

The Topology of Communities of Trust

Mark Alfano

PhD, Associate Professor of Philosophy, Delft University of Technology

Address: Jaffalaan 5, 2628 BX Delft, The Netherlands

E-mail: mark.alfano@gmail.com

Hobbes emphasized that the state of nature is a state of war because it is characterized by a fundamental and generalized distrust. Exiting the state of nature and the conflicts it inevitably fosters is therefore a matter of establishing trust. Extant discussions of trust in the philosophical literature, however, focus either on isolated dyads of trusting individuals or trust in large, faceless institutions. In this paper, I begin to fill the gap between these extremes by analyzing what I call the *topology of communities of trust*. Such communities are best understood in terms of interlocking dyadic relationships that approximate the ideal of being *symmetric*, *Euclidean*, *reflexive*, and *transitive*. Few communities of trust live up to this demanding ideal, and those that do tend to be small (between three and fifteen individuals). Nevertheless, such communities of trust serve as the conditions for the possibility of various important prudential epistemic, cultural, and mental health goods. However, communities of trust also make various problematic phenomena possible. They can become insular and walled-off from the surrounding community, leading to a distrust of out-groups. They can lead their members to abandon public goods for tribal or parochial goods. These drawbacks of communities of trust arise from some of the same mechanisms that give them positive prudential, epistemic, cultural, and mental health value, and so can, at most, be mitigated, not eliminated.

Keywords: trust, distrust, community, flourishing, epistemology, topology, emotion, social epistemology

Introduction

Edwin Curley's 1994 edition of Hobbes's 1668 work emphasized that the state of nature is characterized by a fundamental and generalized distrust. Exiting the state of nature and the conflicts it engenders is therefore a matter of establishing trust. When communities of trust function smoothly, they foster various goods related to the flourishing of humanity. Just how important these communities are becomes clearer when we realize that they are our route out of the state of nature. In the most famous passage of his *Leviathan*, Hobbes compares the state of nature to the state of war (1994: XIII.9):

Whatsoever therefore is consequent to a time of war, where every man is enemy to every man, the same is consequent to the time wherein men live without other security than what their own strength and their own invention shall furnish them withal. In such condition there is no place for industry, because the fruit thereof is uncertain, and consequently, no culture of the earth, no navigation, nor use of commodities that may be imported by sea, no commodious building, no instruments of

moving and removing such things as require much force, no knowledge of the face of the earth, no account of time, no arts, no letters, no society, and which is worst of all, continual fear and danger of violent death, and the life of man, solitary, poor, nasty, brutish, and short.

While we might quibble with the details of Hobbes's portrait, it is clear that it is difficult, if not impossible, to pursue, promote, preserve, and protect certain goods constitutive of flourishing in the state of nature. As Hobbes emphasizes through enumeration, these goods come in several varieties:

- *prudential*: capital ("fruit" of "industry," "commodious buildings," and "instruments of moving and removing"); agriculture ("culture of the earth"); trade ("navigation" and "use of commodities that may be imported"); security (protection against "danger of violent death");
- *epistemic*: science ("knowledge of the face of the earth" and "account of time"); technology ("building," "instruments");
- *mental health*: relief from anxiety ("continual fear");
- *cultural*: "arts"; "letters"; "society."

I contend that Hobbes is right to identify a key psychological mechanism of the state of nature as generalized distrust or what he calls "diffidence" (Hobbes, 1994: XIII.4). The harms and limitations of the state of nature arise in large part because no one can reasonably trust anyone. Life in the state of nature is poor, nasty, brutish, and short, and it is so because it is solitary. If this is right, then exiting the state of nature depends on establishing trust. Unlike Hobbes, I find it implausible that communities of trust are constituted all at once into a Leviathan. Instead, it seems much more naturalistically plausible that they build up from dyads, triads, and other small, tight-knit groups, exactly the small groups theorized here as communities of trust.

Discussions of trust in the philosophical literature, however, focus either on isolated dyads of trusting individuals (Baier, 1986; Jones, 1996, 2012a; McGeer, 2008; Pettit, 1995), or trust in large, faceless institutions or the government (Potter, 2002; Govier, 1997; Townley, Garfield, 2013; Hardin, 2002). In this paper, I begin to fill the gap between these extremes by analyzing what I call *communities of trust*. My focus is on small groups (three to fifteen individuals). Further research will be needed to continue filling the gap in philosophical literature. I focus on the *topology* of communities of trust, understood as directed networks in which the nodes are individual agents and the edges represent dyadic trust. I argue that such communities play an ineliminable role in escape from the state of nature because they furnish the conditions for the possibility of various prudential, epistemic, cultural, and mental health goods.

I describe trust as a dyadic relation between agents, and define the notion of a field of trust, and introduce the concept of a community of trust, which is a set of agents structured into a network through pairwise relations of trust. I explore the topology of communities of trust, addressing extensions and enrichments that make trust in the network more symmetric, reflexive, Euclidean, and/or transitive. I argue that communities of trust

make various problematic phenomena possible. They can become insular and walled-off from the surrounding community, leading to distrust of the out-group, or even outright conflict. They can lead their members to abandon public goods for tribal or parochial goods. These drawbacks of communities of trust arise from some of the same mechanisms that give them positive prudential, epistemic, cultural, and mental health value, and so can at most be mitigated, but not eliminated.

Types of Trust: Dyadic and Communal

Dyadic Trust

Trust is a relation in which a trustor trusts a *trustee* in respect of a *field of trust* (Baier, 1986). For example, when I trust you with the keys to my car, I rely on your competence, goodwill, and good sense in safely and securely operating my car.¹ This involves my believing, or taking it for granted, that you have sufficient skill to drive in ordinary and somewhat extraordinary circumstances. I believe or take it for granted that you can handle my car on both local roads and highways, in sun, rain, wind, and perhaps even with snow and ice. My trusting you with my car also involves believing or taking it for granted that you will refrain from negligent and reckless behavior while driving. You won't get behind the wheel if you're very tired or emotionally overwhelmed. You certainly won't talk on the phone or text while driving, let alone drink and drive. My trusting you with my car also involves a broader assessment of your goodwill and good sense.² I won't give you my keys if I expect you to steal the car or to gamble it away as collateral during a high-stakes poker game. I won't give you the keys if I think you're likely to leave the car unlocked in an area prone to property crime. I won't give you the keys if I think you're the irascible sort who's likely to get into a dispute that leads someone to vandalize it. I won't give you the keys if I think you're disposed to lend them in turn to someone I distrust.

In a trusting relationship, the field of trust is a domain of valued practical concern and activity. In giving you my keys, I trust you with respect to taking care of my valued possession. Trust essentially involves vulnerability. If I didn't care one whit about the car, I couldn't trust you with it. I might lend it to you, give it to you, hand you the keys out of curiosity with where you'll go and what you'll do, and so on. However, it wouldn't be apt to describe any of these as an act of trust unless I also placed some value on the car.

Note that I might say, while giving you my keys, "But don't take it out on the road if the temperature is at or below freezing. You don't have any experience driving in wintry

1. I will take it for granted that the attitude of trust is not simply a belief (Hieronymi, 2008) but something more affect-laden (Faulkner, 2011; Becker, 1996). Because goodwill is a component of trust, "forced trust" (e.g., in loveless shotgun marriages) in which there is no goodwill does not qualify as genuine trust. Thanks to an anonymous re-viewer for raising the issue of forced trust. In addition, I am focusing primarily on trust versus lack of trust, rather than trust versus distrust. For present purposes, we can define distrust analogously to trust, substituting *malice* for *goodwill*.

2. Competence and goodwill are commonly associated with trust in the literature. It will become clear below why I am adding good sense.

conditions.” In doing so, I limit the scope of the field of trust. Negotiating the scope of the field of trust is tricky business for several reasons. First, it’s impossible to spell out all relevant contingencies in advance. If you’re fleeing kidnappers, I will forgive you for driving on icy roads while impaired and using your phone to call or text for help. If someone puts a gun to your head and says, “The car or your life,” I will understand and endorse your giving away the keys. Trying to enumerate and categorize every relevant counterfactual possibility is hopeless. This is one reason why trust also relies generically on the trustee’s goodwill and good sense. It’s also why, in the typical case, the trustee has a rough idea how much the trustor values the field of trust. Secondly, obsessively staking out the boundaries of the field of trust draws attention to the ways I don’t feel comfortable relying on your competence, goodwill, or good sense. This can undermine the warmth of our relationship and lead you to reject my limited trust as offensive or insulting. “Here are my keys. Feel free to take the car out for a spin whenever you like. Unless it’s icy or snowy. Or dark. Or rainy. Or you’re going to a destination you’ve never driven to before. Or you’ve just watched a scary movie. Or . . .” Thirdly, pointing out all the ways in which I might wish to constrain the field of trust may lead you to distrust yourself. Just as trust and hope in another person can be empowering (McGeer, 2008; see also Alfano, 2016; McGeer, Pettit, forthcoming), drawing sharp lines around one’s trust can be disempowering.

The field of trust might involve a valued artifact like a car, but it can also involve human beings, non-human animals, other living things, and more abstract values. I could trust my babysitter to look after my child. I could trust you to feed my pet rabbit while I’m out of town. I could trust you with the password to my email account or my online banking account. I could trust you with my hopes, my secrets, and my sins. In a well-functioning relationship of trust, the trustee *knows* that she’s been entrusted with something valued by the trustor. This is one of the reasons why actively putting your trust in someone is a sign of esteem (Pettit, 1995). Indeed, it is often the case that the trustor knows that the trustee knows that she’s been entrusted with something valuable. Alfano (2016; see also Alfano, 2015; Faulkner, 2011; Nickel, 2012) argues that such interlocking attitudes can lead to self-fulfilling prophecies in which the trustee qualifies as trustworthy in part because she has been trusted, knows it, and knows that the trustor knows it.

The greater the value I place on the field of trust, the more I need to rely on your competence, goodwill, and good sense before initiating such a relation. Maybe I think that you’re generally reliable but prone to bouts of forgetfulness. In that case, perhaps I’d trust you to water my houseplants but not to feed my child or my pet. This indicates that the esteem expressed by trust also comes in degrees: trusting you with something I hold very dear is a sign of higher esteem (or greater desperation) than trusting you with something I value only a little.

In paradigm cases of trust such as deep friendship and marriage, I trust my trustee in an open-ended way that is sometimes referred to as *absolute* or *implicit* trust. In so doing, I trust my trustee to look after my wellbeing and core interests generically and without restriction on the scope of the field of trust. This is presumably one reason why traditional wedding vows go to such lengths to dilate the scope of the field of trust (“in good times

and bad, in sickness and in health”). Such trust is consistent with not trusting my friend or spouse with some peripheral field distinct from my core interests. For instance, I might trust my perpetually-tardy friend to keep my secrets and stand up for my reputation, but not to arrive to a party punctually. I might trust my spouse with just about everything other than the cooking because I know that she is as likely to burn the house down as to prepare something edible and tasty. In cases of absolute trust, I expect my partner often to act in my core interest and rarely to act against it, especially without good reason and without warning me in advance.

Finally, trust has both practical and epistemic aspects. This is clear if we examine cases that might at first blush appear purely epistemic or purely practical. Consider trust in testimony as an example of the former. If I trust you to tell me the truth, I must believe or take it for granted that you are a competent inquirer. However, inquiry is an active and practical process, so trust in testimony also has a practical aspect. Conversely, if I trust you to water my plants while I am out of town, I must believe or take it for granted that you know or will find out how much water is too little, how much is too much, and how much is within an acceptable range. In this case, trusting you to know or find out means that there is an epistemic aspect to my practical trust.

Thus far, I have largely followed the literature inspired by Baier’s seminal article on trust (1986) in viewing it as an attitude that one agent directs towards another agent in respect of a valued field. I have added a good sense criterion to her competence and goodwill criteria; in addition, I have relativized the trust relation to a field of trust defined and delimited by a scope that is sometimes explicitly negotiated, and sometimes only implicitly understood.

Communities of Trust

Now consider cases of trust that involve multiple agents with a common field of practical concern and activity, such as dance partners, a sports team, a happy family, a scientific collaboration, a business, a motorcycle gang, a terrorist cell, or a police department. In such small-scale communities no larger than thirty-five (Dunbar, 1992) and typically fewer than a dozen (Buys, Larson, 1979), multiple people know, rely on, and trust one another in the promotion, preservation, and protection of a mutually-recognized field of trust. At first pass, let us call the agents involved in such a group a *community of trust* with respect to the shared field of practical concern and activity in question, a definition I will refine below. Recent work in social psychology and related fields distinguishes four types of communities (Arrow, 2010):

- (1) cliques, understood as dense clusters of people linked by dyadic bonds,
- (2) coalitions, understood as strategic alignments of people linked by dyadic bonds for mutual advantage,
- (3) comrades, understood as people united through collective action, joint membership in a group, and shared commitments, and
- (4) colleagues, understood as people united by common interests and shared identities.

Since it builds from the notion of dyadic trusting relations, this paper focuses primarily on cliques and coalitions (more so the former than the latter), which are especially amenable to analysis using the tools of social network theory (understanding agents as nodes and relations between them as edges). In their review paper, Katz et al. (2004: 311–12) identify five primary tenets of this approach, four of which are relevant here. First, the network approach is appropriate when people's behavior is best predicted not by their individual characteristics, but by their position in a web of relations. Secondly, the focus of analysis is on relations rather than monadic properties of individuals. Thirdly, conventional assumptions of statistical independence among members of the community do not hold. Fourthly, aggregating dyadic ties alone are not sufficient to explain the behavior of the community or its members: the structure of the network also matters.

Not all cliques and coalitions are communities of trust, but all communities of trust are either cliques or coalitions. Broadly speaking, when such communities of trust function smoothly, everyone does their part in the group's shared field of trust, everyone knows that everyone is disposed to do their part in the group's shared field of trust, everyone trusts everyone to do their part in the group's shared field of trust, and everyone knows that everyone trusts everyone to do their part in the group's shared field of trust. Lack of trust or distrust outside the shared field of concern and activity is compatible with this. I might think my dance partner is a fraud in business or incompetent at the flute. That said, distrust tends to undermine the generic reliance on others' goodwill and good sense, so trusting someone in one sphere tends to be in tension with distrusting them in another.

Such communities needn't be symmetric in all details of their implementation. Symmetry doesn't rule out a division of labor. In a football match, for instance, I can trust the goalkeeper to defend our goal while I rush forward to attack our opponents' goal. However, if I don't trust the goalkeeper, that will lead me to hang back, warping the geometry of our attack and making it less effective. In a scientific collaboration, you can trust me to collect data reliably while I trust you to analyze it carefully. However, if neither of us trusts our third collaborator to write up the results without introducing a huge number of typos and misinterpretations, we are likely to start taking on portions of his task, thereby slowing down the collaboration and potentially leading to a conflict with the third collaborator. In a business, I can trust you to develop new products for our company while you trust me to market them effectively. If we don't trust the driver who delivers our products to not to steal the whole lot of them, our venture will fail for lack of sales. In a police department, you can trust me not to snitch when you kill an innocent civilian, while I trust you not to report my taking bribes and fixing parking tickets, and so on.

Idealizing somewhat, we can say that communities of trust involve a shared understanding of the scope of the field of trust. If I think the scope is much greater than you do, we have a problem. I will expect more of you than you're perhaps prepared or willing to give. You will be likely to do things that flout my expectations, even without failing by your own standards to exhibit competence, goodwill, and good sense. In addition, in communities of trust, we typically have the same level of trust in one another and place the same value on the field of trust (or, barring that, at least appreciate the value placed

on it by the other party). This is important because, as I have explained above, trust typically comes with caveats and exceptions. I trust you to water my houseplants, unless, of course, there is an emergency that keeps you from doing so. What qualifies as a sufficiently urgent problem for you to let my houseplants wither is something that, in the ideal case, we both implicitly or explicitly understand. Unless I know how committed you are to our community and you know how committed I am to it, we're likely to find different caveats appropriate, leading to confusion, conflict, and perhaps even a dissolution of the community.

In addition, members of a community of trust bear one another goodwill, and they often know it and know that they know it. This mutual goodwill, which is established and maintained through activities like social grooming (Dunbar, 1993), laughing together (Dezecache, Dunbar, 2012), singing and dancing together (Dunbar, 2012; see also Slingerland, 2014), or enduring traumatic loss together (Elder, Clipp, 1988), should sometimes help them to resist the temptation to defect in mixed-motive interactions. Moreover, the presence or potential for mutual knowledge of one another's goodwill introduces a second-order motivation to maintain one's reputation (Dunbar, 2005). A one-off defection in a mixed-motive interaction exposes one to loss of reputation, and thereby to exclusion from the benefits of further cooperation and coordination. In a community of (epistemic) trust, reputation-relevant information is likely to travel quickly. Dunbar (1993) estimates that at least 60% of human conversational time comprises of gossip about relationships and personal experiences. When the sense of identity stretches into one's ancestors and descendants, as Richman (2006) argues it does among Jewish merchants in the diamond trade, even the last transaction of one's life redounds not only to one's reputation but also to one's prospects. As Alfano and Robinson (forthcoming) argue, these phenomena make the disposition to gossip well a sort of virtue, being a disposition that protects both oneself and other members of one's community from betrayal while punishing or ostracizing systematic defectors.

This brings us to an epistemic benefit of communities of trust: they are highly effective ways of disseminating knowledge through testimony and other means (Sterelny, 2012). In the case of gossip, the information in question concerns the actions, intentions, and dispositions of another person, but communities of trust are able to transmit not only reputational information, but all kinds of information. In computer science, it has been shown that, depending on the topology of a communicative network, almost everyone gets the message even when the probability of any particular agent gossiping is between 0.6 and 0.8 (Haas, Halpern, Li, 2006). Furthermore, bottlenecks and biases can result in the effective silencing of some people who would like to broadcast a message through the testimonial network (Alfano, Skorburg, 2016). These considerations make the topology of the network an object of essential concern.

Furthermore, in complex inquiries that require specialized intellectual skills, collaboration and teamwork are necessary for knowledge acquisition and transfer. As Hardwig (1991: 694) has emphasized, this means that "the trustworthiness of members of epistemic communities is the ultimate foundation for much of our knowledge." Epistemic

communities need trust in order to coordinate and cooperate effectively in inquiry, to transmit and distribute knowledge amongst themselves, and to communicate it effectively with others.

Mutual knowledge and common knowledge are technical concepts in social epistemology. A proposition is the object of mutual knowledge in a group if everyone knows it. Furthermore, mutual knowledge comes in degrees. We might all know something, but I might not know that you know it. This can lead to breakdowns in coordination and cooperation. Suppose that I am deathly allergic to peanuts. I know about my allergy, as do you, but I don't know that you know about my allergy. Now imagine a scenario in which you kindly bake me some cookies and leave them on my office desk. I arrive at work the next morning to find your thoughtful gift, but I cannot reasonably enjoy it because, for all I know, there are peanuts or peanut oil in the cookies. If I knew that you knew about my allergy and I trusted you not to intentionally or accidentally poison me, I could go ahead and eat the cookies without fear. To coordinate more effectively, we need not just mutual knowledge but also a second degree of mutual knowledge: I need to know that you know about my allergy.³ More complex cases can be constructed in which what is required is third-, fourth-, or even higher degrees of mutual knowledge (Gerbrandy, Groeneveld, 1997). In the limit, we all know, know that we know, know that we know that we know, and so on, out to infinity. This limit is called common knowledge, which, even if it is unrealizable in finite minds, serves as a kind of ideal.

There are three main reasons that communities of trust foster mutual knowledge among their members and help them to achieve multiple orders of mutual knowledge. First, because communities of trust effectively facilitate testimonial knowledge-transfer, they make it likely that any important information in the community eventually makes the rounds. Second, to the extent that the members of the community have at least an implicit understanding of the effectiveness of their own testimonial network, they are in a position to achieve second- and even higher-order levels of mutual knowledge. They can reasonably make "If that were true, I would have heard it by now" judgments of the sort explored by Goldberg (2010), and can go further by making judgments like, "Because that's true, everyone in my community must have heard it by now." Finally, communities of trust tend to have occasional plenary meetings during which they are arranged in an inward-facing circle, making for easy eye-contact among all those in attendance (Chwe, 2001). Such plenary meetings are common evening rituals around the campfire in contemporary small-band hunter-gatherers such as the Ju/'hoansi Bushmen, and were probably also common among paleolithic humans (Wiessner, 2014). Such meetings are ideal settings for public announcements, and can be modeled as bases of mutual knowledge and common knowledge. When a proposition is publicly announced to a group, not only does every member of the group come to know that proposition, but also, because of the publicity of the announcement (e.g., its volume and the fact that everyone can look around and see others paying attention to it), everyone is in a position to know that

3. For a formal study of the epistemic conditions for cooperation, see Aumann and Brandenburger (1995).

everyone knows it, and to know that everyone knows that everyone knows it, and so on (Gerbrandy, Groeneveld, 1997; Baltag, Moss, Solecki, 1999; Pacuit, Parikh, 2004). However, a public announcement can only play this role if it is trusted and everyone in the group expects it to be trusted. If this is right, then communities of trust may be the best and perhaps the only way for humans to systematically generate mutual and common knowledge.

Such communal rituals are also crucially involved in recovery from grief and trauma. For instance, Blustein (2008: 290) argues that rituals of remembrance for family and friends who have died “are essentially and importantly interpersonal. They are undertaken not only in relation to others or in the presence of others, but *with* others in a common action of communalization, and they join the individual in solidarity with others.” If this is right, then there is a bi-directional relation between psychological healing and communities of trust. Therapeutic rituals of remembrance essentially require a community that responds with solidarity and trust. At the same time, as I have mentioned above, communal bonds are created and enhanced through such rituals, and through enduring loss together. The necessity of trust for effective therapeutic rituals is made clearer by Blustein’s discussion of the testimonial interdependence of speaker and hearer. He argues that “bearing witness is associated with finding and registering one’s “voice,” that is, with telling one’s own story *and having it heard in the right way*” (p. 302, emphasis mine). He claims, moreover, that “bearing witness is a type of address to an audience in need, crucially dependent on *trust* in the witness, who has the relevant authority or competence to serve as a witness” (p. 305, emphasis mine).

If the bereaved does not trust his community to respond appropriately to his grief or trauma, he may not bear witness in the first place. This is one reason why victims of sexual assault often do not report their abuse (Sable et al., 2006). By the same token, if the community does not trust the bereaved to tell the story with authority and competence, the ritual is liable to fail and may lead to alienation and re-traumatization. Elizabeth DeVita-Raeburn, the sister of the famous “boy in the bubble” who suffered and died from aplastic anemia in the 1970s, gives poignant expression to this problem in *The Empty Room* (2004). Drawing on dozens of interviews with other bereaved siblings, she quotes “Amber,” whose younger brother died in a car crash when she was nineteen, saying “I felt like there were a lot of outlets for my parents, but I felt like there was nothing for me. People would interrupt my grief to say, ‘Boy, that must be really hard for your parents’” (p. 50). When grief is interrupted in this way or met with “institutional betrayal” of another kind (Smith, Freyd, 2014; see also Shay, 1994, 2003), people who have experienced loss and trauma tend to withdraw into their own suffering. They sever whatever bonds of trust formerly connected them to the community that trivialized or exploited their grief, and the process of mourning must be undertaken alone and often ineffectually.

This pattern is perhaps even more familiar to combat veterans than to civilians. In *Achilles in Vietnam*, Veterans’ Administration psychiatrist Jonathan Shay says that, for his patients, “healing from trauma depends upon communalization of the trauma—being able safely to tell the story to someone who is listening and who can be trusted to

retell it truthfully to others in the community” (1995: 4). Listening in a perfunctory way, he claims, “*destroys trust.*” Veterans in a civilian society face a nearly-insurmountable problem in this regard because people who have never experienced anything like combat tend to find it difficult to empathize, and may respond instead with disbelief, shock, or blame. This all-too-human reaction, like the reaction to Amber quoted in the previous paragraph, leads some veterans to wall themselves off from civilian society. As Shay says, it destroys trust. He goes on to argue that “any blow in life will have longer-lasting and more serious consequences if there is no opportunity to communalize it. This means some mix of formal social ceremony and informal telling of the story with feeling *to socially connected others*” (p. 39).

Although he does not give a detailed explanation for this psychological generalization, Shay does offer a topological insight in *Odysseus in America* (2003), in which he argues for the importance of what he calls “social trust.” Such trust, he says, “requires at least *three* people. Dyadic trust between two people, no matter how many times it is created pair-wise, does not make a community. A community begins with the addition of the third person, and with the belief of *each* individual that, when alone together, the other two will continue to safeguard the interests of each even when that person is *absent*” (pp. 175–6).

It is clear that cultural knowledge, both propositional and practical, depends on communities of trust. We do not have to learn everything from scratch, reinventing the wheel with every generation because we can take advantage of the learning of our contemporaries and predecessors. Naturally, this is most effective when the learner trusts teachers, role models, and other pedagogues. It is even more effective when the teachers trust their students. If, for instance, I don’t trust you not to turn around and kill me with the weapons I show you how to make, I will be, at best, a reluctant teacher. Thus, cultural goods depend on symmetric trust.

Cultural forms related to collective identity also depend essentially on communities of trust. Such identities inform our self-conceptions. They give us a sense of belonging, home, and history (Nietzsche, 1874). They provide us with heroes and villains on whom to model our behavior and practice moral judgments. They help to solidify bonds of trust within a community. The rituals that help to build collective identities often involve the same inward-facing circles mentioned above. Indeed, Wiessner (2014) estimates that upward of 80% of fireside chat among small-band hunter-gatherers concerns topics like *what makes us who we are and how we differ from them.*


Thus, once again, we see that dyadic trust, while an essential building block, is not sufficient to understand how communities of trust work, or fail to work. Let us redefine a community of trust, then, as whatever social structure is the *best path* out of the state of nature. In other words, a community of trust is whatever set or structure of social arrangements is the best way to pursue, promote, preserve, and protect the prudential, epistemic, mental-health, and cultural goods we have been cataloguing in this section. The concept of community as I use is therefore functionally defined in terms of its facility in fostering certain values and goods. When the people involved in such a community

cease to systematically reap the benefits it is meant to furnish, the community is defective. Should such defects persist, it becomes unclear whether there really is a community at all.⁴ In such an eventuality, distrust may be more appropriate than trust (Krishnamurthy, 2015). Furthermore, because relations of trust are interconnected, a breakdown in one part of the network can ramify out into other parts of the network; this makes trust a “fragile commodity” (Dasgupta, 1988). In addition, the normative facet of the concept of community as I use it here is “path-dependent” (Russell, 2015: 105; see also Russell, 2014) rather than “path-independent” in the sense that it takes biological and psychological facts about human animals as a starting point, then attempts to plot a normatively acceptable “best path” (Besser-Jones, 2014) constrained by these facts.

Topology of Communities of Trust

As McLoed points out in her 2015 entry on trust in the *Stanford Encyclopedia of Philosophy*, one of the central philosophical questions about trust is under what conditions it is warranted. McLoed has unidirectional dyadic trust in mind, but we can expand the question to ask what reasons there might be to enrich a minimal network of trust (i.e., a unidirection, or a dyadic relation), either by extending it to a new member, or by increasing its density by establishing new connections within it. To answer this question, we need to reflect on the topology of trust.⁵ Think of individual people as nodes in a directed network, where an edge from node A to node B represents A’s trusting B with respect to some exogenously-given field of trust. Four of the most important properties in such a network are symmetry (being trusted by those you trust), reflexivity (trusting yourself), Euclideanism (your trustees trusting each other), and transitivity (trusting those who are trusted by those you trust). Figure 1 illustrates a minimal network, in which A trusts B, B doesn’t trust A, neither A nor B trust themselves, and no one else is involved.

Figure 1. Basic dyadic trust



In a minimal network like this, the prudential, epistemic, cultural, and mental health goods canvassed above are mostly unavailable. Testimonial knowledge can be transferred from B to A. A may not suffer from the fear of death at B’s hands. A certain amount of coordination and perhaps even cooperation may be possible. Beyond this, however, a minimal network delivers few of the goods that communities of trust are meant to provide.


4. See also Turri (forthcoming) and Miller (2010).

5. Things become more complicated when the field and its scope are endogenously determined. I will not deal with such complex cases in this paper.

Symmetric Trust

One way to enrich the minimal network is to add a link from B to A, making them equal partners (Figure 2).

Figure 2. Symmetric trust



To illustrate this, let the field of trust be navigating from their current location to a place neither has been before by car. A doesn't know how to drive, but does have a facility with reading a map. B is a capable driver and knows how to read a map. It should be obvious how they will proceed: B will do the driving, following the directions given by A. It should also be obvious that this arrangement would not work if B didn't trust A's competence, goodwill, or good sense in giving directions. If B doubted A's competence, she would probably try to do both the navigating and the driving. In the best-case scenario, she would end up a bit stressed-out, but otherwise no worse for wear. We can easily imagine, however, cases in which the increased cognitive load leads B to get lost or to crash the car. Even if neither of these bad outcomes results, A might be insulted by B's lack of trust and withdraw his own trust from B, perhaps leading him to abandon the joint trip. If B doubted A's goodwill (for instance, she thought A might direct her to some other location, such as a destination he preferred, or an isolated spot where he could rob her), B would presumably refuse to embark in the first place. Such a rebuff is liable to lead A to withdraw his trust from B, disintegrating the minimal network. If B doubted A's good sense (for instance, she thought A was likely to be distracted from the task of navigating, tempted by scenic routes, or indulgent towards creepy hitch-hikers), she might decide to embark but resolve to double-check every instruction B gives her. Once again, such an arrangement could work out, but it is fragile. It is not hard to imagine A becoming annoyed or resentful over the incessant double-checking and demonstrations of contempt of his good sense.

Coordination and cooperation between A and B will be greatly enhanced and made more modally robust if they both trust each other, rather than A trusting B but not receiving reciprocated trust. What might warrant B's trusting A? There are three arguments, one for each component of trust (competence, goodwill, and good sense). First, competence in a domain is highly associated with meta-competence in making judgments about competence in that domain (Collins, Evans, 2007: ch. 2). Pianists are better than the average person at judging the expertise of pianists. Medical doctors are better than the average person at judging the expertise of medical doctors, especially within their own specialization. If A trusts B to do the driving, and B thinks that this trust is well-

placed, then B has some evidence that A is capable of identifying competent drives. This in turn suggests that he might know enough about transit to competently read a map.

Secondly, B could have *pro tanto* reason to believe in A's goodwill if they have engaged in some of the bonding activities discussed above, such as laughing together, singing and dancing together, enduring trauma or grief together, or sharing a mutually recognized collective identity. In addition, B could have memories of A's past conduct in similar situations or testimonial knowledge (derived from her network of trusted gossipers) about A's track-record in similar situations with other people.

Thirdly, B might be warranted in trusting A by her own sense of fairness or reciprocity, or by her faith in humanity (Preston-Roedder, 2013). If someone has made it clear that they are willing to make themselves vulnerable to B, then she may naturally feel that it is only fair for her to make herself vulnerable to them. Likewise, the optimism associated with faith in humanity may lead B to think that anyone who, like A, is open with their trust deserves the benefit of the doubt, giving her warrant to return A's trust. These considerations, like the others just considered, clearly are not decisive, and there may be additional considerations that defeat them. Nevertheless, something like default reciprocity or a faith in humanity may also play a role in warranting the construction of symmetric trust.

Reflexive Trust

There is a remarkable near-consensus that, unless you have particular reasons to the contrary, you ought to trust yourself. For example, Pasnau (2015) argues that self-trust justifiably influences how we should react to peer disagreement. Lehrer (1997) argues that self-trust grounds reason, wisdom, and knowledge. Govier (1993) argues that self-trust grounds autonomy and self-respect. Jones (2012b) positively evaluates self-trust from a feminist perspective. Goldberg (2013) argues that self-trust is a good model for trust in others. Completely isolated self-trust may be psychologically possible, but it is unlikely that any human animal has ever developed it. The main questions for us in this section are, first, whether someone's trusting themselves or not provides additional warrant for another person to trust them, and secondly, whether being trusted by someone provides additional warrant to trust yourself.


Let us return to our navigation example above, in which A trusts B to drive to their destination while B trusts A to navigate. Suppose that A experiences a moment of self-doubt. He starts to worry that he isn't a competent navigator after all. In such a case, B's explicitly trusting him could function in a therapeutic way, giving him reason to restore his own self-trust. Next, consider the truism that one's possibilities for action are constrained by one's modal knowledge and beliefs. If you think that something is impossible, even if it's not, you can't try to accomplish it. A's impression of his own possibilities for action (and thus the range of actions he can actually take) is expanded by B's confidence in him. When she signals that she trusts him to do the navigating, she opens up the possibility of navigating for him (Alfano, Skorburg, forthcoming). If this is right, then thera-

peutic trust can be effective by influencing someone's confidence, and thereby their competence. However, it can also work through the goodwill component of trust. As McGeer (2008) reminds us, human motivation is often complicated and confusing. Sometimes we don't know what we really desire, value, or love. In those cases, it is helpful to refer to a normative lodestone, a model of good conduct. According to McGeer (pp. 248–9),

For help in this regard, we are sometimes encouraged to look outside ourselves for role models, finding in others' thoughts and actions laudable patterns on which to fashion our own. And this may serve us pretty well. However, something similar can occur, often more effectively, through the dynamic of hopeful scaffolding. Here we look outside ourselves once again; but instead of looking for laudable patterns in others' behavior, what we find instead are laudable patterns that others see—or prospectively see—in our own. We see ourselves as we might be, and thereby become something like a role model for ourselves. The advantage in this is clear: Instead of thinking, "I want to be like her,"—i.e., like someone else altogether—the galvanizing thought that drives us forward is seemingly more immediate and reachable: "I want to be as she already sees me to be."

Starting from a symmetric network like the one picture in figure 2, therapeutic trust may lead to reflexive trust as pictured in Figure 3.

Figure 3. Symmetric and reflexive trust




Next, consider the question of whether someone's trusting themselves could provide warrant for extending trust to them socially. This is a more difficult case. We can easily imagine a wildly overconfident person who has total self-trust, and who should therefore be avoided rather than trusted. By the same token, it would be imprudent to put your trust in someone who actively distrusted their own competence or strength of will to resist self-serving temptations. In between these extremes, someone who has measured trust in their own abilities and willpower might reasonably inspire trust in others. This perspective fits well with research by Justin Kruger and David Dunning (1999) which indicates that people with middling competence are fairly reliable in estimating their own competence, while highly incompetent people tend to overestimate their own abilities. If this is on the right track, then the second question about reflexivity receives a qualified positive answer.

Euclidean Trust

Consider next the networks of at least three individuals. What reasons might there be by to enrich such a network so that it is (closer to being) Euclidean? We can begin with the simple case illustrated in Figure 4.


Figure 4. Symmetric, but non-Euclidean trust



Under what conditions is it warranted for B to trust C, and C to trust B? At first blush, it might seem like this will be an especially troublesome link to establish. When A forms a symmetric link with C, B might be tempted to ask what makes this new friend so special. Rivalry, competition, and jealousy loom. However, in many cases, B and C will have good reason to establish relations of trust with one another, primarily in order to ensure that A remains trustworthy. To see why, consider the somewhat more complex “star-network” pictured in Figure 5.


In his seminal work on social networks, Freeman (1978) described the “star-network,” in which all other nodes are connected to a single central hub but not each other, and is the most unequal topology possible. Sexual predators and their targets often form a star-network, with the predator at the center and the victims on the points of the star. This keeps the victims from effectively communicating with one another, and coordinating or cooperating against the predator (Fire, Katz, Elovici, 2012). Star-networks are also associated with financial fraud (Šubelj, Furlan, Bajec, 2011), academic fraud (Callaway, 2011), and terrorist activities (Reid et al., 2005; Krebs, 2002). In a recent paper, Savage et al. (2014) categorize star-networks and near-star-networks in an effort to make their online detection more effective. Mutual and common knowledge, especially about problematic behavior by the hub of a community, is greatly enhanced by the gossip that occurs in Euclidean connections in that community.

Figure 5. Star-network



Naturally, not every relationship is in imminent danger of turning into sexual harassment, fraud, or terrorism, but for many sensitive relationships, these considerations suggest that making networks more Euclidean is warranted and may increase warranted trust in other members of the network (e.g., B's trust in A, in Figure 6).


Figure 6. Symmetric and Euclidean trust



Transitive Trust

This leaves us with transitivity, which is extending one's trust out into a network of trust. What reasons might there be to do this? If that network is already characterized by symmetry and reflexivity, then adding transitivity will make trust an equivalence relation in the community in question, as illustrated in Figure 7.

Figure 7. Trust as an equivalence relation




The philosophy of emotion tends not to bother much about logic, but in computer science, trust between computers is formally modeled as an equivalence relation (Kramer, Goré, Okamoto, 2010). As we saw above in Shay's definition of social trust, he believes, on the basis of clinical experience, that recovery from grief and trauma is only possible in a context where trust extends transitively for at least one step. These observations suggest that we have reason to find or construct small communities that have this characteristic.

Small-scale communities in which everyone knows everyone can sustain the literal transitivity of trust. As the community increases, the need for vicarious or mediated trust increases with it. X vicariously trusts Y through M with respect to the field of trust of F, just in case X trusts M with respect to F, M trusts Y with respect to F, and X trusts M's judgment about who is trustworthy with respect to F. Vicarious trust has a distinctive counterfactual signature in the sense that, if X vicariously trusts Y through M, then were X to become directly acquainted with Y, X would continue to trust Y non-vicariously. We can think of this in terms of *delegation* (empowering someone to extend your trust vicariously) and *ratification* (explicitly confirming an instance of delegation). In cases where an acquaintance with Y leads X to withdraw rather than ratify her vicarious trust in Y, she may also begin to doubt M: perhaps M lacks goodwill or good sense after all. To illustrate

this, suppose that I trust you with my secrets. You trust your sibling with your secrets, but I explicitly distrust your sibling. If I find out that you tell your sibling everything, I'm going to be motivated to avoid telling you things I don't trust her with. Or, suppose my boss trusts me to complete a task, and that I sub-contract out a part of that task to someone she distrusts. Again, if she finds out that I've done this, she will most likely withdraw her trust from me, at least regarding this task and perhaps more, generally. She might conclude that I lack goodwill or good sense. Or, she might question my competence to perform the task in question, given that I sub-contracted it to someone she considers untrustworthy.

Shy of such a highly demanding approach to transitivity, we might ask about extending one's trust one or two steps out into a community (Figure 8). What reasons are there for C to trust D, who is trusted by someone she trusts (B)?

Figure 8. Transitive extension of trust



In addition to *delegation*, we might focus on the phenomenon of *vouching*. B vouches for D to C if B makes himself accountable for any failure on D's part to prove trustworthy, either via outright betrayal or via incompetence or lack of good sense. Thus, one reason for C to trust D is if B vouches for D. How is such vouching meant to work? It relies on the (apparent) rationality of extending trust transitively. If C trusts B and B trusts D, then C has a reason to trust D.

This reason needn't be compelling. C can withhold her trust from D even as she gives it to B. As above, my hypothesis is that transitivity provides a *pro tanto* reason to extend trust, not an all-things-considered reason. Why would it do so? There are three arguments, one for each component of trust (competence, goodwill, and good sense). First, as we saw above, competence in a domain is highly associated with meta-competence in making judgments about competence in that domain (Collins, Evans, 2007: ch. 2). If C

trusts B, that means C deems B competent or even an expert with respect to their shared field of trust. It stands to reason, then, that C should expect B to be better than the average person at judging the competence of others in that field. So if B gives his trust to D, C has a reason to think that D is competent.

Secondly and thirdly together, it is psychologically difficult and perhaps even practically irrational to consciously engage in efforts to undermine your own values in the very process of pursuing and promoting those values. Imagine someone locking a door while they're trying to walk through the doorway. Someone could do this by mistake, but then they lack good sense. Someone could do it as a gag, or in pretense, but it is hard to envision a case in which someone does this consciously. Likewise, it is hard to envision a case in which someone genuinely bears you goodwill, and consciously expresses that goodwill by recommending that you put your fate in the hands of someone they expect to betray your trust. They might do so by mistake, as a gag, or in pretense, but a straightforward case is difficult to imagine. If C trusts B, that means C judges that B bears her relevant goodwill. It stands to reason, then, that C should expect B to act on that goodwill in a practically rational way. So if B gives his trust to D, C has a reason to think that D would act consistently with B's goodwill. Putting these together, if C trusts B and B trusts D, then C has a reason to think that D is competent, bears (people like) C goodwill, and has good sense. In other words, C has *pro tanto* reasons to trust D.

Drawbacks of Communities of Trust

Returning to Hobbes's taxonomy of the goods furnished by communities of trust, I shall now argue for a more systematic taxonomy, structured as follows:

- *prudential*: coordination; cooperation;
- *epistemic*: testimonial knowledge; collaborative inquiry; mutual and common knowledge;
- *mental health*: recovery from grief and trauma;
- *cultural*: art & letters; collective identity.

I understand coordination in terms of approximation of Nash's 1951 work on equilibrium and cooperation in terms of mutual benefit (refusing to defect) in mixed-motive contexts. Ignoring transaction costs, people can only benefit from coordination, whereas cooperation is beneficial on the whole but sometimes requires partners to incur costs. It should be immediately clear from what has already been said that communities of trust foster both coordination and cooperation. As I have explained above, members of communities of trust have a shared understanding of the scope of the field of trust and confidence in one another's competence in respect of the field of trust. This makes it possible for them to coordinate their activities harmoniously, in turn allowing them to reap such benefits of coordination as Hobbes mentions (capital, agriculture, trade, or security).

However, communities of trust make various problematic phenomena possible, such as collective betrayals and alleged collective betrayals. They can become insular and walled-off from the surrounding community, leading to distrust of the out-group. They

tend towards nepotism, cronyism, and other forms of parochial altruism. These drawbacks of communities of trust arise from some of the same mechanisms that give them positive prudential, epistemic, cultural, and mental health value, and so can at most be mitigated, though not eliminated.

Perhaps the best fictional reflections on the dynamics, functions, and pitfalls of communities of trust are zombie apocalypses. This is because in a zombie apocalypse, the characters lose access to the infrastructural and institutional benefits (and harms) that pervade contemporary life. They return to a state of nature which is exceedingly dangerous. There are all of the ordinary difficulties associated with acquiring access to scarce resources like water, shelter, and food. There are ordinary threats like bad weather, wild animals, and disease. There are, of course, the zombies themselves, which elevate the level of existential threat and engender distrust even among comrades who have to worry at all times that one among them has been bitten (and either doesn't know it or is hiding the injury). Perhaps most importantly, there are other bands of survivors who have the same fears and compete for the same resources. Should outsiders be incorporated into the group, potentially strengthening it while increasing its needs and vulnerability to betrayal? Should other groups be resisted, merged with, tolerated as independent entities, or raided for resources? These questions arise naturally in this genre of fiction, which can help us to vividly imagine what it is like to be in and trying to escape from the state of nature.

Furthermore, these questions revolve around the question of trust, so it should come as no surprise that *The Walking Dead* (a television series based on a critically-acclaimed series of graphic novels), a recent exemplar of the genre of the zombie apocalypse, explicitly references the necessity and the difficulty of establishing trust within a community and of extending trust to outsiders. For example, in Season 4, Episode 1 (Gimple, Nicotero, 2013), a central character named Rick Grimes institutes a new ritual. Potential recruits of his beleaguered community must satisfactorily answer three questions: "How many walkers [zombies] have you killed? How many people have you killed? Why?" As I have pointed out above, trusting someone involves believing or taking it for granted that they have competence, goodwill, and good sense. In the context of a zombie apocalypse, allying with someone who is incompetent at killing zombies could be deadly for oneself and one's community. Such a person is liable to undermine coordinated efforts at mutual defense. Beyond that, they are more likely to become infected themselves, putting the community at risk from an internal threat. The first question Grimes asks addresses competence. If you have haven't killed enough zombies, you're probably incompetent.

The second and third questions address goodwill and good sense. Murder is *prima facie* wrong and contrary to goodwill, but in a state of nature such an act may be prudentially required to protect oneself, or morally required to protect members of one's community from raiders. In a state of nature, someone who claims not to have killed any other people faces suspicion of lacking good sense, or of lying. Either way, they're not trustworthy. In a state of nature, someone who cannot provide an inter-subjectively acceptable answer to the question "Why?" faces suspicion of lacking goodwill. Note that

the answer might be deemed unacceptable in virtue of its content (e.g., “He looked at me funny.”), or in virtue of the manner of its delivery (e.g., responding too quickly, too slowly, too tersely, too longwindedly, or with shifty eyes). As Slingerland (2014) argues, hard-to-fake embodied signals are essential to establishing trust in interpersonal contexts.

In Season 5, Episode 12 of *The Walking Dead* (Powell, Nicotero, 2015), the tables are turned: Deanna Monroe, the leader of a well-resourced community called Alexandria, is considering whether to recruit Rick Grimes’s group. Here is a telling exchange between the two leaders:

RICK GRIMES: You put up the wall [around Alexandria]?

DEANNA MONROE: Well, there was this huge shopping mall being built nearby. And my husband Reg is a professor of architecture. . . . He got the first plates up with our sons. And after a few weeks, more people arrived and we had help. We had a community.

RICK GRIMES: You’ve been behind these walls this entire time?

DEANNA MONROE: We need people who have lived out there. Your group is the first we’ve even considered taking in for a long time.

RICK GRIMES: You should keep your gates closed.

DEANNA MONROE: Why?

RICK GRIMES: Because it’s all about survival now. At any cost. People out there are always looking for an angle. Looking to play on your weakness. They measure you by what they can take from you. By how they can use you to live. So bringing people into a place like this now . . .

DEANNA MONROE: Are you telling me not to bring your people in? Are you already looking after this place? Aaron [a recruiter for Alexandria] says I can trust you.

RICK GRIMES: Aaron doesn’t know me. I’ve killed people. I don’t even know how many by now. But I know why they’re all dead. They’re dead so my family, all those people out there, can be alive, so I could be alive for them.

DEANNA MONROE: Sounds like I’d want to be part of your family.

Even though he hasn’t been asked the three questions, Grimes volunteers answers to the second (“How many people have you killed?”) and third (“Why?”). Perhaps without meaning to, he convinces Monroe that he is trustworthy. In doing so, he induces her to extend her trust and, to the extent that she has power over it, the trust of her whole community to Grimes and his group.

Thus, zombie apocalypse stories are perhaps the best fictional reflections on the dynamics, functions, and pitfalls of communities of trust in a hostile world. In the Alexandria plot, the community is literally walled off from the outside world. This is a natural reaction for communities of trust. After all, one of the benefits of their coordination and cooperation is the accumulation of capital, including movable capital. This makes them a target of opportunity for other individuals and communities. A successful community of trust, unless its coordination and cooperation somehow produces only infinite or intrinsically non-transferrable goods, will face this challenge. Assuming they recognize the

challenge, they will inevitably have reasons to distrust outsiders and make preparations for their mutual defense.

Next, as I have mentioned above, one of the psychological mechanisms that relies on communities of trust and enhances the bonds within them is a sense of collective identity. Collective identities tend to be contrastive. They primarily concern what we are like, but they often also concern what *they* are like, where the third-person pronoun references one or more out-groups. For obvious reasons, such contrasts tend to be invidious, leading to or reinforcing biases against out-groups and their members. In addition, as Dunbar (1993) points out, one of the main aims of cooperation in communities of trust has been military success, from the Roman maniple to the current day. Violent inter-group conflicts inevitably lead to shared traumatic loss, a key psychological mechanism for building communities of trust. Such conflicts naturally also lead to and reinforce hostility to other groups. Finally, recent works by De Dreu, Balliet, and Halevy (2014) suggest that one of the psycho-biological mechanisms of trust-formation within a group is the hormone neuropeptide oxytocin.⁶ Oxytocin is associated with both parochial cooperation, which is oriented towards the in-group and explicitly exclusive towards others, and derogation of and hostility towards members of rivaling out-groups (*ibid.*: 4).

If these cautionary points are on the right track, then communities of trust, despite their desirability, should be approached with ambivalence. They make possible a basket of essential prudential, epistemic, cultural, and mental health values, but the mechanisms that foster them ineliminably engender regrettable conflicts. As such, the benefits of communities of trust can at most be mitigated, but not eliminated.

Acknowledgements

I am grateful to Jelle de Boer, J. Adam Carter, John Protevi, Karen Jones, Jason D'Cruz, Moti Gorin, David Lumsden, Philip Nickel, Nicole Huijts, Sabine Roeser, and Philip Pettit for written feedback on this work. I am also grateful to the members of the philosophy departments at the University of Edinburgh, University of Waikato, Australian National University, and the Eindhoven University of Technology for the opportunity to present a work-in-progress at their respective colloquia. Finally, I am grateful to the organizers of the 1st annual "Connecting Virtues" conference in Genoa, Italy, where I had the opportunity to present some of this research.

References

Alfano M. (2015) How One Becomes What One Is Called: On the Relation between Traits and Trait-Terms in Nietzsche. *Journal of Nietzsche Studies*, vol. 46, no 1, pp. 261–269.

6. Though see Nave, Camerer, and McCullough (2015) for a critical review of this research program.

- Alfano M. (2016) Friendship and the Structure of Trust. *From Personality to Virtue: Essays in the Psychology and Ethics of Character* (eds. J. Webber, A. Masala), Oxford: Oxford University Press, pp. 186–206.
- Alfano M., Robinson B. (forthcoming) Gossip as a Burdened Virtue. *Ethical Theory and Moral Practice*.
- Alfano M., Skorburg A. (forthcoming) Extended Knowledge, the Recognition Heuristic, and Epistemic Justice. *Extended Knowledge* (eds. D. Pritchard, J. Kallestrup, O. Palermos, J. A. Carter), Oxford: Oxford University Press.
- Alfano M., Skorburg A. (forthcoming) The Embedded and Extended Character Hypotheses. *Handbook of Philosophy of the Social Mind* (ed. J. Kiverstein), London: Routledge.
- Arrow, H. (2010). Cliques, Coalitions, Comrades, and Colleagues: Sources of Cohesion in Groups. *Social Brain, Distributed Mind* (eds. R. Dunbar, C. Gamble, J. Gowlett), Oxford: Oxford University Press, pp. 269–281.
- Aumann R., Brandenburger A. (1995) Epistemic Conditions for Nash Equilibrium. *Econometrica*, vol. 63, no 5, pp. 1161–1180.
- Baier A. (1986) Trust and Antitrust. *Ethics*, vol. 96, no 2, pp. 231–260.
- Baltag A., Moss L., Solecki S. (1999) The Logic of Public Announcements, Common Knowledge, and Private Suspicions. *TARK'98: Proceedings of the 7th Conference on Theoretical Aspects of Rationality and Knowledge*, San Francisco: Morgan Kaufmann, pp. 45–56.
- Backer L. (1996) Trust as Noncognitive Security about Motives. *Ethics*, vol. 107, no 1, pp. 43–61.
- Besser-Jones L. (2014) *Eudaimonic Ethics: The Philosophy and Psychology of Living Well*, New York: Routledge.
- Blustein J. (2008) *The Moral Demands of Memory*, Cambridge: Cambridge University Press.
- Buys C., Larsen K. (1979) Human Sympathy Groups. *Psychological Report*, vol. 45, no 2, pp. 547–553.
- Callaway E. (2011) Report Finds Massive Fraud at Dutch Universities. *Nature*, no 479, p. 15.
- Chwe M. (2001) *Rational Ritual: Culture, Coordination, and Common Knowledge*, Princeton: Princeton University Press.
- Collins H., Evans R. (2007) *Rethinking Expertise*, Chicago: University of Chicago Press.
- Dasgupta P. (1988) Trust as a Commodity. *Trust: Making and Breaking Cooperative Relations* (ed. D. Gambetta), London: Blackwell, pp. 49–72.
- De Dreu C., Balliet D., Halevy N. (2014) Parochial Cooperation in Humans: Forms and Functions of Self-Sacrifice in Intergroup Conflict. *Advances in Motivation Science*, Vol. 1 (ed. A. J. Elliot), Waltham: Academic Press, pp. 1–47.
- DeVita-Raeburn E. (2004) *The Empty Room: Surviving the Loss of a Brother or Sister at Any Age*, New York: Scribner.

- Dezecach G., Dunbar R. (2012) Sharing the Joke: The Size of Natural Language Groups. *Evolution & Human Behavior*, vol. 33, no 6, pp. 775–779.
- Dunbar R. (2012) On the Evolutionary Function of Song and Dance. *Music, Language and Human Evolution* (ed. N. Bannan), Oxford: Oxford University Press, pp. 201–214.
- Dunbar P. (1991) Functional Significance of Social Grooming in Primates. *Folia Primatologica*, vol. 57, no 3, pp. 121–131.
- Dunbar R. (1992) Neocortex Size as a Constraint on Group Size in Primates. *Journal of Human Evolution*, vol. 22, no 6, pp. 469–493.
- Dunbar R. (1993) Coevolution of Neocortical Size, Group Size and Language in Humans. *Behavioral and Brain Sciences*, vol. 16, no 4, pp. 681–735.
- Dunbar R. (2005) Gossip in Evolutionary Perspective. *Review of General Psychology*, vol. 8, no 2, pp. 100–110.
- Elder G., Clipp E. (1988) Wartime Losses and Social Bonding: Influences across 40 Years in Men's Lives. *Psychiatry*, vol. 51, no 2, pp. 177–198.
- Faulkner P. (2011) *Knowledge on Trust*, Oxford: Oxford University Press.
- Fire M., Katz G., Elovici Y. (2012) Strangers Intrusion Detection: Detecting Spammers and Fake Profiles in Social Networks Based on Topology Anomalies. *Human Journal*, vol. 1, no 1, pp. 26–39.
- Freeman L. (1978) Centrality in Social Networks Conceptual Clarification. *Social Networks*, vol. 1, no 3, pp. 215–239.
- Fricker M. (2007) *Epistemic Injustice: Power and the Ethics of Knowing*, Oxford: Oxford University Press.
- Gerbrandy J., Groeneveld W. (1997) Reasoning about Information Change. *Journal of Logic, Language and Information*, vol. 6, no 2, pp. 147–169.
- Gimple S., Nicotero G. (2013) *30 Days Without an Accident: The Walking Dead Television Series Episode*, Senoia: AMC Studios.
- Goldberg S. (2010) *Relying on Others: An Essay in Epistemology*, Oxford: Oxford University Press.
- Goldberg S. (2013) Self-Trust and Extended Trust. *Res Philosophica*, vol. 90, no 2, pp. 277–292.
- Govier T. (1993). Self-Trust, Autonomy, and Self-Esteem. *Hypatia*, vol. 8, no 1, pp. 99–120.
- Govier T. (1997) *Social Trust and Human Communities*, Montreal: McGill-Queen's University Press.
- Hardin R. (2002) *Trust and Trustworthiness*, New York: Russell Sage Foundation.
- Hardwig J. (1991) The Role of Trust in Knowledge. *Journal of Philosophy*, vol. 88, no 12, pp. 693–708.
- Haas Z., Halpern J., Li L. (2006) Gossip-Based Ad Hoc Routing. *IEEE/ACM Transactions on Networking (TON)*, vol. 14, no 3, pp. 479–491.
- Hieronymi P. (2008) The Reasons of Trust. *Australasian Journal of Philosophy*, vol. 86, no 2, pp. 213–236.
- Hobbes T. (1994) *Leviathan* (ed. E. Curley), Indianapolis: Hackett.

- Jones K. (1999) Second-Hand Moral Knowledge. *Journal of Philosophy*, vol. 96, no 2, pp. 55–78.
- Jones K. (2012a) Trustworthiness. *Ethics*, vol. 123, no 1, pp. 61–85.
- Jones K. (2012b) The Politics of Intellectual Self-Trust. *Social Epistemology*, vol. 26, no 2, pp. 237–251.
- Katz N., Lazer D., Arrow H., Contractor N. (2004) Network Theory and Small Groups. *Small Group Research*, vol. 35, no 3, pp. 307–332.
- Kramer S., Goré R., Okamoto E. (2010) Formal Definitions and Complexity Results for Trust Relations and Trust Domains Fit for TTPs, the Web of Trust, PKIs, and ID-Based Cryptography. *ACM SIGACT News Archive*, vol. 41, no 1, pp. 75–98.
- Krebs V. (2002) Mapping Networks of Terrorist Cells. *Connections*, vol. 24, no 3, pp. 43–52.
- Krishnamurthy M. (2015) (White) Tyranny and the Democratic Value of Distrust. *The Monist*, vol. 98, no 4, pp. 391–406.
- Kruger J., Dunning D. (1999) Unskilled and Unaware of It: How Difficulties in Recognizing One's Own Incompetence Lead to Inflated Self-Assessments. *Journal of Personality and Social Psychology*, vol. 77, no 6, pp. 1121–1134.
- Lehrer K. (1997) *Self-Trust: A Study of Reason, Knowledge, and Autonomy*, Oxford: Oxford University Press.
- McGeer V. (2008) Trust, Hope, and Empowerment. *Australasian Journal of Philosophy*, vol. 86, no 2, pp. 237–254.
- McGeer V., Pettit P. (forthcoming) The Empowering Theory of Trust. *The Philosophy of Trust* (eds. P. Faulkner, T. Simpson), Oxford: Oxford University Press.
- McLeod C. (2015) Trust. *The Stanford Encyclopedia of Philosophy*. Available at: <http://plato.stanford.edu/archives/fall2015/entries/trust/> (accessed 1 September 2016).
- Lewis D. (2002) *Convention*, London: Blackwell.
- Miller S. (2010) *The Moral Foundations of Social Institutions: A Philosophical Study*, Cambridge: Cambridge University Press.
- Nash J. (1951) Non-cooperative Games. *Annals of Mathematics*, vol. 54, no 2, pp. 286–295.
- Nave G., Camerer C., McCullough M. (2015) Does Oxytocin Increase Trust in Humans? A Critical Review of the Research. *Perspectives on Psychological Science*, vol. 10, no 6, pp. 772–789.
- Nickel P. (2012) Trust and Testimony. *Pacific Philosophical Quarterly*, vol. 93, no 3, pp. 301–316.
- Nietzsche F. (1997) *Untimely Meditations* (ed. D. Breazeale), Cambridge: Cambridge University Press.
- Pacuit E., Parikh R. (2004) The Logic of Communication Graphs. *Declarative Agent Languages and Technologies II* (eds. J. Leite, A. Omicini, P. Torroni, P. Yolum), Berlin: Springer, pp. 256–269.
- Pasnau R. (2015) Disagreement and the Value of Self-Trust. *Philosophical Studies*, vol. 172, no 9, pp. 2315–2339.
- Pettit, P. (1995) The Cunning of Trust. *Philosophy and Public Affairs*, vol. 24, no 3, pp. 202–225.

- Pettit P. (2008) Freedom and Probability: A Comment on Goodin and Jackson. *Philosophy & Public Affairs*, vol. 36, no 2, pp. 206–220.
- Pettit P. (2015) *The Robust Demands of the Good: Ethics with Attachment, Virtue, and Respect*, Oxford: Oxford University Press.
- Potter N. (2002) *How Can I be Trusted? A Virtue Theory of Trustworthiness*, Lanham: Rowman & Littlefield.
- Powell C., Nicotero G. (2015) *Remember: The Walking Dead Television Series Episode*, Senoia: AMC Studios.
- Preston-Roedder R. (2013) Faith in Humanity. *Philosophy and Phenomenological Research*, vol. 87, no 3, pp. 664–687.
- Reid E., Qin J., Zhou Y., Lai G., Sageman M., Weimann G., Chen H. (2005) Collecting and Analyzing the Presence of Terrorists on the Web: A Case Study of Jihad Websites. *Intelligence and Security Informatics* (eds. P. Kantor, G. Muresan, F. Roberts, D. D. Zeng, F. Y. Wang, H. C. Chen, R. C. Merkle), Berlin: Springer, pp. 402–411.
- Richman B. (2006) How Community Institutions Create Economic Advantage: Jewish Diamond Merchants in New York. *Law & Social Order*, vol. 31, no 2, pp. 386–420.
- Russell D. (2014) Aristotle on Cultivating Virtue. *Cultivating Virtue Perspectives from Philosophy, Theology, and Psychology* (ed. N. E. Snow), Oxford: Oxford University Press, pp. 17–48.
- Russell D. (2015) From Personality to Character to Virtue. *Current Controversies in Virtue Theory* (ed. M. Alfano), London: Routledge, pp. 92–106.
- Sable M., Danis F., Mauzy D., Gallagher S. (2006) Barriers to Reporting Sexual Assault for Women and Men: Perspectives of College Students. *Journal of American College Health*, vol. 55, no 3, pp. 157–162.
- Savage D., Zhing X., Yu X., Chou P., Wang Q. (2014) Anomaly Detection in Online Social Networks. *Social Networks*, vol. 39, no 4, pp. 62–70.
- Shay J. (1994) *Achilles in Vietnam: Combat Trauma and the Undoing of Character*, New York: Simon & Schuster.
- Shay J. (2003) *Odysseus in America: Combat Trauma and the Trials of Homecoming*, New York: Scribner.
- Slingerland E. (2014) *Trying Not to Try: The Art and Science of Spontaneity*, New York: Crown.
- Smith C. P., Freyd J. (2014) Institutional Betrayal. *American Psychologist*, vol. 69, no 6, pp. 575–587.
- Sterelny K. (2012) *The Evolved Apprentice: How Evolution Made Humans Unique*, Cambridge: MIT Press.
- Strawson P. (1962) Freedom and Resentment. *Proceedings of the British Academy*, vol. 48, pp. 1–25.
- Šubelj L., Furlan U., Bajec M. (2011) An Expert System for Detecting Automobiles Insurance Fraud Using Social Network Analysis. *Expert Systems Applications*, vol. 38, no 1, pp. 1039–1052.

- Townley C., Garfield J. (2013) Public Trust. *Trust: Analytic and Applied Perspectives* (eds. P. Makela, C. Townley), Amsterdam: Rodopi Press, pp. 95–108.
- Turri J. (forthcoming) Sustaining Rules: A Model and Application. *Knowledge First* (eds. J. Carter, E. Gordon, B. Jarvis), Oxford: Oxford University Press.
- Wiessner P. (2014) Embers of Society: Firelight Talk among the Ju/'hoansi Bushmen. *Proceedings of the National Academy of Sciences*, vol. 111, no 39, pp. 14027–14035.

Топология сообществ доверия

Марк Альфано

PhD, доцент департамента ценностей, технологий и инноваций Политехнического института в Дельфте

Адрес: Jaffalaan 5, Delft, The Netherlands 2628 BX

E-mail: mark.alfano@gmail.com

Гоббс считал, что естественное состояние — это состояние войны, поскольку оно определяется фундаментальным и обобщенным недоверием. Выход из естественного состояния и конфликтов, которые оно неизбежно влечет, предполагает установление доверия. Однако обсуждение «доверия» в философии чаще всего фокусируется либо на изолированных диадах индивидов, доверяющих друг другу, либо на доверии в больших безликих институтах. Задача данной статьи — заполнить пробел между этими двумя крайностями путем анализа «топологии сообществ доверия». Такие сообщества можно лучше всего представить как взаимосвязанные диадические отношения, которые приближены к идеалу поскольку они симметричны, рефлексивны и транзитивны. Немногие сообщества доверия соответствуют этому описанию, и те, что приближаются к нему, это, как правило, небольшие сообщества (от 3 до 15 индивидов). Именно в таких сообществах доверия возможно появление различных важных рациональных, познавательных, культурных благ и душевного спокойствия. Однако в сообществах доверия становятся возможными также многие сомнительные феномены. Сообщества могут изолироваться от окружающего общества, и тем самым стать источником недоверия со стороны окружения. Они могут вынудить своих членов отказаться от общественных благ в пользу родовых или частных благ. Этими недостатками сообщества доверия обязаны тем же самым механизмам, которые придают им положительную, рациональную, эпистемическую, культурную и психическую ценность, поэтому последствия работы этих механизмов можно только смягчить, но не исключить целиком.

Ключевые слова: доверие, недоверие, сообщество, процветание, эпистемология, топология, эмоции, социальная эпистемология