay more attention to how reputation circulates than in high trust societies. In particular, when social trust is low, people tend to maximise the flow of reputation in their relationships because they want to know about the behaviour of others (premise 2).

In the remainder of this thesis, we will refer to the trust between a given pair of individuals as the *dyadic* trust between them. This term describes their *specific* propensity to trust each other, by opposition to the notion of social trust, or simply trust, which represents the belief that an average unknown partner will cooperate. Our idea is that the easy circulation of information between two individuals increases their dyadic trust, compensating for the lack of social trust. We present two models with two different approaches.

In the first model, social trust is represented by a trust threshold above which individuals decide to cooperate. In the second model, social trust is represented by the weight that individuals give to trust over natural affinity when calculating the utility of cooperating with someone.

The codes of the two models used for the simulations are available on github at the following adress <https://github.com/RoxaneVDB/Cogmaster-M2-Internship>, together with example videos of the simulations.

Subsection 2.1 defines reputation flow, Subsections 2.2 and 2.3 present the first and second models respectively. Subsection 2.4 presents the measures used to study the network. Finally, Subsection 2.5 presents how the empirical data were collected from previous studies.

2.1 Reputational Flow

We introduce the notion of *reputational flow* as a measure of gossip within a social network. Imagine we have two people in a social network: Ilona and Jade. We want to quantify the information about Ilona that reaches Jade and vice versa. This information is about Ilona's propensity to cooperate or to cheat. Since Ilona will not tell Jade how often she cheats, Jade can only get this information indirectly: through her mutual friends, who may or may not have been cheated by Ilona.

We assume that the information circulates in both directions in the same way (e.g. if Jade knows about Ilona, then Ilona knows about Jade).

Gossip can sometimes go further than just direct mutual friends. For example, Ilona and Jade might not have a friend in common, but have friends who know each other. In this situation, Jade could still hear about Ilona's reputation through her friends, but the probability of the information getting to Jade becomes smaller. We write α as the probability that the information travels via a direct link (we assume α to be a constant). The probability that the information travels via an indirect link with k intermediaries is then α^{k+1} . Figure 1 illustrates indirect links with their lengths as defined below.

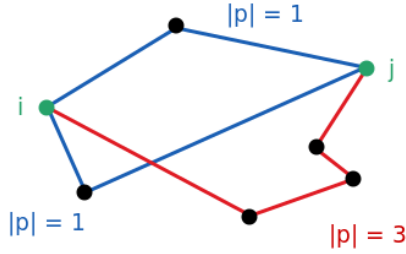


Figure 1: Connexions of length 1 (in blue) and 3 (in red) between i and j

It is also intuitive to consider that if an indirect connection is too long, no information will travel through it. We have chosen a limit of l_{max} for the length of connections we will consider.

Now we can define the *reputational flow* between Ilona and Jade as the sum over all the indirect links between them of the probability of information travelling through the link.

To write it formally: i (Ilona) and j (Jade) are two vertices in a graph representing the social network. A connection between Ilona and Jade with k people in between is a path $p : i \rightarrow \dots \rightarrow j$ in the graph and we define its length $|p|$ as the number of intermediate people (i.e. $|p| = k$). We only consider simple paths, i.e. paths without loops. So we have:

$$RepFlow(i, j) = \sum_{\substack{p : i \rightarrow \dots \rightarrow j \\ |p| < l_{max}}} \alpha^{|p|+1}$$

2.1.1 Algorithmic difficulties

To compute $RepFlow(i, j)$, the main difficulty is to count the number of paths of each length between i and j . Suppose we have a function $NbrPaths_k(i, j)$ that returns the number of paths with k intermediaries from i to j , then we have

$$RepFlow(i, j) = \sum_{k=1}^{l_{max}} NbrPaths_k(i, j) \times \alpha^{k+1}$$

The function $NbrPaths_k$ can be computed by adapting the Floyd-Warshall algorithm, which originally computes the length of the shortest paths between two points. In the Floyd-Warshall algorithm [Floyd, 1962], the shortest paths between all vertices in the graph are computed simultaneously using the following idea: if there is no path of length k or less between i and j , and there is a path of length k or less between i and h and a direct connection between h and j , then the path from i to h with the additional connection $h - j$ is a shorter path between i and j . The algorithm then runs through all values of k , i , j and h to find the length of the shorter paths between each pair of points. We modify this algorithm to store the number of paths for each length.

The strength of the Floyd-Warshall algorithm is that it does not need to remember the paths, only their lengths. This is important because storing all possible paths from one point to another takes up an exponential amount of memory.

However, we mentioned above that we only want to count paths without loops. The reason for this is that, intuitively, information should not go through loops. If Ilona tells Henry that Jade is a cheater, Henry will not repeat it to Jade. Furthermore, Jade is not going to tell anyone else that she is a cheater just because she heard it from Henry. Therefore we only need to count simple paths in the calculation of $RepFlow(i, j)$.

One difficulty is that the Floyd-Warshall algorithm cannot distinguish between non-simple paths and simple paths, precisely because it does not store the paths. Therefore, when adding a new vertex to a path, it is impossible to check that it is not already in the path. The complexity of `NbrPathsk` is therefore much higher if loops are to be deleted. [Roberts and Kroese, 2007] showed that counting simple paths is a #P-complete problem. In practice it is still possible to compute `NbrPathsk` without loops. We have created a greedy algorithm for this, which does not use exponential memory, but exponential time instead. This choice of algorithmic complexity allows the algorithm to run on a computer without overloading the memory. However, in the context of this internship with limited time and resources, it could only be run on very small populations.

The compromise we chose was $l_{max} = 1$. In this particular case, paths with loops can be removed manually, as they all have the form $i \rightarrow h \rightarrow i$. This corresponds to the approximation that information only passes through one intermediary: Jade will only hear about Ilona through her direct common friends. We could explore higher l_{max} in small populations in further work.

2.2 Model one

In this model, the level of dyadic trust between two individuals is entirely determined by the reputational flow between them. At each time step, we calculate the reputational flow between each pair of individuals in the group. We obtain a weighted graph, which we call *reputational flow graph*, represented by a matrix such that the cell in row i and column j contains $RepFlow(i, j)$. This matrix is symmetric and the graph is undirected, since information flows equally in both directions.

This weighted graph represents the *dyadic* trust for each pair of individuals, since dyadic trust is defined as the reputation flow between them.

In this model, social trust is represented by a threshold T . If the dyadic trust between two individuals is higher than this threshold, they cooperate. If their dyadic trust is lower than T , they will not cooperate, even if they used to. In this logic, a low-trust society would correspond to a high trust threshold, because individuals will demand a high level of dyadic trust in order to cooperate. On the other hand, a high-trust society would correspond to a low T . If $T = 0$, this means that the society is so confident that individuals do not care about reputation flow at all.

We call *cooperation graph* the graph of the social network with edges representing cooperative relationships. We assume that information travels through cooperative links.

Finally, we limit the number of connections per person to 5. Limiting the number of friends is realistic because it is impossible to cooperate with an infinite number of people at the same time. We chose 5, even though it is a small number, because the network itself is small and we want to avoid the population of the simulation to form one large unique group of friends.

To summarise, at each time step we calculate the reputation flow depending on the current relationships. Then we update the cooperation by checking for each pair of individuals whether the reputation flow between them is higher than T . We add noise by adding and removing links with a probability defined by the parameter `noise`. Finally, we remove excess friends for individuals who have more than 5 cooperative relationships.

The initial graph is randomly computed with several parameters. Initial links have a maximum length `max_length` to ensure some locality in the initial links and to facilitate the appearance of groups (this parameter is varied in robustness checks). We can vary the probability of making a connection when it respects the maximum length to vary the number of initial connections with the parameter `proba_init_connction`. This parameter does not give the real probability of the presence of a connection, because the number of connections per person must be kept below the maximum number of friends `max_friends`.

Figure 2 and 3 show some steps of example simulations for a very low and a very high trust threshold respectively. Figure 3 shows a case of *quasi-collapse*, a problem that is very common in model 1 and that is discussed in the results.

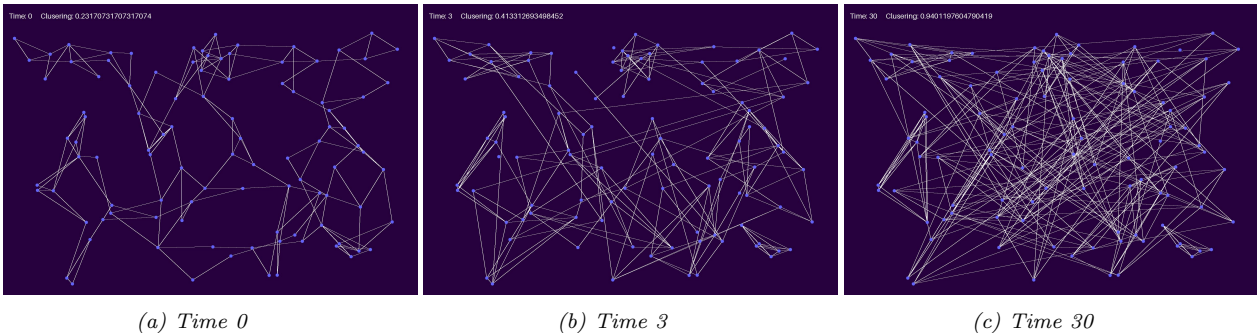


Figure 2: Visualisation of steps 0, 3 and 30 of an example simulation for model 1 with parameters: $T = 0$ and **noise** = 0.001.

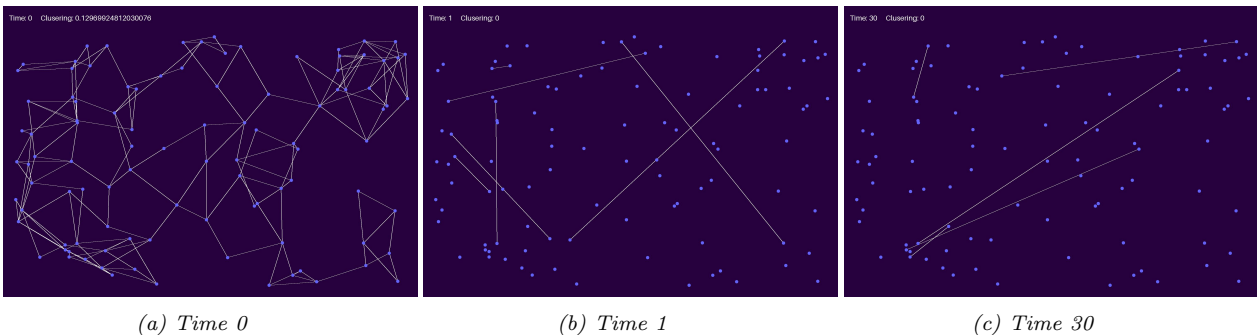


Figure 3: Visualisation of steps 0, 1 and 30 of an example simulation for model 1 with parameters: $T = 0.98$ and **noise** = 0.001.

2.2.1 Why this model is meta

This model simulates meta-cooperation in the sense that we do not model every interaction. We assume that if the reputational incentive is high enough, individuals will trust each other and be reliable because they risk too much by cheating others. On the other hand, we assume that if the reputational incentive is too low, individuals will stop cooperating, either because they no longer trust each other or because one of them cheated in the absence of an incentive to be reliable.

This model can therefore be interpreted in two ways: either we consider that individuals have a concrete preference for high reputational flow and choose their cooperators with regard to the reputational flow between them. Or we can think that individuals have no real preference for high reputational flow, but it just happens that if the incentive to behave nicely is too low, relationships end up breaking down. Reality is probably a mixture of these two views.

2.2.2 Systematic test of parameters

We vary T to observe the consequences of different social trust levels on the shape of social networks. We test different levels of noise (**noise** takes the values: 0, 0.001, 0.002, 0.003, 0.004, 0.005, 0.01, 0.1, 0.2). The maximum reputation flow for $l_{max} = 1$ is $\frac{5}{4} = 1.25$. We test values for T between 0 and 0.8. There is no

need to test higher T because cooperation breaks down at 0.8 (or earlier).

For model 1, no robustness check was performed for two reasons: a first version of this model was implemented with many robustness checks, but it happened to contain an algorithmic bug, and after correcting this bug, there was not much time left for many simulations. More importantly, the results of model 1 were not satisfactory and we preferred to work on model 2, which is more realistic and promises better results. The limiting problems of model 1 are explained in Subsection 3.1.1 of the Results Section.

2.2.3 How to limit the number of friends

A natural idea when choosing which friends to remove if they are too numerous is to remove the less trusted ones. Although this idea is intuitive, it has the artefactual effect of increasing clustering even if the trust threshold T is equal to 0.

An alternative is to randomly remove friends if they are too numerous. We tried this solution, but perfectly randomising the order of friends would have taken a lot of time, and we preferred to work on model 2 instead, which was more promising. The reason why randomising the order of limiting friends is complex is that it is not enough to take individuals in a random order and limit their number of friends. In fact, individuals who are first to choose which friends to keep will end up with more connections, and unlucky individuals who are chosen last will often lose all their connections. This effect makes the model very unbalanced, especially in the random case. We therefore decided to keep the idea of deleting connections with lower reputation flow when limiting friends. This was a more intuitive solution, with fewer arbitrary artefacts. More importantly, we wanted to move quickly to model 2, so we did not try to improve this model too much.

2.3 Model two

This second model aims to improve on the main limitations of model 1. In this model, individuals calculate the utility of each potential relationship and try to form a new relationship with the person who has a higher utility for them. There is no longer a threshold, and each individual tries to create one new link per time step. The number of friends is still limited to 5.

The utility function takes into account two factors: the natural affinity between individuals and their level of dyadic trust. Natural affinity represents all the reasons why two people might want to cooperate: they like each other, they like similar activities, they have a common goal that they could achieve together, etc. Trust is made up of two elements: the flow of reputation and the duration of existing cooperative relationships.

In this model, social trust is represented by the weight given to trust in the utility calculus. We call this weight κ . The weight given to affinity is then $1 - \kappa$. When $\kappa = 1$, social trust is the lowest and individuals only care about dyadic trust when choosing cooperators. When $\kappa = 0$, social trust is at its highest and individuals only care about affinity when choosing cooperators.

2.3.1 Utility

The utility function is defined as:

$$U(i, j) = \kappa \times Trust(i, j) + (1 - \kappa) \times Affinity(i, j)$$

where trust is defined as:

$$Trust(i, j) = (1 - \beta) \times RepFlow^*(i, j) + \beta \times Duration^*(i, j)$$

and $affinity(i, j)$ is a random value between 0 and 1 that is set independently for each pair (i, j) at the start of the simulation.

The weight given to duration over reputation flow β is an arbitrary parameter. We want to use $\beta = 0.5$ (same weight for both). Different values of beta are tested to ensure that its value does not change the results.

$RepFlow^*(i, j)$ and $Duration^*(i, j)$ are normalised versions of $RepFlow(i, j)$ and $Duration(i, j)$. Normalisation was necessary because $Affinity$, $RepFlow$ and $Duration$ take values on different scales. We also considered that reputation flow and friendship duration are affected by an effect of diminishing returns. Therefore, we use the normalisation function $f(x) = \frac{x}{k+x}$ where $k = 5$ for duration and $k = 0.5$ for reputation flow. This normalisation function has a slope of approximately $1/k$ when x is small and tends towards one when x is large. A kind of saturation is reached at $x = 2k$, which we used to determine the k s we used: we considered that cooperating for 10 time steps is a cooperation long enough to inspire maximum trust in terms of duration. For the reputation flow, we calculated the maximum reputation flow and fixed the saturation of the reputation flow near this maximum value. The maximum possible reputation flow between two individuals is 1.25 (5 connections with an intermediate friend). The flow between two friends in a clic of size 6 is equal to 1 and corresponds to a case of maximum clustering. We used 1 as the saturation value because it corresponds to a special case of very intense reputation flow and is close to 1.25.

2.3.2 Algorithmic steps

To summarize the algorithmic steps of the simulation: The social network is initialized randomly at the beginning of the simulation and a random affinity matrix is drawn. At each time step, all individuals calculate the utility function of a potential connection with each other individual depending on reputational flow, duration of a cooperation and natural affinity.

Individuals that have less than the maximal number of friends try to befriend the person with maximal utility between the individuals that are not yet their friend. Individuals that have the maximal amount of friends find the person with maximal utility outside of their friends. If this person has a higher utility than their friend with lower affinity, they befriend this new person and let go their friend with lowest affinity.

Individuals will try to befriend people that do not have the maximal amount of friend already. They will only try to befriend “busy” people with a probability 0.1.

At this moment of the simulation, individuals might have more friends than allowed because others tried to befriend them. Similarly as in the first model, we then limit the number of friends. Individuals are enumerated in a random order and they all remove their friends with lowest utility until they have the allowed number of friends or less. This time there is no unwanted effect because they use the utility function to decide which connection to remove, and the utility already includes the variations in social trust.

Finally, reputational flow and duration are actualised for the next time step.

Several steps of the simulation that must be performed on each individual are done in a random order to avoid favouring the first individuals by always letting them make their choices first.

Noise is added to the simulation by adding an error in the utility calculation that is different for each pair (i, j) and at each time step. The parameter `noise` allows you to adjust the amount of noise in the simulation. The error in each $Utility(i, j)$ is a random value between $-\text{noise}$ and noise .

The initialisation of the network is similar to model 1.

Figure 4, 5 and 6 show some steps of example simulations for a very low, a middle and a very value of κ respectively.

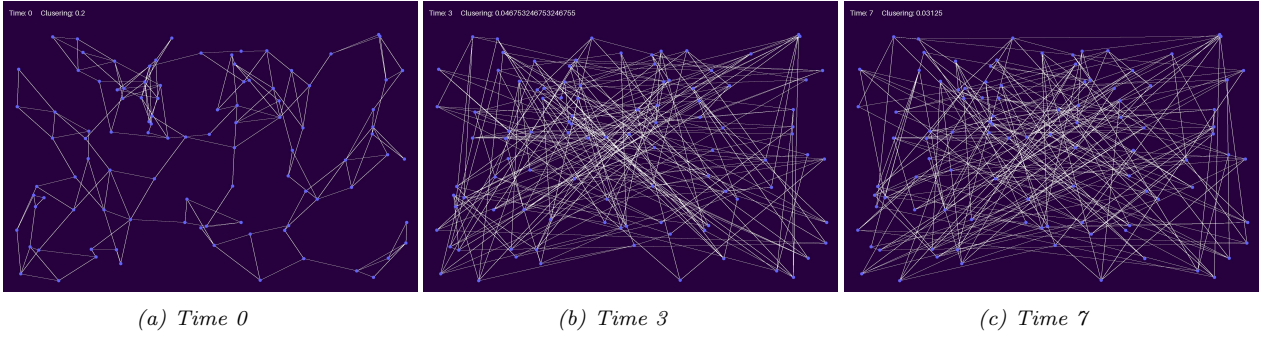


Figure 4: Visualisation of steps 0, 3 and 7 of an example simulation for model 2 with parameters: $\beta = 0.5$, **noise** = 0.1 and $\kappa = 0$.

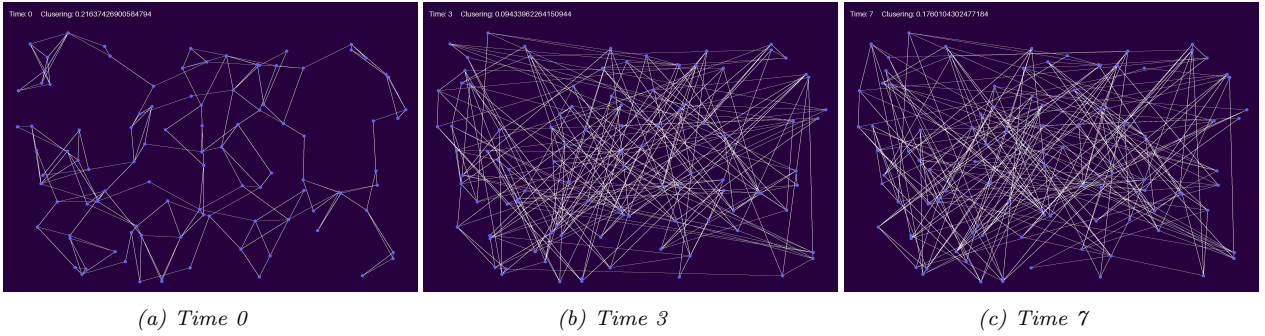


Figure 5: Visualisation of steps 0, 3 and 7 of an example simulation for model 2 with parameters: $\beta = 0.5$, **noise** = 0.1 and $\kappa = 0.5$.

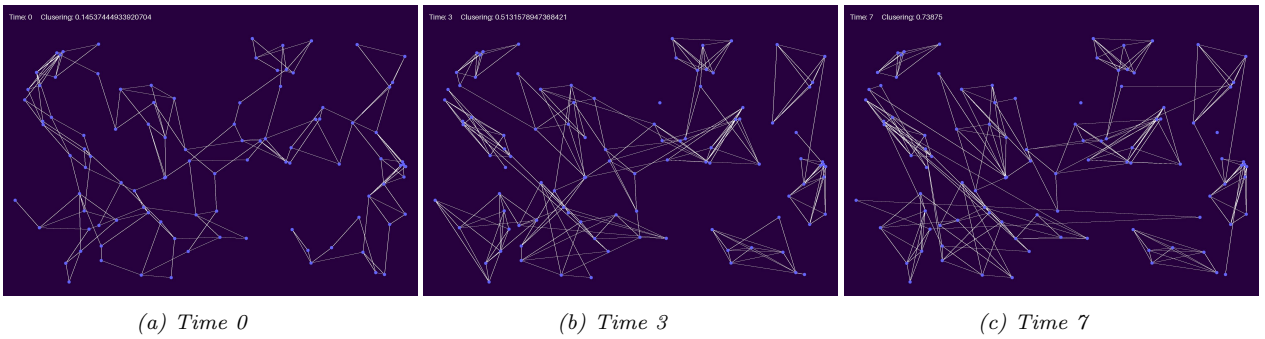


Figure 6: Visualisation of steps 0, 3 and 7 of an example simulation for model 2 with parameters: $\beta = 0.5$, **noise** = 0.1 and $\kappa = 1$.

2.3.3 Systematic test of parameters

The main parameter we want to vary is κ , because it represents social trust. κ can take any value between 0 and 1, and we will systematically sample at least 11 points (from 0 to 1 with an interval of 0.1). For each point, we run 30 simulations to average the measures and stabilise the results. Previous tests showed that 30 simulations is the minimum number of simulations to obtain a stable average (with 20 simulations the results are still unstable, and more than 30 is useless because 30 is already satisfactory).

For each set of parameters, we test the values $\beta = 0.3$, $\beta = 0.5$ and $\beta = 0.7$, coupled with the values $noise = 0$, $noise = 0.1$, $noise = 0.2$ and $noise = 0.5$. Testing three values for β allows us to see how the result depends on the weight given to the reputation flow. Testing different values of $noise$ allows to make sure that the effect is neither a noise effect nor an effect of too much determinism. $noise = 0$ allows you to see the mechanism when the noise has no effect at all, while $noise = 0.5$ is by definition a very high noise, which should be very disruptive to the simulation and affect the results. $noise = 0.1$ and $noise = 0.2$ are more realistic situations where people make some mistakes in their utility calculations and sometimes change their mind about someone, but their calculations are still a good approximation of reality.

Several initialisations are tested. In particular, we want to test the effect of the initial clustering on the simulation, so we run initialisations with different initial clusterings. We vary parameters such as the maximum number of friends and the initialisation to check the robustness of our results. We also test $\beta = 0$ and $\beta = 1$ to check consistency with the other results.

2.4 Measures on social networks

2.4.1 Clustering coefficient

To study the shape of the network, we use a standard mathematical measure on graphs: the clustering coefficient. This measure allows us to quantify how individuals form closely related clusters. In practice, it counts the proportion of triangles in pairs of friends of the same individual. Figure 7 shows the two types of triplets (closed and open triangles). The clustering coefficient counts the proportion of closed triangles in all triplets.

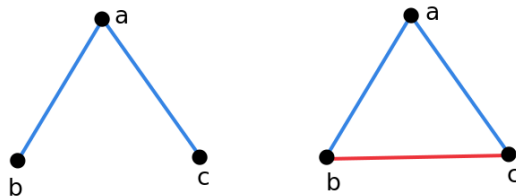


Figure 7: Different triplets: open triangle (left) and closed triangle (right)

We use the standard notion of *global clustering coefficient* rather than the local version because we want to measure clustering on the whole network. The formal definition for a graph g is:

$$ClustCoef(g) = \frac{3 \times \text{number of closed triangles}}{\text{number of all triplets}}$$

where closed triangles are counted as three different triplets: a with two common friends, b with two common friends and c with two common friends. This is why the number of closed triangles in the fraction is multiplied by 3.

The clustering coefficient varies between 0 and 1. If $ClustCoef(g) = 1$, this means that all triplets are closed triangles. Intuitively, this means that for each individual, all their friends are also friends between them. If the clustering coefficient is 0, this means that there are no triangles. Either individuals cooperate without forming clusters, or they simply form too few boundaries. Figure 8 shows examples of both extreme cases.

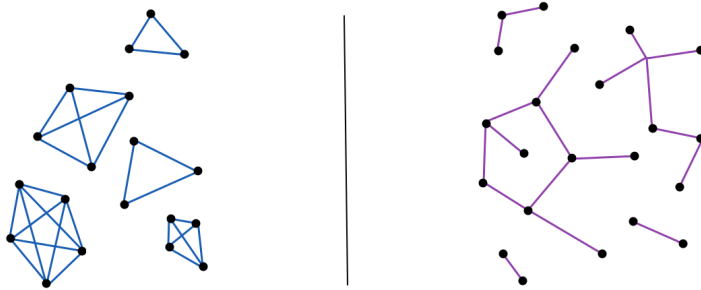


Figure 8: Networks with clustering coefficient equal to 1 (left) and 0 (right)

We have taken the convention that if there is no triplet in the network, then $ClustCoeF(g) = 0$ to avoid division by 0.

We measure the clustering coefficient at the end of the simulation, when its variation has stabilised. We also make some examples where the clustering is measured throughout the simulation, in particular to determine the stabilisation time.

The idea of the algorithm that calculates the clustering coefficient is to enumerate all individuals i . For each i we enumerate all its friends j and for each j we enumerate all its friends h . For each such triplet $i - j - h$ we count one triplet in the graph. If i and h are connected, we also count a closed triangle. Finally, we divide the number of closed triangles counted by the number of triplets counted. In this way of finding closed triangles, each closed triangle is already counted three times (as $i - j - h$, $j - h - i$ and $h - i - j$), so there is no need to multiply the result by 3.

2.4.2 Size of groups

We want to study the variation in the size of the social circles and measure the number of individuals in groups in our simulations.

To do this, we first need to define what we call a group. Many definitions of groups in graphs have been described in the literature, along with algorithms for finding them in a network. This question gave birth to the field of *community detection*, which was nicely summarised in [Fortunato and Hric, 2016] and [Fortunato, 2010]. It is possible to define a group by specifying the minimum number of connections between members of the group, or the maximum distance between members of the group, or the number of groups we want to find, and so on. We can assume that groups cannot intersect, or we can accept that they might intersect. There are many plausible definitions, and none of them seems more suited to our subject than others.

One intuitive notion that one might want to use as a definition for a group is the mathematical notion of *clique*. A *clique* in a graph is a set of vertices of the graph where all the links between each pair of vertices in that set exist in the graph. In other words, a *clique* is a complete subgraph. More precisely, the natural idea would be to consider *maximum cliques*, i.e. cliques that are not themselves included in any other clique and are maximal with respect to inclusion. However, there are many reasons to criticise this definition of a group. The main criticism is that all members of a clique have a very symmetrical role. There is no central individual, etc. More generally, the definition of cliques imposes too strong a constraint on the definition of social groups.

The definitions proposed in community detection can be divided into two categories: definitions that are relaxed versions of the definition of cliques, and definitions that try to maximise the number of internal links and minimise the number of external links.

In the first categories we find the *n-clubs* and the *n-cores*, which may be interesting definitions. A *n-club* is a subgraph such that any pair of vertices in this subgraph can be connected by a path of length at most n

(*n-clique*), with the additional conditions that the path must be inside the subgraph (*n-clan*) and that it is maximal with respect to inclusion (*n-club*). The *k-kore* is a subgraph such that each vertex of the subgraph has a degree of at least k within the subgraph, which means that it has at least k connections within the subgraph. The notion *k-plex* is very close: every vertex must be connected to every other vertex of the subgraph, except at most k of them.

In the second categories, the definitions of *strong communities* and *weak communities* are based on the so-called internal and external degree:

The *internal* and *external* degrees of a vertex v in a subgraph \mathcal{C} are respectively k_v^{int} the number of connections of v with vertices in \mathcal{C} and k_v^{ext} the number of connections of v with vertices outside \mathcal{C} . The internal and external degree of \mathcal{C} are defined respectively as:

$$k_{\mathcal{C}}^{int} = \sum_{v \in \mathcal{C}} k_v^{int} \quad \text{and} \quad k_{\mathcal{C}}^{ext} = \sum_{v \in \mathcal{C}} k_v^{ext}$$

Strong communities are defined as subgraphs such that every vertex in that subgraph has a higher internal degree than its external degree. Formally, \mathcal{C} is a strong community if and only if:

$$\forall v \in \mathcal{C}, k_v^{int} > k_v^{ext}$$

Weak communities are defined as the subgraphs whose internal degree is higher than their external degree, which is a weaker constraint. Formally, \mathcal{C} is a weak community if and only if:

$$k_{\mathcal{C}}^{int} > k_{\mathcal{C}}^{ext}$$

Many other definitions exist, some are based on similar ideas, others are defined along with the corresponding algorithm and have no formal definition.

All of these definitions seem interesting and deserve to be studied. However, because we had limited time to work with, we decided to use the definition from graph partitioning, which is an additional possible definition, because it was intuitive and because some Python libraries that compute graph partitions already existed (we used PyMetis).

The groups are computed in graph partitioning as follows: the number of groups to be found is fixed as a parameter k . We assume that groups cannot overlap and that each vertex must be in a group (i.e. we create a partition). The computer will return the groups such that the number of connections between two groups is minimal, using the idea of the min-cut in graph theory. Our first idea was to test different numbers of groups k and keep the partition with the smallest cut, or which maximises the internal degrees of the groups and minimises the external degrees of the groups. We found that idea cannot work because larger and less numerous groups will always have a smaller cut and better internal and external degrees. So this method cannot be used to measure the evolution of the size of groups. Another definition from the ones described above could be used instead. This will be investigated in further work.

2.4.3 Cost of social mobility

We plan to investigate the costs of social mobility in further work. We want to study the impact of social trust and clustering on relational mobility. To do this, we plan to simulate the arrival of an individual in a mature (i.e. stabilised) social network and measure how this individual manages to integrate this network and form new connections. This could easily be done by reusing the code from model 2.

2.5 Empirical data

We wanted to test the predictions of our models with empirical data. To do this, we collected data from two existing studies: *The Social Capital Project* [The social capital project, 2017] and *Social Capital I and II* by Chetty et al. [Chetty et al., 2022]. Both had open access data sources on the United States. The first had measures of social trust in each US state, and the second had measures of clustering in each county. This

allowed us to calculate the correlation between social trust and clustering in each US state and compare it with the prediction of our model. Our hypothesis is that social trust is negatively correlated with the level of clustering.

2.5.1 Measure of social trust and trust in neighbors

[The social capital project, 2017] includes two variables of interest: social trust in each US state and people’s trust in their neighbours. Trust in one’s neighbours is a useful variable because the data on social trust were only partial, whereas this variable was more complete. We believed that social trust and trust in neighbours would be highly correlated, which would allow us to use trust in neighbours as a proxy for social trust if needed.

Social trust was first measured by the *American National Election Study* (ANES), which asks the standard method used to measure generalized trust: participants are asked to give a rating on a Likert scale from “Most people can be trusted” to “One can’t be too careful”.

Trust in neighbours was first collected by the *Volunteer Supplement to the November 2013 Current Population Survey*. They measured the “share of adults reporting that they can trust all or most of their neighbours”.

These two datasets were presented in [The social capital project, 2017], from which we extracted them.

2.5.2 Measure of clustering

[Chetty et al., 2022] collected Facebook data from American users to measure clustering in their individual social networks. In each US county, the authors measured what proportion of people’s friends were also friends with each other, which is exactly what the clustering coefficient represents. They included friends inside and outside the county. The final variable is the average across all individuals in the county. The database also provides the population size in each county. Since social trust was measured at the state level and not at the county level, we calculated clustering at the state level by calculating the average clustering across all counties within a state, using data on county population size to properly weight the average.

2.5.3 Statistical analysis

We test the correlations between the three variables, including the correlation between social trust and trust in neighbours. We compute the linear regression for the three parameter pairs and report each corresponding r^2 (see Subsection 3.3 for the results).

3 Results

This Section presents the results of the two models (Subsection 3.1 and 3.2). Subsection 3.3 presents the statistical analysis of the empirical data.

3.1 Model one

Figure 9 shows the evolution of the average clustering coefficient as the social trust increases (*i.e.* the trust threshold T decreases) for model 1. The following parameters have been used:

```

Parameters

max_time = 30 # Duration of the simulation
number_of_people = 100 # Population size
proba_init_connexion = 1 # We initialise with as many connections as possible
alpha = 0.5 # Exponential decrement for indirect connexions in reputational flow
max_length = 1 # Longest indirect connexions taken into account in reputational flow
max_friends = 5 # Maximal number of connections per individual
max_dist = width/7 # Maximal length of initial connections
    
```

We see a thresholding effect: clustering is higher when social trust is lower (and T is higher), but it takes the form of a staircase function. Noise reduces the effect of social trust on clustering, as expected. The orange and blue curves represent the proportion of collapse and quasi-collapse, respectively.

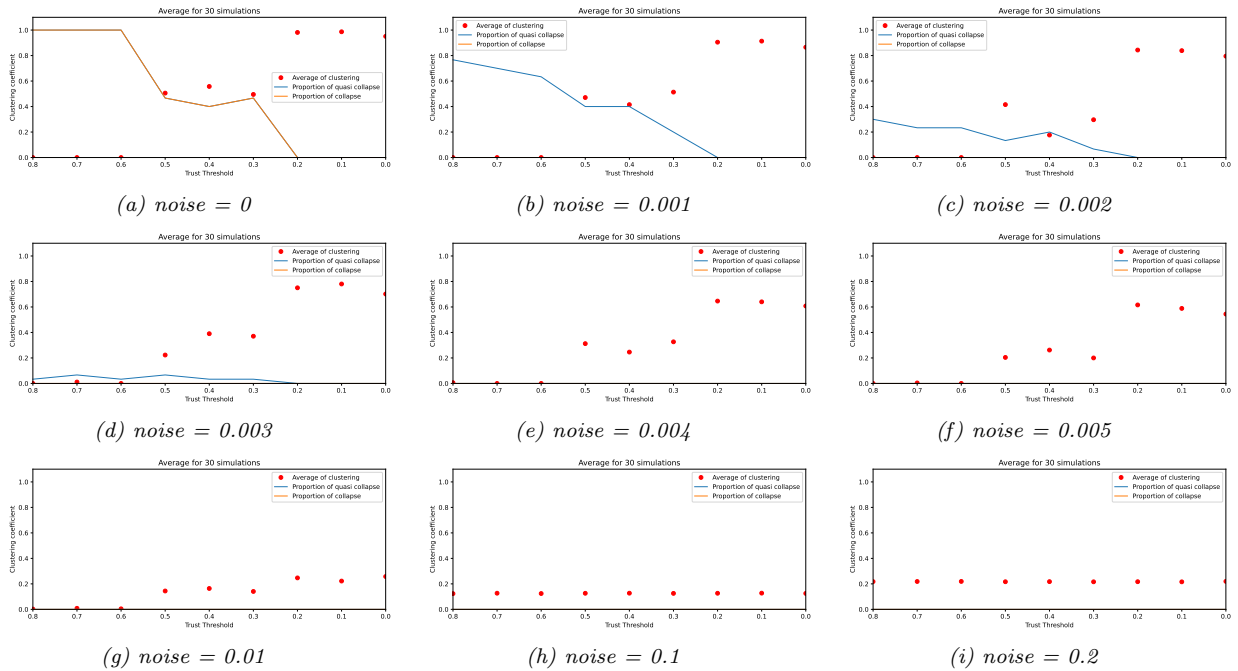


Figure 9: Model 1. Average of the clustering coefficient (red points) depending on social trust. Each point is an average on 30 simulations. On each subfigure, T is decreasing from left to right, corresponding to an increase in social trust. From left to right and top to bottom, subfigures correspond to increasing noise.

The proportion of collapse is the fraction of simulations where cooperation stopped completely (formally,

there is no connection), and the proportion of quasi-collapse is the fraction of simulations where cooperation almost stopped and each individual has at most one relationship (formally, there is no triplet).

We see that when trust is too low and the threshold is too high, cooperation simply stops. In the intermediate phase, between $T = 0.5$ and $T = 0.3$, half of the simulations collapse and the other half reach a high clustering coefficient. The low clustering coefficient in the low trust simulations corresponds to a decrease in cooperation rather than a change in the style of cooperation, with cooperation stopping completely in some cases. We have generated more detailed figures to confirm this impression of a staircase function (Figure 10).

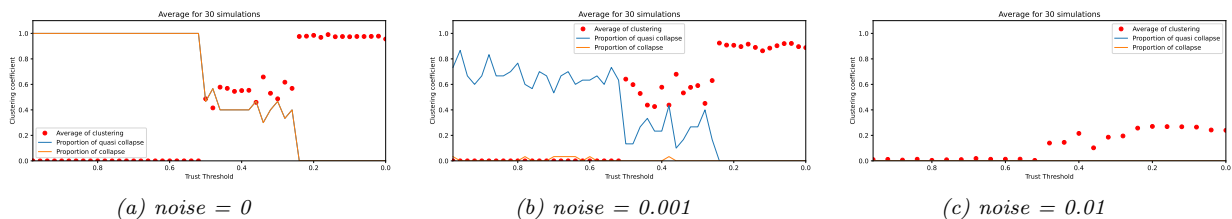


Figure 10: Model 1, with more points on the curve. Average of the clustering coefficient (red points) depending on social trust. Each point is an average on 30 simulations. On each subfigure, T is decreasing from left to right, corresponding to an increase in social trust. From left to right, subfigures correspond to increasing noise.

We see on figure 10 that any increase or decrease in clustering is accompanied by a symmetrical shape in the orange or blue curve for noise 0 and 0.001. This means that the measured variation in clustering is only induced by simulations with clustering equal to 0, which significantly reduces the average of clustering on the 30 simulations. It should be noted that the clustering in this model does not vary as a function of social trust: only the collapse fraction is affected. In the case of noise 0.01, the noise allows cooperation to continue by randomly adding relationships, but it has no realistic meaning. Figure 11 shows exactly this phenomenon by showing the evolution of clustering in time for each of the 30 simulations, for each tested trust threshold.

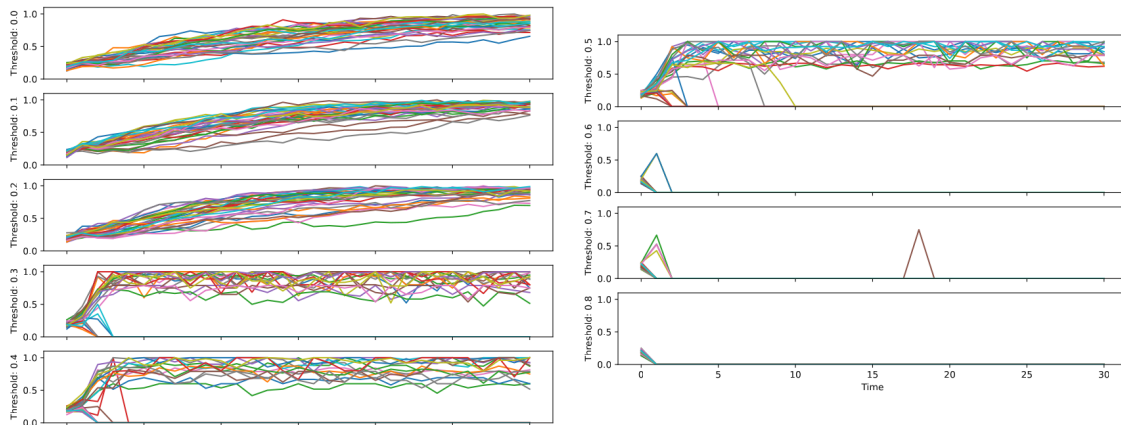


Figure 11: Evolution of the clustering coefficient across time. Each line represent a simulation, each subfigure corresponds to a different trust threshold from 0 (top right) to 0.8 (bottom left), noise = 0.001

Overall, we find that clustering increases regardless of the level of social trust. This is an artefact of the function used to limit the number of friends, which keeps the friends with the higher reputation flow, thus increasing clustering. The only effect this model shows is that too little trust leads to no cooperation at all, which is very questionable. This model has other limitations, which are discussed in the following Subsection.

3.1.1 Limits of model 1

This model has many limitations. The main limitation is that there is nothing to ensure that all individuals will have some connections, and very often a significant proportion of the population has no connection at all, with no chance of forming a new link. This situation, where sometimes more than half of the population stops cooperating, is not realistic.

Another limitation of this model is the threshold effect, which is hardly realistic. The fact that T is a constant, common to all individuals regardless of their situation, is also problematic.

Finally, the fact that interpersonal trust is defined by reputation flow means that many important factors are neglected, such as the duration of friendships. It is also important to note that, in reality, people do not only cooperate with those they trust the most, but also choose others according to their own interests (similar tastes, goals, etc.).

Overall, the effects shown by this model are mainly artefactual and not significant. Model 2 has been designed to address all these problems. It is more realistic and seems more promising in terms of providing representative results.

3.2 Model two

Figure 12 shows the evolution of the clustering with respect to the social trust. The following parameters have been used:

Parameters

```
max_time = 7 # Simulation duration
number_of_people = 100 # Population size
proba_init_connexion = 1 # We initialise with as many connections as possible
alpha = 0.5 # Exponential decrement for indirect connexions in reputational flow
max_length = 1 # Longest indirect connexions taken into account in reputational flow
max_friends = 5 # Maximal number of connections per individual
top_seniority = 10 # Saturating value for normalising friendship duration in trust
max_dist = width/7 # Maximal length of initial connections
```

The number of time steps is very small. Previous tests have shown that simulations stabilise around 5 time steps and are fully stabilised at 7 time steps, allowing many simulations to be run and more data to be collected.

In Figure 12 we see a clear decrease in clustering as κ , the weight given to trust in the utility, decreases. We have varied β and we see that when β is smaller, i.e. the weight given to reputation flow over the duration of the relationship decreases, the effect of social trust decreases and clustering is lower overall. We also see that noise reduces the effect of social trust, leading to very low clustering coefficients when noise is high. The proportion of collapse and quasi-collapse is now constantly equal to zero, proving that the problem of collapse in cooperation in Model 1 has been successfully solved.

This model is conceptually more realistic and gives the expected results: clustering increases when social trust is lower. Model 2 allows us to model the subtlety that this clustering effect is related to how much individuals consider reputational flow when estimating trust. If they give less weight to reputational flow, then the effect of low social trust on clustering is smaller.

To ensure that these results are not artefactual and do not depend on arbitrary parameters, we performed several robustness checks. The corresponding figures can be found in Appendix A. We ran the same simula-

tions with a different maximum number of friends: 10 instead of 5. The results were very similar to Figure 12 (see Figure 16 in the appendix). We also changed the initialisation parameters to see whether the results depended on the fact that there was some initial clustering in the graph, or some initial types of groups that allowed clustering to appear easily. We set `max_dist` to $2 \times \text{width}$, where `width` is the width of the rectangle in which individuals are assigned random positions. This resulted in very low initial clustering, as anyone could form initial links with anyone. Figure 17 in the Appendix shows that the results remain exactly the same. A more extreme test of the effect of initialisation was to initialise with no connection at all, and still the results remained similar (see figure 22 in the Appendix for the results on clustering and Figures 19, 20 and 21 for the visualisation of corresponding example simulations). The curve takes a slightly different shape, but remains decreasing and around the same values.

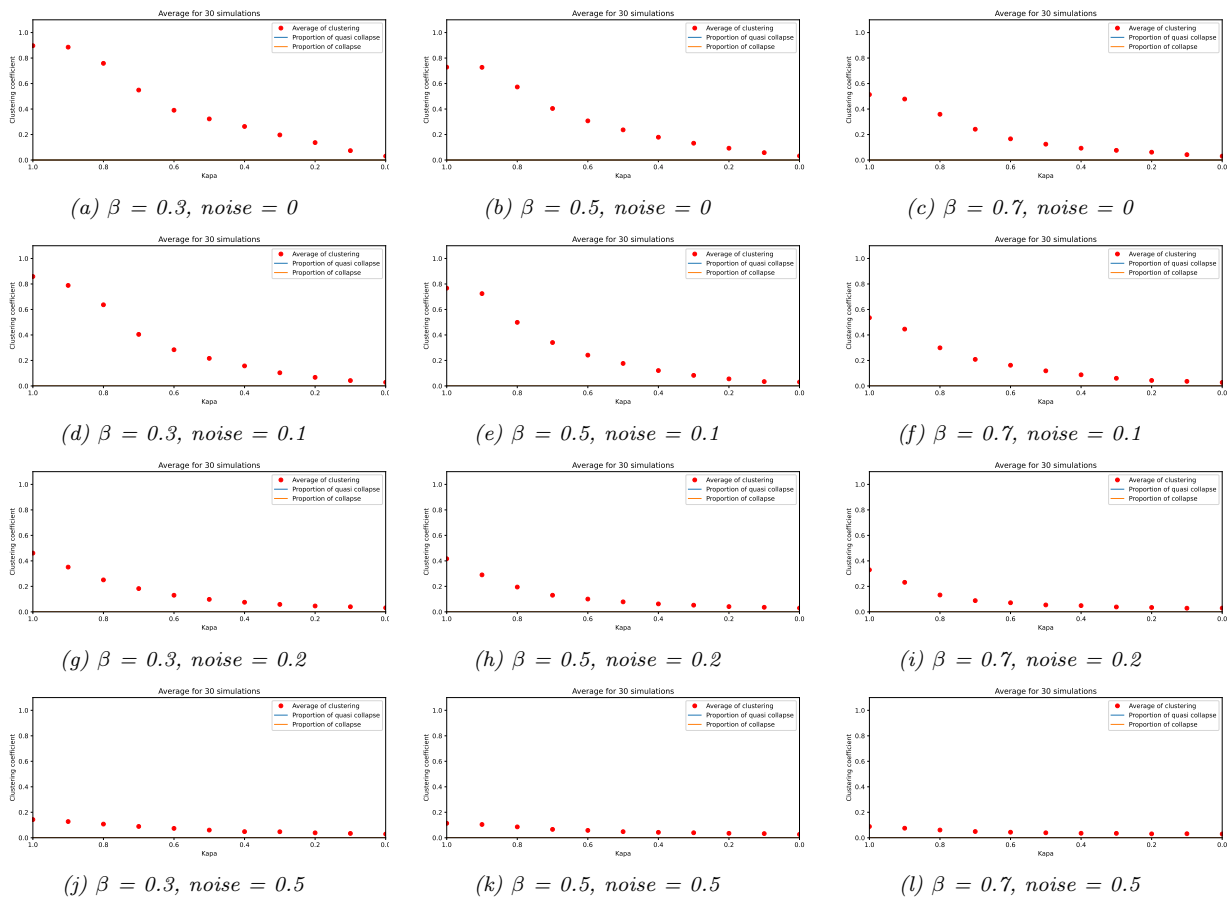


Figure 12: Average of the clustering coefficient (red points) depending on social trust. On each subfigure, the weight of trust in utility decreases in the left direction. Each point is an average on 30 simulations. From left to right, subfigures correspond to increasing β , ie. increasing weight given to duration of cooperation over reputational flow in the calculus of trust. From top to bottom, subfigures correspond to increasing noise.

We also tested whether the effect would be very strong for $\beta = 0$ (when trust is only computed out of reputation flow) and absent for $\beta = 1$ (when trust is only computed out of relationship duration). The Figures 18 in the Appendix confirm this almost perfectly: there is only a tiny effect of decreasing clustering as social trust increases for $\beta = 1$ and $\text{noise} = 0$ (figure 13a). This effect was caused by the initial clustering being preserved in low trust simulations because individuals cared more about the duration of relationships, leading to more conservatism. This intuition was confirmed by the last robustness test on figure 23 in the appendix, where we tested $\beta = 0$ and 1 with an initialisation with no connection (empty initialisation). The

tiny effect disappears as expected when there is no initial clustering, as Figure 13b shows.

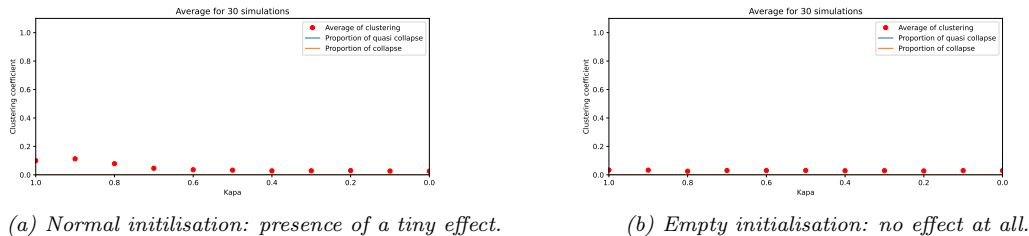


Figure 13: Tiny effect of conservatism on clustering coefficient ($\beta = 1$, noise = 0)

Overall, this model shows a clear causal relationship between social trust and the clustering coefficient, assuming that individuals include reputation flow when choosing their collaborators. This result is robust to changes in several parameters and to changes in initialisation. It is also shown that without the assumption that individuals consider reputational flow when choosing their cooperators, low social trust can still induce higher clustering coefficients when initial clustering is high, simply because defiance induces more conservatism in relationships. The limits of this model are discussed in Section 4.

3.3 Empirical data

The table with the extracted empirical data is to be found in Appendix B. We obtain the following correlation coefficients:

Statistics		
Trust in neighbors / Clustering:	Correlation: 0.359	r^2 : 0.129
Generalized trust / Clustering:	Correlation: -0.371	r^2 : 0.138
Trust in neighbors / Generalized trust:	Correlation: 0.327	r^2 : 0.107

The r^2 correspond to the linear regression on Figure 14.

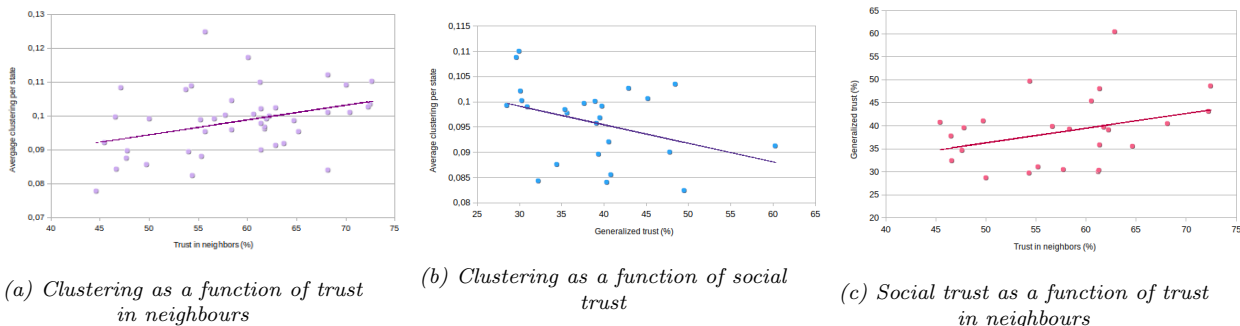


Figure 14: Linear regressions between the 3 variables: social trust, trust in neighbours and social clustering. Each point is a state of the US.

We see that there is an effect (correlation > 0.3), but this effect is limited because the correlation is still small. We see that social trust is negatively correlated with clustering, as expected. We also see that social trust and trust in neighbours are positively correlated, as expected. However, trust in neighbours and clustering are positively correlated, which we did not expect. This is easy to understand because local relationships naturally tend to increase clustering.

The negative correlation between clustering and social trust supports the prediction of our model. As the correlation is small, it is expected that several factors influence clustering and social trust would be one of them. Further studies could investigate these other factors.

4 Discussion

In summary, our exploration of social trust and its impact on social clustering provided different insights through two models. While model one proved to be inconclusive and uninformative, model two demonstrated a clear causal relationship between trust and the degree of clustering within social networks. This finding is particularly significant given that empirical data also supports the correlation between trust and clustering, suggesting that trust is a critical factor in the formation of social groups. However, the empirical correlation, while present, remains modest, suggesting that other factors also contribute to social clustering.

The robustness of model two to a wide range of parameter variations further confirms the strength of the observed effects. This consistency underscores the model’s reliability in reflecting the dynamics of trust and clustering, and provides a solid foundation for further investigation.

A key insight of our findings is that the macroscopic structure of social networks can emerge from individual micro-decisions. Individuals do not intend to influence the structure of social networks. Instead, they seek to protect themselves from cheaters by fostering local trustworthy environments that take advantage of mutual connections. In doing so, however, they inadvertently create a macroscopic pattern in which social networks are fragmented into small, isolated groups. However, this result only occurs when the overall level of trust is low; otherwise, individuals do not need the protection provided by mutual connections.

Understanding the origins of clustering in social networks is crucial due to its extensive socio-economic implications. Clustering affects everything from information dissemination and innovation diffusion to social cohesion and economic inequality. By elucidating the role of trust in fostering clustering, our research contributes to a deeper comprehension of social network structures and the factors that shape them. Future research will build on these findings, aiming to explore additional determinants of clustering and their complex interactions within social systems.

4.1 The big picture

Understanding the factors that influence the shape of the social network is an essential step in understanding the functioning of contemporary societies, because the shape of the social network has many consequences. Therefore, understanding what influences this shape would allow us to identify an additional step in the path leading to social consequences and consequences for cooperation. For example, [Chetty et al., 2022] in the US showed that one’s social capital has a large impact on one’s chances of economic advancement, especially homophily between groups of different socio-economic status (SES).

As mentioned above, many studies support the idea that low social trust may be causally related to low SES. [Bjørnskov, 2007] tried to identify the factors that influence social trust. He showed that social polarisation, such as economic inequalities and cultural diversity, tended to reduce social trust, as did being a post-communist society, while Protestantism and monarchy tended to increase social trust. [Nettle, 2015]’s field observations suggested that deprivation could be a factor in reducing social trust.

Several theoretical arguments support [Nettle, 2015]’s observation. First, it is intuitive that disadvantaged individuals can afford to lose less because they have fewer resources. It is therefore expected that they will behave more cautiously and demand more trust in cooperation because they can afford less risk [de Courson et al., 2024]. On the other hand, the needs and preferences of the disadvantaged are different from those of the wealthy. The trade-off between being honest and benefiting from cheating is obviously not the same if you want a chance of eating today as if you just want some extra money without really needing it. This idea has even been shown to be intuitively integrated, as anger has been shown to be inversely proportional to the actual benefit of the other party. It is therefore normal to trust people less if you know they have an interest in cheating. This idea can be formalised with time preferences [Lie-Panis and André, 2022, Mell et al., 2021]. Their idea is that disadvantaged individuals have a short-term view of their interests because they are trying to survive from day to day or to maintain a fragile situation, while wealthy individuals have a long-term view. This could be explained by the fact that deprived individuals have to spend too much time and energy on their basic needs and have no resources left to think about the distant future. In the long run, it is ad-

vantageous to maintain a good reputation and it is affordable for wealthy individuals to cooperate more. On the other hand, disadvantaged individuals may have more pressing needs than cooperating with many people to build a nice and broad reputation. [Bac, 2009] showed in a theoretical study that wealth makes individuals more confident (but less trustworthy) and poverty makes individuals less confident (but more trustworthy).

Our work focuses on one of a number of causal links that together form a global framework described in Figure 15.

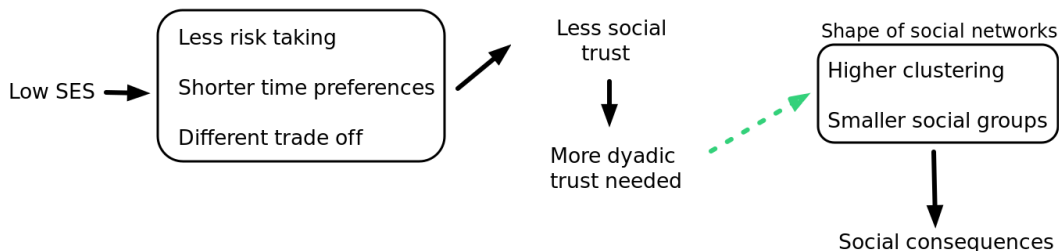


Figure 15: The global framework in which this work fits. Our study focuses on the green dotted arrow.

To summarise the framework, social trust varies across place and time, and this variation can be explained, at least in part, by differences in SES. When social trust is low, individuals must compensate with interpersonal trust through reputation and gossip (premise 2). In a Bayesian view, when individuals’ prior about the reliability of others is too low, they need more information about them in order to cooperate with them. This need for a positive belief that others will cooperate corresponds to reputation-based cooperation (premise 1). We study the consequences of this preference for known cooperators in low trust societies for the shape of the social network. Understanding this effect is crucial because the shape of social networks is known to have socio-economic implications for the population.

Other factors such as culture, the structure of institutions and law enforcement could affect social trust and [Bjørnskov, 2007] identified several of them. As described above, clustering itself has an impact on social trust. But this phenomenon of trust affecting clustering and clustering affecting trust is not a simple loop because many other factors affect social trust, although it does imply a degree of endogeneity. Overall, as the empirical correlation between low social trust and high social clustering shows, many different factors need to be identified.

4.2 Limits

This study presents a model with significant limitations that warrant discussion. Firstly, the simplicity of the model is evident in its small scale, being applied to only 100 individuals. While this sample size allows for manageable computation and clear visualisation of trends, it does not capture the complexities and variabilities present in larger, more diverse populations. Therefore, conclusions drawn from this model may not be directly generalisable to real-world scenarios without further validation on larger datasets.

Moreover, the observed effect, while notable, appears to be an almost trivial consequence of the assumptions made in the model. The primary novelty lies in the transition from individual-level dynamics to macro-level outcomes. This transition, while informative, does not significantly extend our understanding of trust beyond what is already implied by the model’s initial assumptions. The results confirm the expected finding that individual behaviours, when aggregated, produce collective patterns which, in this context, manifest as a macroscopic reflection of trust.

A further point of concern is the degree of clustering observed in the simulations, which is significantly higher than that typically measured empirically. This discrepancy suggests that while the model captures some aspects of trust, it may overestimate the tendency for individuals to form tightly-knit groups. Such

over-clustering could be a consequence of the model’s simplistic assumptions and limited scope, suggesting a need for refinement to better reflect real-world social dynamics.

Future work aims to explore less trivial effects by measuring group size as a function of trust. Group size is a critical metric that can provide richer insights into the social structures and interactions fostered by trust within a community. By focusing on this aspect, we expect to uncover more complex and meaningful implications of trust, such as its impact on the cohesiveness and functionality of social groups. Such analyses could provide a more nuanced understanding of trust dynamics and potentially reveal emergent properties that are not apparent at the individual level.

In summary, while the current model provides a fundamental step in exploring the macroscopic implications of trust, it is inherently limited by its scale and the simple nature of its findings. The next phase of this research will aim to address these limitations by extending the model and exploring more complex phenomena, thereby contributing a more comprehensive perspective on the role of trust in social systems.

4.3 Further work

Future research should focus on three main points. First, we plan to investigate group size as a function of trust. Understanding how trust influences group size will provide deeper insights into the social structures and dynamics within communities. This analysis would lead to less direct effects with a greater theoretical scope.

Second, the effect of clustering on relational mobility could be studied as a direct continuation of this work. Relational mobility, which refers to the ease with which individuals form and dissolve relationships, has significant implications for psychological well-being and behavioural patterns [Yuki and Schug, 2020, Thomson et al., 2018]. By examining how clustering influences relational mobility, we aim to understand the broader consequences of social network structures on individual psychology and behaviour. This line of research is particularly important because relational mobility affects various aspects of social life, including trust, cooperation and social support. The current simulations could easily be used for this purpose.

In addition, the acquisition of more empirical data would be crucial in testing these theoretical ideas. More comprehensive surveys, such as those conducted by [Igarashi and Hirashima, 2021], would provide valuable data to test the robustness of our model and its assumptions.

Appendices

A Robustness tests

A.1 Maximal number of friends is 10 instead of 5

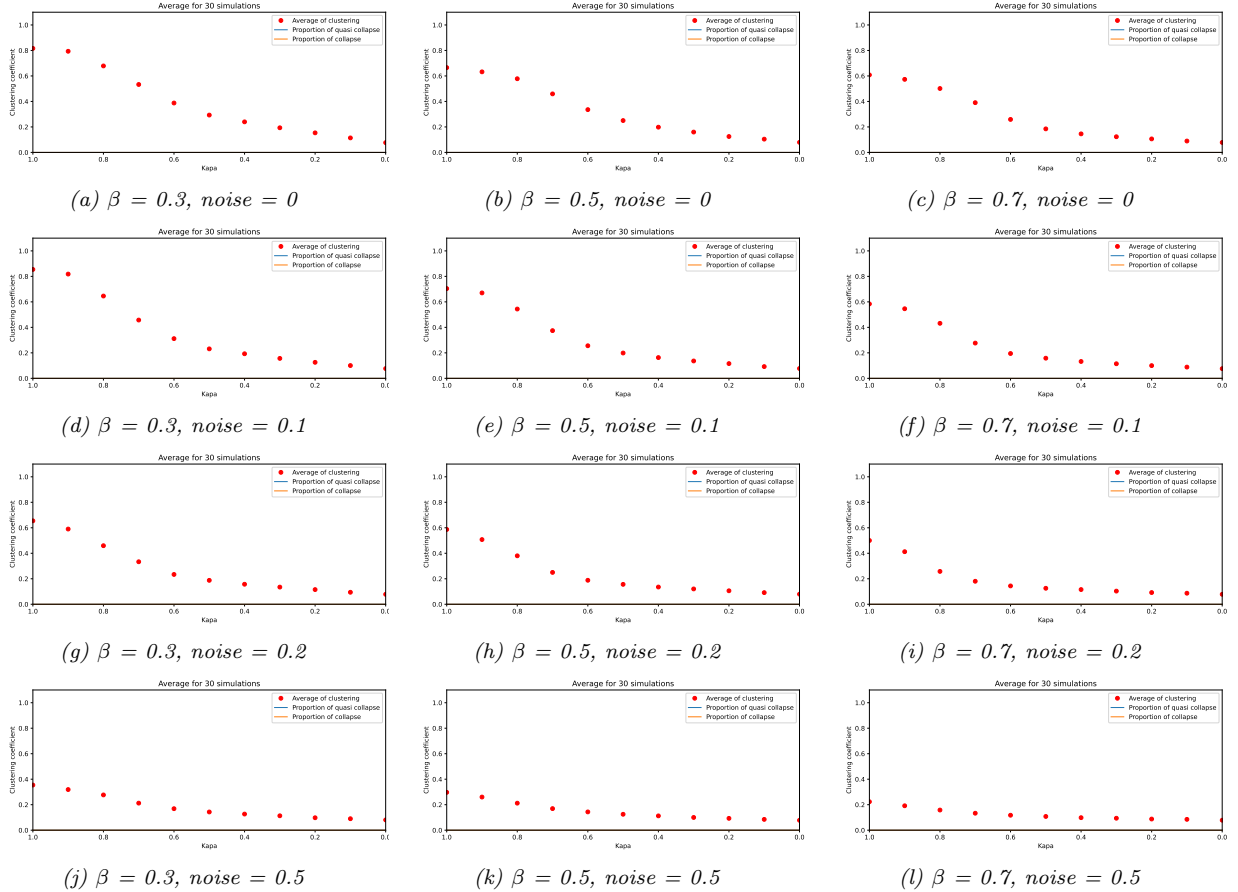


Figure 16: Robustness test 1: Maximal number of friends is 10. Average of the clustering coefficient depending on social trust. On each subfigure, the weight of trust in utility decreases in the left direction. Each point is an average on 30 simulations. From left to right, subfigures correspond to increasing β , i.e. increasing weight given to duration of cooperation over reputational flow in the calculus of trust. From top to bottom, subfigures correspond to increasing noise.

A.2 Lower initial clustering

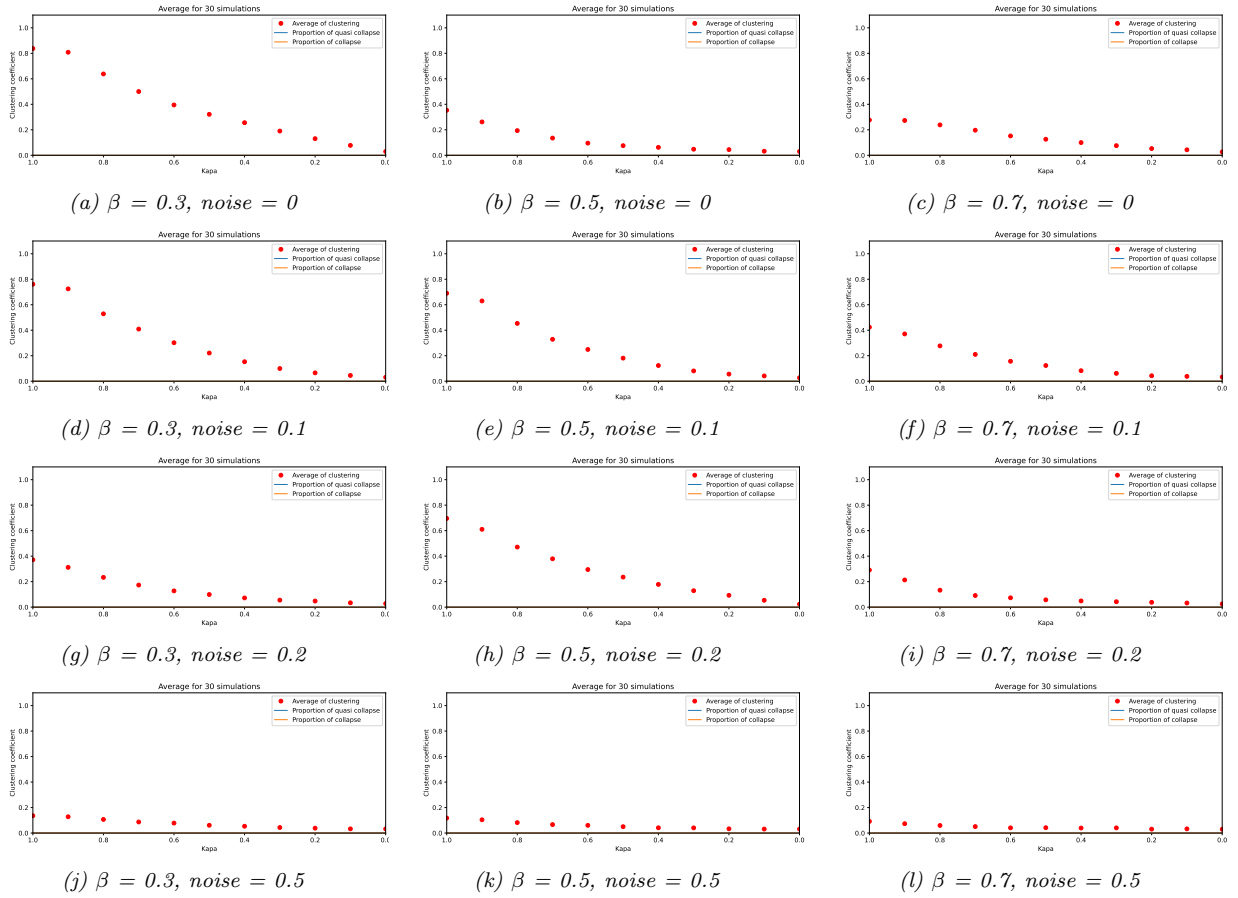


Figure 17: Robustness test 2: Lower initial clustering. Average of the clustering coefficient depending on social trust. On each subfigure, the weight of trust in utility decreases in the left direction. Each point is an average on 30 simulations. From left to right, subfigures correspond to increasing β , ie. increasing weight given to duration of cooperation over reputational flow in the calculus of trust. From top to bottom, subfigures correspond to increasing noise.

A.3 $\beta = 1$ and $\beta = 0$

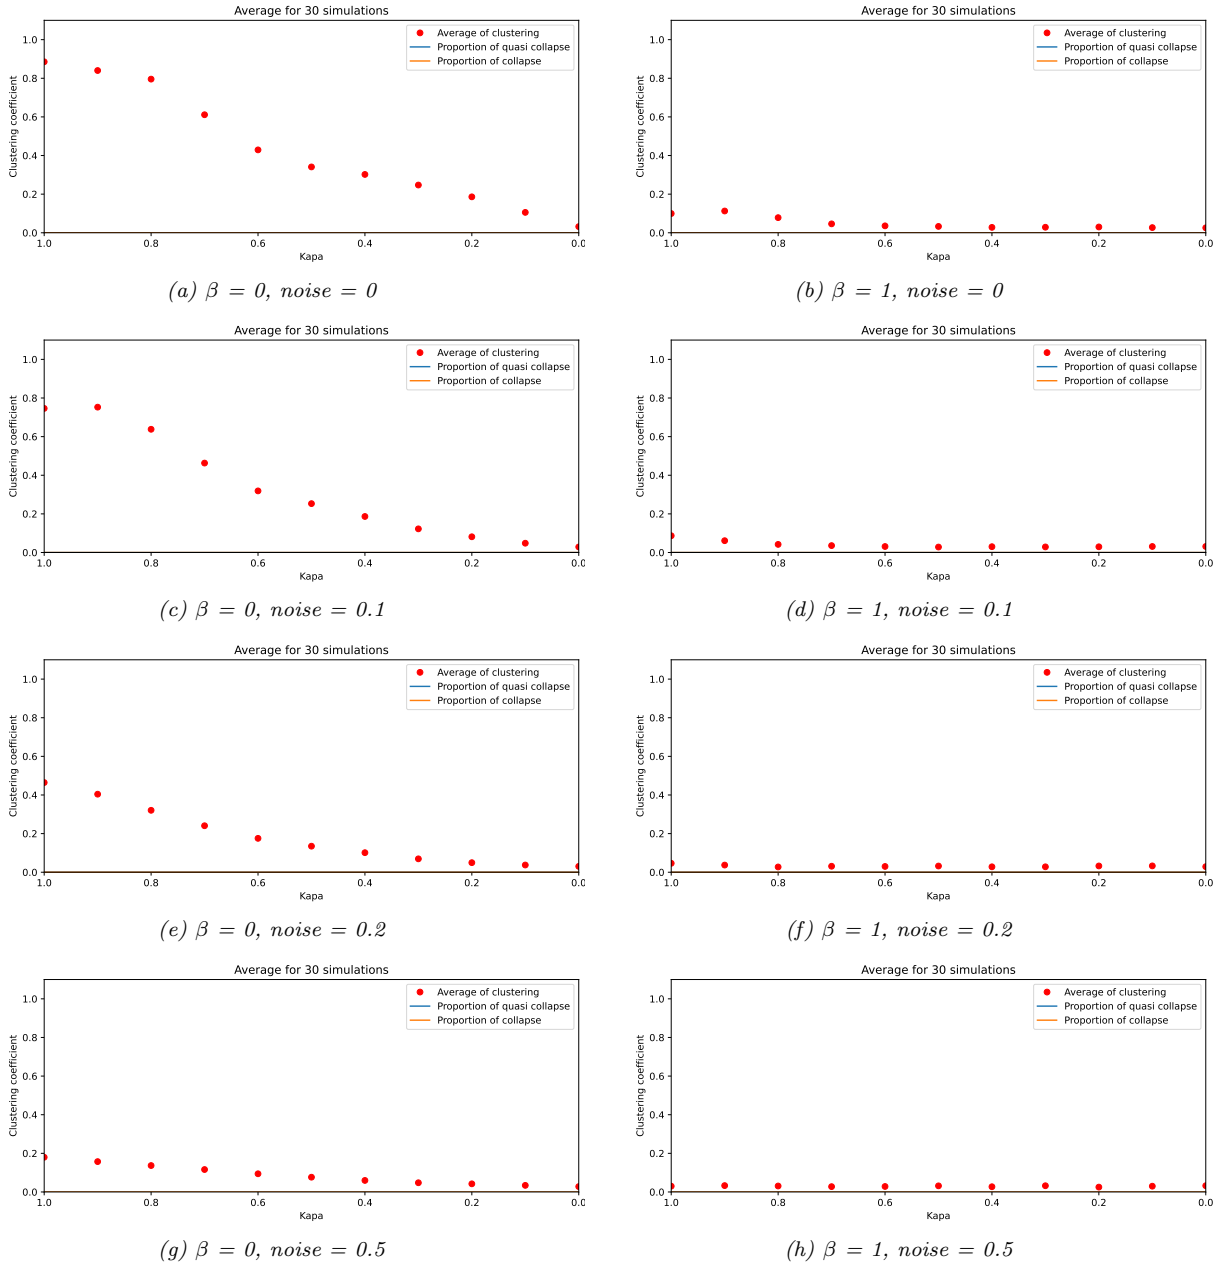


Figure 18: Robustness test 3: $\beta = 0$ (left column) and $\beta = 1$ (right column). Average of the clustering coefficient depending on social trust. On each subfigure, the weight of trust in utility decreases in the left direction. Each point is an average on 30 simulations. From top to bottom, subfigures correspond to increasing noise.

A.4 Initialisation with no initial connection

Figure 4, 5 and 6 show some steps of example simulations for a very low, a middle and a very value of κ respectively.

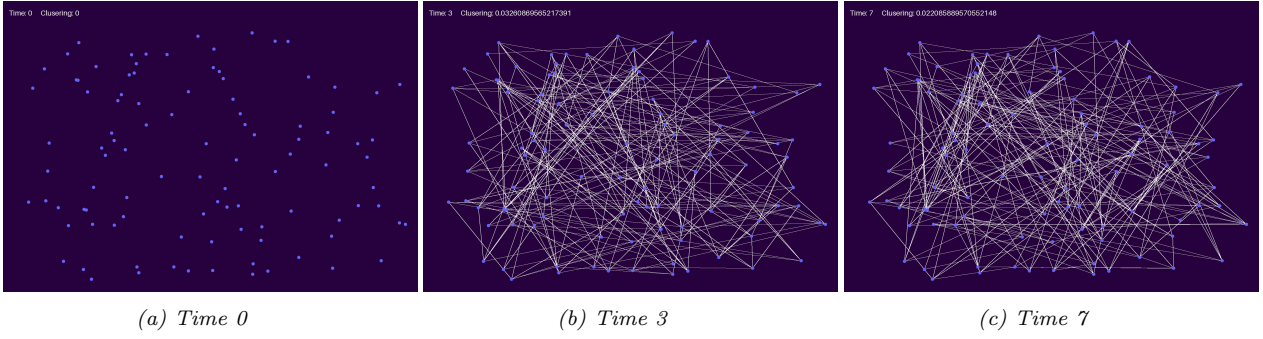


Figure 19: Visualisation of steps 0, 3 and 7 of an example simulation for model 2 with parameters: $\beta = 0.5$, **noise** = 0.1 and $\kappa = 0$ and an empty initialisation.

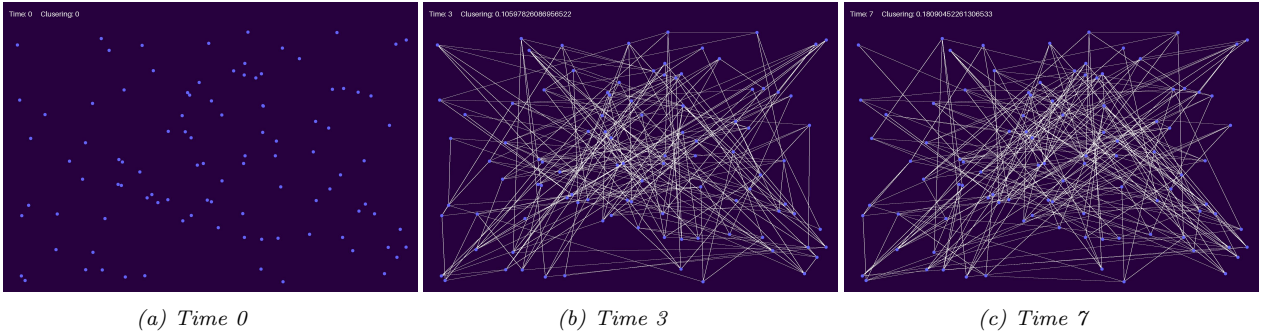


Figure 20: Visualisation of steps 0, 3 and 7 of an example simulation for model 2 with parameters: $\beta = 0.5$, **noise** = 0.1 and $\kappa = 0.5$ and an empty initialisation.

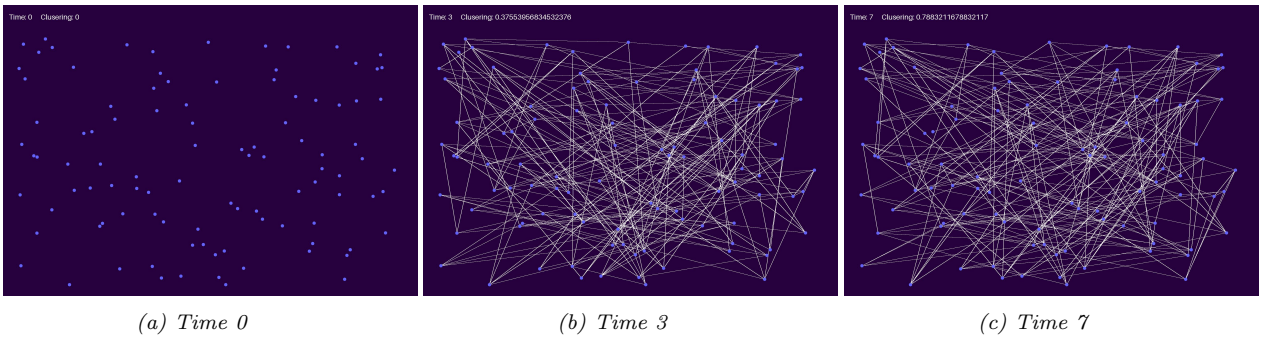


Figure 21: Visualisation of steps 0, 3 and 7 of an example simulation for model 2 with parameters: $\beta = 0.5$, **noise** = 0.1 and $\kappa = 1$ and an empty initialisation.

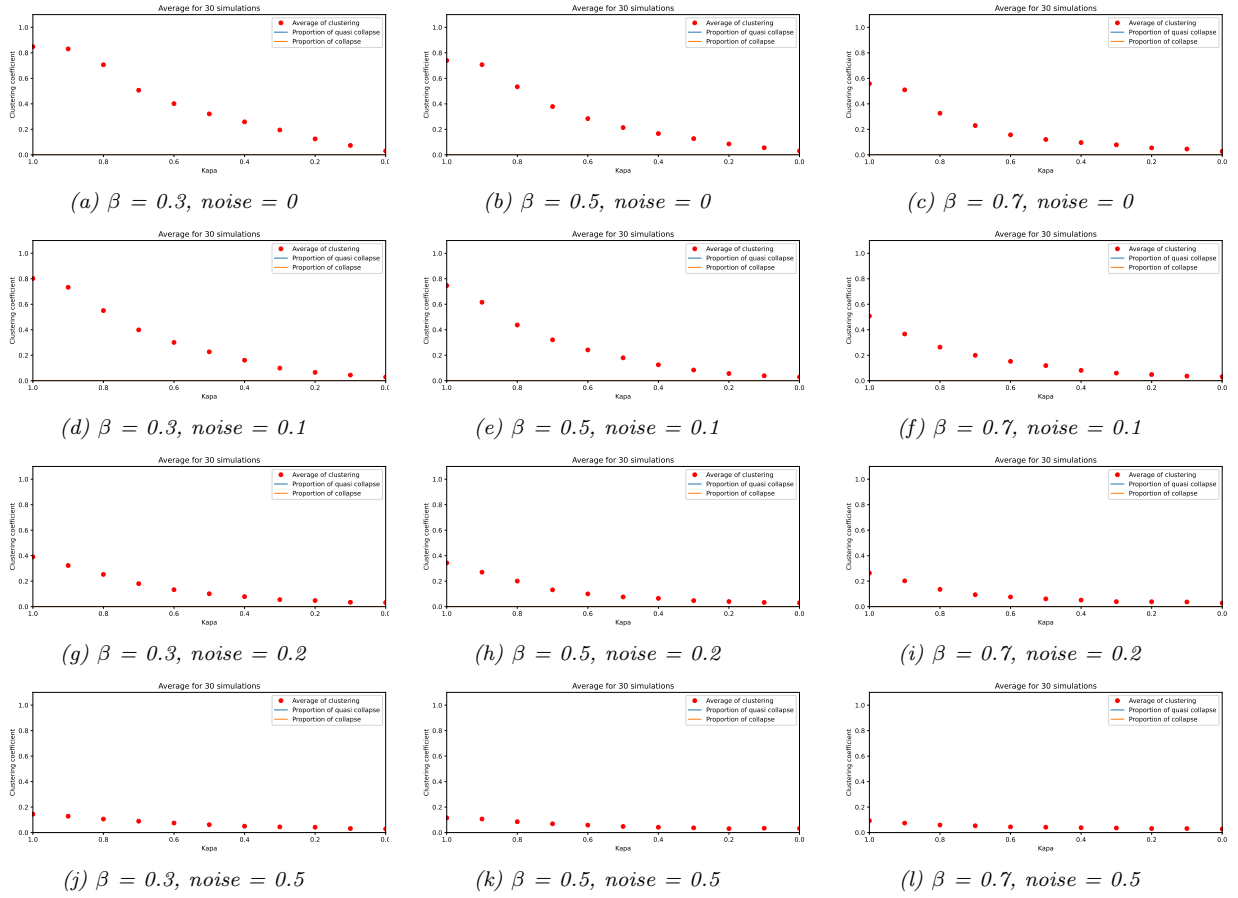


Figure 22: Robustness test 4: Initialisation with no connection. Average of the clustering coefficient depending on social trust. On each subfigure, the weight of trust in utility decreases in the left direction. Each point is an average on 30 simulations. From left to right, subfigures correspond to increasing β , ie. increasing weight given to duration of cooperation over reputational flow in the calculus of trust. From top to bottom, subfigures correspond to increasing noise.

A.5 $\beta = 1$ and $\beta = 0$ and initialisation with no initial connection

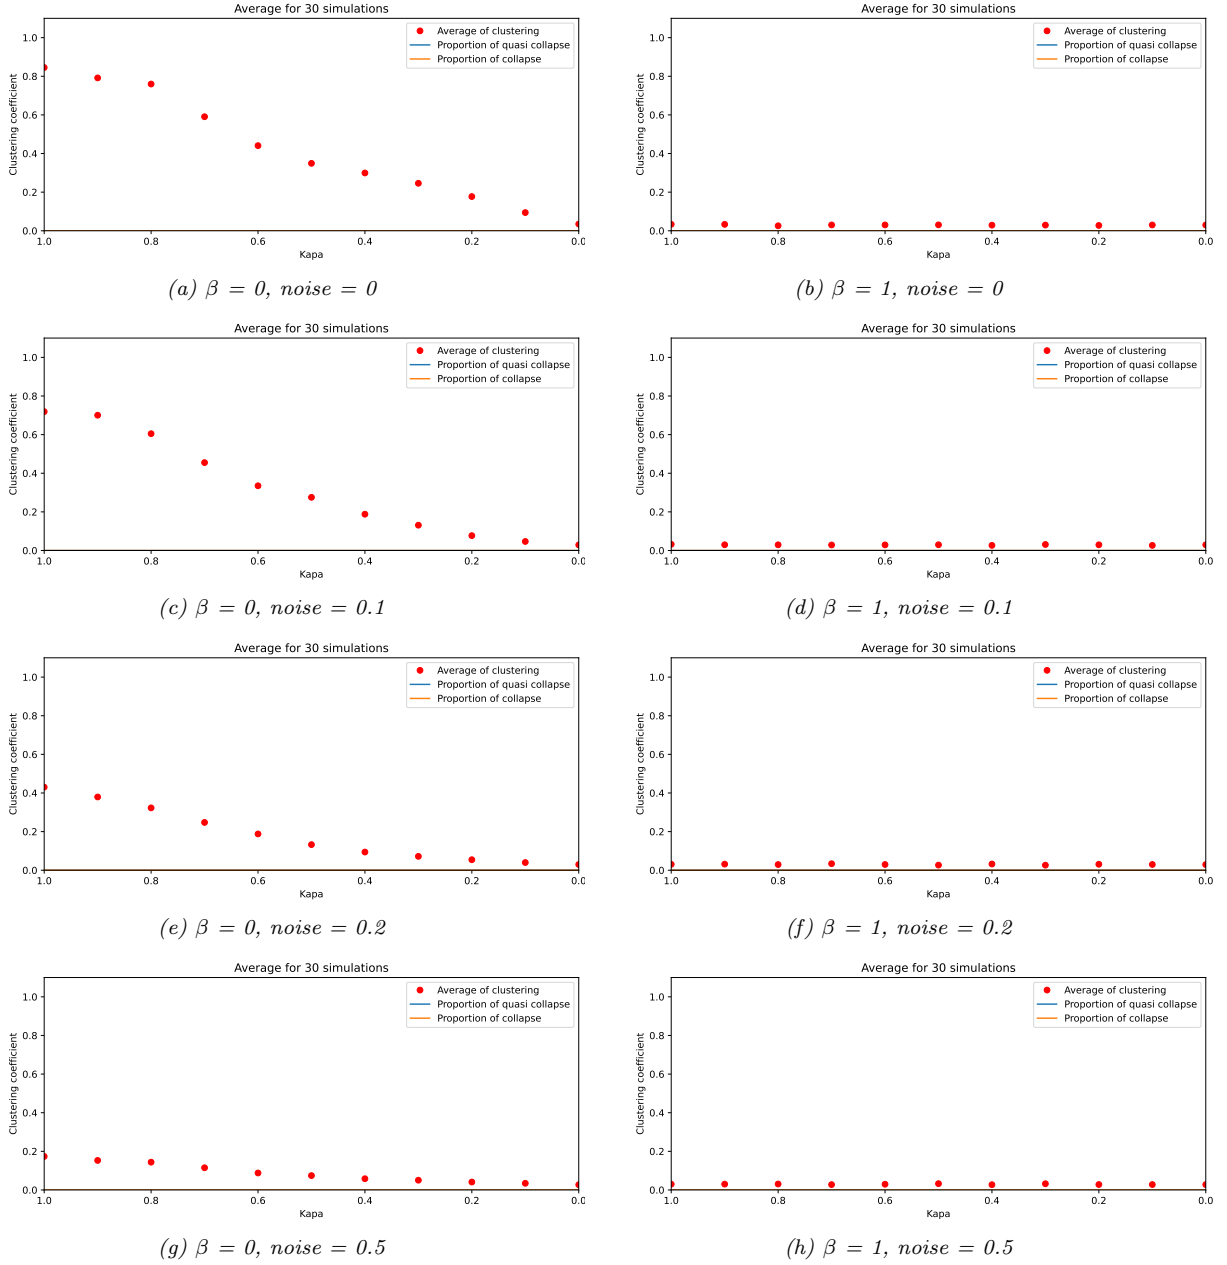


Figure 23: Robustness test 5: $\beta = 0$ (left column) and $\beta = 1$ (right column) and no initial connection. Average of the clustering coefficient depending on social trust. On each subfigure, the weight of trust in utility decreases in the left direction. Each point is an average on 30 simulations. From top to bottom, subfigures correspond to increasing noise.

B Empirical data

Alt FIPS Code	State	% who trust all/most neighbors	% saying people can be trusted most/all of time	average clust
1	Alabama	61,3	30,0	0,109908678192145
2	Alaska	64,2		
4	Arizona	54,5	49,5	0,0823596388537043
5	Arkansas	47,2		0,108195629692707
6	California	49,7	40,9	0,0854348096516008
8	Colorado	68,2	40,4	0,0839756072659212
9	Connecticut	61,9		0,0989799902074064
10	Delaware	54,1		0,0894507730620614
11	District of Columbia	44,6		0,077699997
12	Florida	46,6	32,3	0,084257376002022
13	Georgia	46,5	37,7	0,0995583735371655
15	Hawaii	61,8		0,0960894519252479
16	Idaho	63,8		0,0918146386956322
17	Illinois	61,8	39,6	0,096657409440303
18	Indiana	61,4	30,2	0,102075729297319
19	Iowa	68,2		0,111949239847198
20	Kansas	62,9		0,102418364091478
21	Kentucky	54,3	29,7	0,108671408796828
22	Louisiana	53,8		0,107808537290855
23	Maine	72,7		0,110215528270616
24	Maryland	47,8	39,4	0,089464198446303
25	Massachusetts	61,4	35,7	0,0976294884715934
26	Michigan	60,6	45,3	0,100470473794523
27	Minnesota	72,5	48,5	0,103348770362347
28	Mississippi	55,7		0,124663875032034
29	Missouri	64,7	35,5	0,0984178722068127
30	Montana	68,1		
31	Nebraska	68,2		0,100931676486691
32	Nevada	41,2		
33	New Hampshire	70,4		0,100899439736353
34	New Jersey	56,7	39,8	0,0990497502587158
35	New Mexico	55,4		0,0878628655675534
36	New York	45,4	40,6	0,0919848148468672
37	North Carolina	55,3	31,0	0,0988437416473884
38	North Dakota	71,4		
39	Ohio	57,8	30,3	0,100075226181719
40	Oklahoma	55,7		0,0952584279990383
41	Oregon	62,9	60,3	0,0911963346328829
42	Pennsylvania	62,3	39,0	0,099968868337994
44	Rhode Island	65,2		0,0953251989002623
45	South Carolina	58,4		0,104404255039448
46	South Dakota	73,1		
47	Tennessee	50,0	28,6	0,0991791703026396
48	Texas	47,7	34,5	0,0874235085558394
49	Utah	76,8		
50	Vermont	70,1		0,108944336105002
51	Virginia	58,4	39,2	0,095702860028063
53	Washington	61,4	47,9	0,0898953437794544
54	West Virginia	60,1		0,11720577802519
55	Wisconsin	72,3	43,0	0,102624598982404

Figure 24: Extracted data from [The social capital project, 2017] for social trust and trust in neighbors, and from [Chetty et al., 2022] for clustering.

C ChatGPT

I asked ChatGPT what he was thinking about my internship subject:

Roxane: *What is generalized trust ?*

ChatGPT: Generalized trust refers to the level of trust individuals place in people in general, including those who are not part of their immediate social circle, such as strangers or the broader society. It is a measure of how much people believe that most others in society can be trusted and are generally well-intentioned and honest. This contrasts with particularized trust, which is the trust individuals have in specific people they know personally, like family and friends.

Generalized trust is an important concept in social sciences because it is linked to various positive social outcomes. Higher levels of generalized trust are often associated with stronger social cohesion, better economic performance, more effective governance, and higher levels of social capital. It contributes to the smooth functioning of societies by facilitating cooperation and reducing the need for extensive monitoring and enforcement mechanisms.

Roxane: *How does generalized trust impact the shape of social network ?*

ChatGPT: Generalized trust significantly impacts the shape and dynamics of social networks in several ways:

Network Size and Diversity: High levels of generalized trust tend to result in larger and more diverse social networks. Individuals who trust others broadly are more likely to interact with a wider range of people, including those outside their immediate circles. This leads to more extensive networks that cross various social, economic, and cultural boundaries.

Network Density: In societies with high generalized trust, networks are often more interconnected, with multiple overlapping ties between individuals. This density facilitates the flow of information and resources and fosters mutual support among network members.

Bridge Formation: Generalized trust encourages the formation of bridging ties, which connect different social groups. These bridges help integrate disparate parts of the network, promoting social cohesion and facilitating the spread of ideas and innovations across the network.

Social Mobility: Trust in others can enhance social mobility by reducing barriers to entry into different social circles. Individuals with higher generalized trust are more likely to seek opportunities and collaborations beyond their immediate social groups, potentially leading to upward social and economic mobility.

Cooperation and Reciprocity: Networks shaped by generalized trust are characterized by higher levels of cooperation and reciprocity. When individuals believe that others are generally trustworthy, they are more likely to engage in cooperative behaviors, leading to stronger, more resilient networks.

Conflict and Fragmentation: Conversely, low levels of generalized trust can lead to fragmented networks with tight-knit clusters that have little interaction with each other. This can result in increased social division and conflict, as people may be more suspicious of those outside their immediate groups.

Overall, generalized trust fosters inclusive, expansive, and cohesive social networks that enhance collective action, social integration, and the effective functioning of society.

Roxane: *How does generalized trust impact clustering in social networks ?*

ChatGPT: Generalized trust impacts clustering in social networks by influencing the degree to which individuals form tightly knit groups or clusters. Here are the key ways generalized trust affects clustering:

Reduced Clustering: Higher levels of generalized trust generally lead to reduced clustering. When people trust others broadly, they are more likely to form connections outside their immediate social circles, resulting in more open and expansive networks. This openness decreases the tendency to form tightly knit, insular groups, leading to lower overall clustering.

Increased Interconnectivity: With generalized trust, individuals are more willing to interact with diverse groups of people. This increased interconnectivity creates more bridges between different clusters, integrating various parts of the network. These connections reduce the isolation of clusters and promote a more integrated network structure.

Balanced Clusters: While generalized trust can reduce the overall level of clustering, it can also lead to more balanced clusters. Individuals may still form close-knit groups, but these groups are more likely to have connections to other groups, facilitating a network with both strong internal ties and extensive external links.

Enhanced Information Flow: Networks with lower clustering due to generalized trust often experience enhanced information flow. Trusting individuals are more likely to share information and collaborate with others outside their immediate groups, promoting the dissemination of knowledge and resources across the network.

Diverse Social Capital: Generalized trust supports the development of diverse social capital by encouraging connections across different social, economic, and cultural boundaries. This diversity within the network reduces the homogeneity of clusters and enriches the overall social fabric.

Social Cohesion: Although high generalized trust can reduce clustering, it simultaneously enhances overall social cohesion. The broader connections fostered by generalized trust can help bridge gaps between different social groups, reducing social fragmentation and fostering a sense of community and solidarity across the network.

In summary, generalized trust typically leads to less pronounced clustering in social networks by encouraging connections beyond immediate social circles. This results in more integrated, interconnected networks with better information flow and enhanced social cohesion.

References

- [Acemoglu et al., 2020] Acemoglu, D., Fergusson, L., and Johnson, S. (2020). Population and conflict. *The Review of Economic Studies*, 87(4):1565–1604.
- [Bac, 2009] Bac, M. (2009). Generalized trust and wealth. *International Review of Law and Economics*, 29(1):46–56.
- [Barclay and Willer, 2007] Barclay, P. and Willer, R. (2007). Partner choice creates competitive altruism in humans. *Proceedings of the Royal Society B: Biological Sciences*, 274(1610):749–753.
- [Bjørnskov, 2007] Bjørnskov, C. (2007). Determinants of generalized trust: A cross-country comparison. *Public choice*, 130:1–21.
- [Chetty et al., 2022] Chetty, R., Jackson, M., et al. (2022). Social capital and economic mobility.
- [Coleman, 1988] Coleman, J. S. (1988). Social capital in the creation of human capital. *American journal of sociology*, 94:S95–S120.
- [David-Barrett, 2023] David-Barrett, T. (2023). Clustering drives cooperation on reputation networks, all else fixed. *Royal Society Open Science*, 10(4):230046.
- [de Courson et al., 2024] de Courson, B., Frankenhuis, W., and Nettle, D. (2024). Poverty is associated with both risk avoidance and risk taking: an empirical test of the desperation threshold model. Technical report, Center for Open Science.
- [Floyd, 1962] Floyd, R. W. (1962). Algorithm 97: shortest path. *Communications of the ACM*, 5(6):345–345.
- [Fortunato, 2010] Fortunato, S. (2010). Community detection in graphs. *Physics reports*, 486(3-5):75–174.
- [Fortunato and Hric, 2016] Fortunato, S. and Hric, D. (2016). Community detection in networks: A user guide. *Physics reports*, 659:1–44.
- [Glowacki and Lew-Levy, 2022] Glowacki, L. and Lew-Levy, S. (2022). How small-scale societies achieve large-scale cooperation. *Current Opinion in Psychology*, 44:44–48.
- [Hauert et al., 2002] Hauert, C., De Monte, S., Hofbauer, J., and Sigmund, K. (2002). Volunteering as red queen mechanism for cooperation in public goods games. *Science*, 296(5570):1129–1132.
- [Henrich et al., 2010] Henrich, J., Ensminger, J., McElreath, R., Barr, A., Barrett, C., Bolyanatz, A., Cardenas, J. C., Gurven, M., Gwako, E., Henrich, N., et al. (2010). Markets, religion, community size, and the evolution of fairness and punishment. *science*, 327(5972):1480–1484.
- [Horita and Yamazaki, 2022] Horita, Y. and Yamazaki, M. (2022). Generalized trust rather than perception of relational mobility correlates with nominating close friends in a social network 1. *Japanese Psychological Research*.
- [Igarashi and Hirashima, 2021] Igarashi, T. and Hirashima, T. (2021). Generalized trust and social selection process. *Frontiers in Communication*, 6:667082.
- [Kesternich et al., 2020] Kesternich, I., Smith, J. P., Winter, J. K., and Hörl, M. (2020). Early-life circumstances predict measures of trust among adults: evidence from hunger episodes in post-war germany. *The Scandinavian Journal of Economics*, 122(1):280–305.
- [Kuperman and Risau-Gusman, 2012] Kuperman, M. and Risau-Gusman, S. (2012). Relationship between clustering coefficient and the success of cooperation in networks. *Physical Review E*, 86(1):016104.
- [Lie-Panis and André, 2022] Lie-Panis, J. and André, J.-B. (2022). Cooperation as a signal of time preferences. *Proceedings of the Royal Society B*, 289(1973):20212266.

- [Mell et al., 2021] Mell, H., Baumard, N., and André, J.-B. (2021). Time is money. waiting costs explain why selection favors steeper time discounting in deprived environments. *Evolution and Human Behavior*, 42(4):379–387.
- [Mell et al., 2022] Mell, H., Safra, L., Demange, P., Algan, Y., Baumard, N., and Chevallier, C. (2022). Early life adversity is associated with diminished social trust in adults. *Political Psychology*, 43(2):317–335.
- [Nannestad, 2008] Nannestad, P. (2008). What have we learned about generalized trust, if anything? *Annu. Rev. Polit. Sci.*, 11:413–436.
- [Nettle, 2015] Nettle, D. (2015). *Tyneside neighbourhoods: Deprivation, social life and social behaviour in one British city*. Open Book Publishers.
- [Pacheco et al., 2006a] Pacheco, J. M., Traulsen, A., and Nowak, M. A. (2006a). Active linking in evolutionary games. *Journal of theoretical biology*, 243(3):437–443.
- [Pacheco et al., 2006b] Pacheco, J. M., Traulsen, A., and Nowak, M. A. (2006b). Coevolution of strategy and structure in complex networks with dynamical linking. *Physical review letters*, 97(25):258103.
- [Pacheco et al., 2008] Pacheco, J. M., Traulsen, A., Ohtsuki, H., and Nowak, M. A. (2008). Repeated games and direct reciprocity under active linking. *Journal of theoretical biology*, 250(4):723–731.
- [Powers and Lehmann, 2014] Powers, S. T. and Lehmann, L. (2014). An evolutionary model explaining the neolithic transition from egalitarianism to leadership and despotism. *Proceedings of the Royal Society B: Biological Sciences*, 281(1791):20141349.
- [Raihani and Barclay, 2016] Raihani, N. J. and Barclay, P. (2016). Exploring the trade-off between quality and fairness in human partner choice. *Royal Society open science*, 3(11):160510.
- [Reiter et al., 2023] Reiter, A. M., Hula, A., Vanes, L., Hauser, T. U., Kokorikou, D., Goodyer, I. M., Investigators, N. C. N. P., and staff, N. (2023). Self-reported childhood family adversity is linked to an attenuated gain of trust during adolescence. *Nature Communications*, 14(1):6920.
- [Roberts and Kroese, 2007] Roberts, B. and Kroese, D. (2007). Estimating the number of st paths in a graph. *Journal of Graph Algorithms and Applications*, 11(1):195–214.
- [Roberts et al., 2021] Roberts, G., Raihani, N., Bshary, R., Manrique, H. M., Farina, A., Samu, F., and Barclay, P. (2021). The benefits of being seen to help others: indirect reciprocity and reputation-based partner choice. *Philosophical Transactions of the Royal Society B*, 376(1838):20200290.
- [Robinson and Acemoglu, 2012] Robinson, J. A. and Acemoglu, D. (2012). *Why nations fail: The origins of power, prosperity and poverty*. Profile London.
- [Roth, 2024] Roth, F. (2024). Is generalized trust stable over time? Technical report, Hamburg Discussion Papers in International Economics.
- [Takács et al., 2021] Takács, K., Gross, J., Testori, M., Letina, S., Kenny, A. R., Power, E. A., and Wittek, R. P. (2021). Networks of reliable reputations and cooperation: a review. *Philosophical Transactions of the Royal Society B*, 376(1838):20200297.
- [The social capital project, 2017] The social capital project, J. e. C. R. (2017). An overview of social capital in america.
- [Thomson et al., 2018] Thomson, R., Yuki, M., Talhelm, T., Schug, J., Kito, M., Ayanian, A. H., Becker, J. C., Becker, M., Chiu, C.-y., Choi, H.-S., et al. (2018). Relational mobility predicts social behaviors in 39 countries and is tied to historical farming and threat. *Proceedings of the National Academy of Sciences*, 115(29):7521–7526.

- [Turchin et al., 2022] Turchin, P., Whitehouse, H., Gavrilets, S., Hoyer, D., François, P., Bennett, J. S., Feeney, K. C., Peregrine, P., Feinman, G., Korotayev, A., et al. (2022). Disentangling the evolutionary drivers of social complexity: A comprehensive test of hypotheses. *Science Advances*, 8(25):eabn3517.
- [West et al., 2011] West, S. A., El Mouden, C., and Gardner, A. (2011). Sixteen common misconceptions about the evolution of cooperation in humans. *Evolution and human behavior*, 32(4):231–262.
- [Yuki and Schug, 2020] Yuki, M. and Schug, J. (2020). Psychological consequences of relational mobility. *Current opinion in psychology*, 32:129–132.