Alternative conceptions of generalized trust (and the foundations of the social order)

Marc A. Cohen, Seattle University

Abstract. Generalized trust is widely said to be essential for social and economic cooperation, but—despite the large empirical literature—there is disagreement and confusion over how to understand generalized trust. This paper develops the conceptual options that can be drawn from the social science literature—with attention to the moral dimension in each, and with some attention to the different ways that generalized trust can serve as a foundation for the social order.

Introduction

Generalized trust is widely said to be essential for social and economic cooperation and for " 'good' social, political, and economic outcomes in society."[1] But, despite extensive empirical work (described very briefly below), there are competing conceptions of generalized trust and there is disagreement about what it means to trust. This disagreement is often noted at the beginning of published papers—in passing—and after which authors quickly move on to their data. As Peter Nannested notes in his review: "Much of the recent empirical work on [generalized] trust... does not seem to proceed from any clear account of what is meant by trust in the first place. Rather, [generalized] trust is taken to be what is measured by one or more survey questions or by subjects' observable behavior in certain experimental games. Thus, the results do not normally tell much—if anything—about the merits of one theoretical concept of trust versus the merits of another."[2]

The present paper addresses this conceptual question about how to understand generalized trust (or some use the terminology 'social trust'), which is in part a question of how to distinguish generalized trust from particularized trust.

Generalized trust is most commonly defined—and at the same time distinguished from particularized trust—with reference to the trusted parties: we have generalized trust in

out-group members (including people met for the first time, those of another religion, those of another nationality) versus particularized trust of in-group members (including family, neighbors, and those known personally). For example, as Markus Freitag and Richard Traunmüller put it, generalized trust is "a rather abstract attitude toward people in general, encompassing those beyond immediate familiarity, including strangers (people one randomly meets in the street, fellow citizens, foreigners, etc.)," while particularized trust is an attitude toward family members, friends, close neighbors, and close co-workers.[3] So, they describe the two forms of trust as operating in different "spheres." The present paper breaks from this approach. When, for example, I lend money to my sister and trust her to return it, when I trust my neighbor to take care of my children for an afternoon, when I trust a colleague to help me with some task—these are all cases of particularized trust of the form A trusts B to do x, with different specific commitments filling out the content of 'to do x.' When I trust an unknown taxi driver to take me to my destination and not, say, drive me to a dangerous part of the city where I can be robbed, or when I trust the (unknown) owners of a restaurant to insure that food is safe to consume, Freitag and Traunmüller would categorize these two examples as generalized trust in strangers. But these latter two examples have the same structure as the ones above, A trusts B to do (a particular) x—so the present paper considers these to be particularized trust in out-group members. And as a result, Frietag and Traunmüller's reference to "spheres" fails to capture the conceptual space that generalized trust could fill as distinct from particularized trust, and they narrow consideration of the roles played by generalized trust in sustaining the social order.[4]

So, where Freitag and Traunmüller (and others) distinguish generalized trust from particularized trust in terms of the "sphere," the present paper begins with the assumption that—instead—generalized trust and particularized trust differ in more fundamental ways. The material that follows maps and develops the alternative conceptions that can be drawn from the social science literature—discussing each in turn—with particular attention to the moral dimension in each. On the first, generalized trust is a willingness to engage in a particularized trust relationship (with particular content) with members of the out-group— meaning that generalized trust is itself not a trust relationship. This conception is implicit in existing empirical papers, and the material below makes the conception explicit and more precise. On the second, generalized trust is said to be a (general) expectation or a belief that others will not knowingly do harm or exploit our vulnerabilities—'general' in the sense that it does not make reference to a specific task and applies to all persons. Generalized trust is

2

often described in these terms, but I argue that we should reject this second conception. On the third, generalized trust relationships have the following structure: A trusts B means that A relies on B to act in accordance with some (specified) general or background moral obligation, where A can assume that B is committed to acting in that way because of the character of the obligation. The prototypical example might be an obligation to not willingly harm others. On this account B could be an in-group or an out-group member, and generalized trust is distinguished from particularized trust in terms of the content of the trust relationships. This third conception represents my own attempt at repairing problems with the second, extending my own previous work. And fourth, with sociologist Harold Garfinkel, we might talk of generalized trust in terms of the fundamental constitutive practices that make a social order possible (assuming we could identify candidate fundamental constitutive practices, a question discussed below).

As mentioned, these four conceptions represent the options in the literature (though, of course, others could be proposed). Setting aside the second, the remaining three are not in conflict—we might say that they describe different forms of generalized trust and that, on each of these conceptions, generalized trust plays a different role in supporting the social order (though these points about the social order are only partly developed given space limitations).

The present paper is important for three reasons. First, this paper provides a foundation for further empirical work, which could directly measure the different forms of generalized trust distinguished/developed here. Moreover, there are widespread claims (across academic disciplines) that generalized trust is essential for sustaining the social order, as noted above, but the empirical evidence for those claims is mixed and inconclusive[5]—and we might make progress with a clearer understanding of (and with clearer measures of) generalized trust in each form. For example, generalized trust as a propensity for particularized trust (the first form considered below) might be important for building cooperative organizations, even if generalized trust in the other forms does not play an important role in that context. And there is another, related empirical project described below: the third conceptualization of generalized trust makes reference to background moral obligations; these obligations will differ across cultures and contexts, the content of these obligations is an empirical matter. Second, this paper invites further philosophical analysis directed at the role of generalized trust—in its different forms—in the social order. Third, this paper provides grounds for a substantive claim about the foundations of the social order:

3

the second form of generalized trust should be rejected, and the remaining forms are moral in a substantive sense; so, to claim that the social order depends on generalized trust (in one or all of its forms) is to suggest that the social order depends on moral relationships between the persons involved. This line of thought is not developed here, and this paper doesn't address the further implications, but the analysis supports this claim about the normative foundation for the social order.

Two further points before proceeding. The distinction between generalized trust and particularized trust is not made in the philosophical literature, so the present paper is grounded in the (inter-disciplinary) social science literature.[6] That empirical literature widely relies on data gathered using this survey question: "Generally speaking, would you say that most people can be trusted or that you can't be too careful in dealing with people?" In practice, in that literature, generalized trust is what this survey question measures, so the empirical project can proceed without a stated way of understanding generalized trust. (This is the first point.) The question has been asked across a number of surveys (including the World Values Survey) and there is longitudinal data (at the national level) going back to 1948. The ambiguities in the question are widely noted. If we think of a trust relationship in terms of two persons, A trusts B to do x, the question's reference to "most people" is ambiguous as to who counts as B. Cross-country comparisons of generalized trust assume that there is no systematic difference in who counts as "most people" across countries, but Nannestad's findings casts doubt on this assumption: he reports unpublished data showing that, for example, 47.1% of Danish immigrants reported trusting most people, but of that group 26.2% indicated not trusting Danish Jews, and there were similar results for other social and ethic groups.[7] Moreover, the content of the trust relationship—x in the schematic formulation—is not specified at all, and this introduces further ambiguity. Comparisons of the level of generalized trust across populations must assume that persons respond to the survey with (approximately) the same content of trust in order to make those comparisons meaningful. And, these surveys assume that respondents conceptualize trust in the same way, so researchers do not have to specify and commit to one or another conception—so, again, the empirical project can proceed without a stated way of understanding generalized trust.

Despite these ambiguities, results are consistent across multiple surveys at the aggregate (national) level (that is, there is test-retest stability). This indicates that, at least within countries, respondents bring the same assumptions about the meaning of "most people" and also about the content of trust. And, survey respondents do not seem to find the

question difficult or ambiguous (very few respondents skip the survey question or report not understanding it). Moreover, aggregate levels of generalized trust correlate with other measures of trustworthiness—so, for example, aggregate levels of generalized trust are higher in countries where there is less corruption and where there is less violent crime. And aggregate levels of generalized trust correlate with experimental results: for example, generalized trust is higher in countries where persons are more likely to return a lost wallet intact.[8] So the survey question apparently measures something, and empirical work concerns the individual, social, and psychological factors that support and/or damage generalized trust. Of the conceptions considered here, the first three are plausible interpretations of the survey question. But we would need further empirical work—most likely in the form of additional survey questions—in order to say that the data measures a particular one of the forms of generalized trust described here. (This is the second point.)

1. Generalized Trust as a Propensity for Particularized Trust

Most directly, first, we might think that generalized trust does not have specific content; it is instead a willingness to engage in particularized trust relationships—of the form A trusts B to do x—with members of the out-group. So, when generalized trust is high persons are more likely (more willing) to trust, and vice versa. The general survey question could, perhaps, be most plausibly interpreted in these terms—as long as "most people" in the general survey question is taken to refer to the out-group (more on this in a moment). There is still a further question about how to best understand particularized trust, and in particular about whether particularized trust is a moral relationship, but we can set that question aside for the moment. This is most certainly a plausible conception of generalized trust. And we could make sense of broad, intuitive claims about the role of generalized trust in these terms: we might think social order, the possibility of cooperation, economic activity, etc., all depend on the readiness with which persons can trust one another and on that basis interact.

Though not stated explicitly, Jan Delhey, Ken Newton, and Christian Welzel's recent paper is best understood as adopting this conception of generalized trust. The fifth wave of the World Values Survey included a newly developed set of questions, ones designed to distinguish between a willingness to trust in-group members (family, neighbors, and those known personally) and out-group persons (people met for the first time, those of another religion, those of another nationality). Using this additional data, Delhey et al. were able to calculate a "radius of trust," meaning that aggregate (national)-level measures of generalized

trust could be compared on two dimensions—the strength of that trust and also the radius or scope of who counts as "most people." They found that for forty-one countries, responses to the older survey question about trust in "most people" tracked responses to the new question about trusting those in the out-group—so in those forty-one countries "most people" was taken to refer to out-group members. But in eight countries this was not the case: China and South Korea, for example, have a very high level of generalized trust according to the original survey question but, according to Delhey et al.'s analysis, the trust radius in both countries is very narrow. This confirmed Francis Fukayama's suggestion that Confucian countries have great difficulty trusting outside of the family (though, according to Delhey et al., Taiwan is a notable and unexplained exception). So in short, Delhey et al. were able to address ambiguities in the survey question and measure levels of trust in out-group members. In doing so, we can take them to assume that generalized trust is a willingness to engage in particularized trust relationships with out-group members.[9]

2. Generalized Trust as Generalized Expectations

Generalized trust is often conceptualized in terms of generalized expectations or beliefs about the trusted party. For example, according to Jan Delhey and Ken Newton generalized trust is "the belief that others will not deliberately or knowingly do us harm, if they can avoid it, and will look after our interests, if this is possible."[10] This is different from Freitag and Traunmüller's account, mentioned in the introduction, in that the content of the trust relationship here—'x' in A trusts B to do x—is general, as opposed to making reference to some specific task or action. This is an account of generalized trust we should reject.

Delhey and Newton present this account as "close to" the definition proposed by Italian sociologist Diego Gambetta.[11] Gambetta defined trust in terms of expectations expressed as subjective probabilities: "trust (or symmetrically, distrust) is a particular level of the subjective probability with which an agent assesses that another agent or group of agents will perform a particular action, both *before* he can monitor such action... *and* in a context in which it affects *his own* action."[12] Here the expectations could be grounded on a number of factors, including evidence (inference from past behavior), personal identification with the trusted party (i.e., I trust my sister), emotional factors, or a "leap" beyond rational considerations (with Guido Moellering).[13]

Gambetta's account serves as a starting point for much of the social science literature on trust, particularly in sociology and management—including Roger C. Mayer, James H.

Davis, and F. David Schoorman's very widely-cited definition of trust.[14] Gambetta's account concerned particularized trust, and underlying his account is a certain way of thinking about the role particularized trust plays in society, a way of thinking about the problem trust can solve: individuals must interact and coordinate their actions in order to manage/ secure their own interests; interaction of this sort depends in a fundamental way on our having clear expectations about the actions of others; we trust when we have positive expectations about the behavior of others (meaning, expectations with high probability, not normatively positive expectations about the other party's character); and then on that basis we can determine our own actions. Generalized trust conceptualized in these terms would serve the same role: expectations that others will not deliberately or knowingly do us harm would support social and economic interaction.

But, as an account of particularized trust, the problem is that Gambetta cannot explain how *A trusts B to do x* is different from *A expects B to do x*. In a widely-cited paper, Annette Baier argues that trust is distinguished by the possibility of betrayal,[15] and we can use Baier's requirement to think about what is missing from Gambetta's account of trust as expectations: with Gambetta, when B fails to do x, we could say that B's actions surprised A, that B's actions disappointed A's expectations and also frustrated A's plans. But—because B has no obligation to act as A expects (even when the expectation is epistemically warranted) or, put another way, because A's expectations do not give A any sort of right with respect to B's behavior—on Gambetta's account there is no conceptual space for identifying a trust violation as involving betrayal. Gambetta's account is therefore incomplete. To be sure, when A trusts B to do x, A will hold expectations about B's behavior; the point here is that trust is more than merely holding an expectation—and Gambetta's definition has not captured the additional factor(s) involved.

The same line of argument applies to accounts of generalized trust as general expectations about persons: the fact that A expects B do to x does not itself explain the possibility of betrayal or give substance to claims about generalized trust. So we should reject conceptualizations of generalized trust as general expectations about other people. (Where Gambetta refers to subjective probabilities, others refer to expectations, and Delhey and Newton formulate their account in terms of beliefs—but there is no essential difference between those respective terminologies.) *But to emphasize, the critical argument is narrow: the point is that expectations cannot do the work required.* We can, however, read a different sort of account—and a plausible one—into Delhey and Newton's suggestion, an account

focused on the content of the expectation. This alternative is developed in the next section.

3. Generalized Trust as Reliance on Generalized Moral Obligations

The conception of generalized trust developed in this section extends my own account of particularized trust, which is very briefly described here as preparation, and which focuses on the commitment binding the parties: A trusts B to do x when A relies on B's commitment to do x.[16] The commitment here creates an obligation on B's part, and if B fails to do x then A has a legitimate grievance. The obligation is not derived from general or background moral obligations; for example, if A trusts B to pick up A's mail during a trip abroad (do x), that trust relationship is given structure by a commitment on B's part to do x— that commitment creates the obligation—one we cannot derive from more general moral obligations on B's part to be honest, act with goodwill, or the like. Nevertheless, there is a sense in which B owes that action to A in the trust relationship, giving A standing to rebuke B if B fails (to use Margaret Gilbert's language).[17] So there is a moral dimension to trust on this conception, even though (as just mentioned) the obligations are not derived from or explained in terms of general moral ones. And, on this conception, to say that social interaction depends on particularized trust is to say that social interaction relies on a network of commitments and obligations binding the persons involved.

The commitment involved here could be implicit or explicit. When explicit we can think of the commitment in terms of a promise, we trust when we rely on another's promise, A (above) trusts B to pick up his mail when he relies on B's promise to do so.[18] When implicit the commitment could be embedded in roles (for example, doctors make a role-based commitment to exercise due care for patients), or the commitment could be embedded in identity (B might have particular obligations because A and B are both members of the same organization, live in the same neighborhood, are part of the same ethnic group, etc., and on that basis hold certain implicit obligations to each other), or the commitment could have some other source—no single factor will account for the presence of implicit commitments across contexts and situations. Moreover, the commitments and the particular acts of trust could be separated by a great deal of time (for example, the doctor makes an on-going commitment without knowing, in advance, who the patients will be). And the commitments involved could evolve over time; this sort of evolution accounts for the way that trust relationships can get deeper with repeated interaction, involving more complex, demanding obligations on B's part and greater vulnerability on A's part: A might trust B to pick up his

8

mail, and then later—as the relationship deepens, perhaps with many intermediate stages—then trust B to pay A's bills during a trip abroad, giving B access to A's bank account. To be sure, there could be *much* ambiguity in practice as to whether B holds the obligation(s) in question. Trust is misplaced (in one sense) when there is confusion about whether there are obligations binding A and B. (In another sense, perhaps the more common sense, trust will be misplaced when we trust those we shouldn't, when our reasons for trusting are inadequate, mistaken, etc. But, and to emphasize, the reasons for trusting are not at issue here: on the account developed in this section, trust could be rational/calculative, or emotional, or a leap beyond rational considerations, or some mix of all of these—the point here is about the trust relationship itself and the way obligations provide the structure, not the reasons for trusting.)[19]

This commitment conception of trust makes sense of our ordinary practice of trusting, and in particular it makes sense of the complaint "but I trusted you!" after a trust violation. So this commitment conception satisfies Baier's requirement. Other definitions in the social science literature do not offer an alternative that can meet this requirement. And much of the philosophical literature is problematic on exactly this point. That literature widely treats trust as a first person attitude or stance, as a readiness to feel betrayed; for example, Philip J. Nickel defines trust as follows: "if one person [A] trusts another [B] to do something, then she [A] takes him [B] to be obligated to do that thing." But Nickel's point concerns the "conceptual requirements of the attitude of trust," and he is careful to say that no actual obligation exists, A acts *as if* B is obligated. So Nickel can explain why A might *feel* betrayed if B fails to act, but on Nickel's view A's feeling of betrayal cannot be justified: B is not obligated just because A acts *as if* B is obligated. As a result Nickel's account cannot meet Baier's requirement. We must move on to think about the actual obligations that give structure to trust relationships (beyond trust as an attitude).[20]

There is much more to say to defend this account of particularized trust. But the goal here is to systemize the conceptions of generalized trust—and we can make progress without settling those details. So, first, we could go back and apply this account of particularized trust to the first conception of generalized trust above: we can take generalized trust to be a willingness to rely on another's particular commitment—where this other is an out-group member—within a certain scope, that is, within limits on the level of vulnerability involved (a separate ambiguity we would need to address but cannot do so here). Generalized trust so-conceived collapses back into particularized trust and, as mentioned, generalized trust

wouldn't itself have content though it does express a trusting attitude of a sort.

Alternatively: Above (in section two) I argued that we should reject the expectations-conception of particularized trust, because holding expectations about another party cannot account for the possibility of betrayal. In forming a particularized trust relationship, in contrast, A will rely on B's commitment, and we can talk of betrayal or wrong-doing with reference to the obligations put in place by that commitment. We could extend this account to the case of general or background moral obligations—where these background moral obligations could include Delhey and Newton's obligation to not deliberately or knowingly do harm to others, or an obligation to respect others' property and not steal, or a general obligation to be honest. When A relies on B to fulfill a background moral obligation this is generalized trust. But those background moral obligations do not have to be put in place by commitments among the parties involved, they can be presupposed or assumed. So, for example, person A could rely on person B's commitment to be honest because that background moral obligation is in place. That said, we should be careful to distinguish a commitment in this sense, which represents the terms of the social order, from commitment in the sense of really-meaning-to-do-it: B might be committed to being honest in the first sense even though B is a persistent liar.

There are three important things to see here. First, on this conception, generalized trust and particularized trust have the same conceptual structure, in that A (the trustor) relies on a commitment by B (the trustee), and when the trustee fails to act on this commitment we can account for the wrongness in a trust violation.

Second: the difference, though, as just mentioned, is that the commitments in generalized trust can be presupposed or assumed as part of our background moral obligations, meaning that generalized trust *recognizes* and relies on the background moral obligations binding persons. But for particularized trust, commitments must put obligations in place, and these commitments change the relationship between A and B. This way of putting the point might seem too strong; when A trusts B to be honest in some negotiation, where this is an instance of generalized trust—making reference to a background moral obligation to be honest or to act with integrity—we might want to say that the relationship between A and B is changed because A now relies on B's obligation. That point makes sense; the relationship is changed because, now that A relies on B, the particular obligation is relevant—and the obligation in effect—in a way that it wasn't before, even though the generalized trust relationship did not constitute that obligation.

Third, the account here is structural, meaning that it concerns the conceptual structure of generalized trust in terms of background moral obligations and presumed commitments to honor those obligations—but without making a claim about what those background moral obligations are. Some such obligations, like Delhey and Newton's obligation to not deliberately or knowingly do harm to others, seem to be obvious, and we might expect them to be universal, though even that obligation has boundaries and limitations, places/persons to whom it does and doesn't apply. But the content of generalized trust will vary across societies, with different fundamental commitments binding citizens. As an example, consider the following from Fukuyama. Employment in large firms in Japan is essentially lifetime employment, and Fukuyama notes that this sort of arrangement would seem to invite free riding: employees are compensated on the basis of seniority rather than job performance, so "any increased benefits arising from superior performance are in effect a public good with respect to the company as a whole, giving an individual an incentive to shirk his part of the burden."[21] But there is an especially strong work ethic among those employees, and Fukuyama explains this in terms of a reciprocal moral obligation: firms provide stable employment and steady advancement, while workers provide their best efforts. This sense of reciprocal obligation is deeply embedded in Japanese society as a result of its feudal traditions, and mutual acceptance of this norm is implicit in the offering of and the acceptance of employment. Employers rely on employees' commitment when offering and continuing employment, employees rely on the organizations' commitment in accepting work. The arrangement therefore depends on generalized trust.

So, where Freitag and Traunmüller addressed the boundaries of generalized trust in terms of the trusted party (in-group versus out-group), there is a different *empirical* research question about the content of generalized trust in different societies.[22] To claim that good social outcomes depend on generalized trust (in the sense here) is to say that good social outcomes depend on an underlying set of commitments and obligations binding persons— generalized trust is a matter of recognizing and relying on those commitments, which give shape to the relationships between persons—and those commitments and obligations must be specified. The social commitments and obligations in Japanese work arrangements are not present in American ones.

4. Garfinkel's Conception of Generalized Trust

Sociologist Harold Garfinkel's work is best understood as offering a related

conception of generalized trust, even though he does not use that term or distinguish his conception of trust from instances of particularized trust. This section outlines his account and draws out the difference between it and the preceding ones, with particular attention to the normative dimension.[23]

In the context of Garfinkel's work, trust is not about expectations with respect to some particular action; trust is not related to general moral obligation(s); and trust is not associated with integrity, or goodwill, or anything like that. Instead, for Garfinkel, A trusts B when A expects that B will act in "compliance" with the norms or basic rules (where those are synonyms) of the social situation.[24] Relying on these basic rules is a "condition for 'grasping' the events of games," and by extension a condition of making sense of the events in daily life.[25] The expectations here are "constitutive" in that they provide the background against which B's actions (B's moves in a game) have meaning—and in these terms, A trusts B when A holds constitutive expectancies about B's behavior. This is a conception of generalized or basic trust because it applies to all persons and because it concerns expectations that constitute the possibility of action, rather than expectations connected with particular joint activities.

Garfinkel uses the game of ticktacktoe as an example: given the rules of the game— "Play is conducted on a three by three matrix by two players who move alternatively. The first player makes a mark in one of the unoccupied cells...."—or, against the background of these rules, we can interpret/understand what it means to put an X in a cell, and we can understand what is wrong with marking an X on one of the lines instead. The "constitutive expectancies" here are the assumption that the other player is acting in compliance with these rules. Garfinkel was able to expose the role played by these expectancies in an experiment: student-experimenters played ticktacktoe games with subjects; the subjects moved first, then the experimenter erased the subject's mark and moved it to another cell. The subjects experienced confusion because the student-experimenters violated their expectancies; this makes the basic rules and the constitutive order visible. Garfinkel argues that the situations in everyday life are also governed by constitutive expectancies, even if the rules are not specified as precisely as those governing games like ticktacktoe, and he is able to show this in other quasi-experimental demonstrations (for example, subjects entered stores and treated another customer as if he or she were an employee, demanding service, documenting the confused, and sometimes hostile, reactions).

Social interaction clearly depends on constitutive expectancies: individual action in

specific contexts is given structure by the basic rules; acting according to those basic rules renews and reinforces the rules; and persons interpret the actions of others against the background of those rules—as a result, according to Garfinkel, these actions, observed phenomena, are second-order in the sense that they cannot be explained and understood apart from or independent of the underlying constitutive rules.[26] Garfinkel's methodological point is that sociological research ought to focus on those constitutive rules, this is his ethnomethodology research program.

The further question here is whether we ought to talk of trust at all in this context. The argument in section two suggested that merely holding expectations is not to trust and, further, we should only talk of trust when there is obligation/ a commitment in place. Could we argue that persons have an obligation to act in compliance with Garfinkel's basic rules?— if so then talk of trust would be appropriate. (And if so, then Garfinkel's account would be a variant of the one developed in section three making reference to generalized moral obligations.)

At points in his account, Garfinkel makes explicit reference to the normative dimension of the social order. For example, in the (1963) essay he defines trust as follows, "To say that one person 'trusts' another means that the person seeks to act in such a fashion as to produce through his action or to respect as a condition of play actual events that accord with normative orders of events depicted in the basic rules of play."[27] But Garfinkel uses the term "normative order" to refer to the set of basic rules or norms we *expect* persons to follow—there is no obligation. In Garfinkel's ticktacktoe example, for example, there is certainly no sense in which persons have a moral obligation to play that game; Garfinkel is clear about that, describing the commitment to the constitutive practices in terms of *choice*: the rules are binding "for persons who seek to act in compliance with them (a player)," leaving space for those who choose otherwise.[28] We should note, further, that a normative order in Garfinkel's sense—the collection of basic rules—could involve decidedly and unambiguously immoral "moral orders" (think of the normative orders that defined interaction in the American antebellum south).[29] So Garfinkel's own references to a normative order do not support reading him in terms of obligations and commitments, and as a result talk of trust is not warranted. To emphasize, though, this is not a critique of the ticktacktoe experiment described above, nor is it a critique of Garfinkel's more general claim about the role of constitutive expectancies in ordinary life—the point is only that acting on constitutive expectancies is not yet to trust.

13

But sociologist Anne Warfield Rawls suggests otherwise: for Garfinkel, there is a set of *fundamental* basic rules that apply across all social settings, or there is a set of what we might call fundamental constitutive practices ("fundamental" is my term, used to make reference to a set of rules that are deeper and more general than, for example, the rules of a particular game like ticktacktoe or the norms in a particular social interaction). These practices, she suggests, make possible mutual intelligibility, joint activity, even what John Rawls calls primary goods, and perhaps even personhood *across all situations*. This is—plausibly—what Garfinkel had in mind even if much more would have to be said to develop the suggestion.[30] And if so, given that stability and coherence in the social order depend on mutual commitment to fundamental constitutive social practices, persons would have an obligation to be committed to these rules (whether or not they are committed to the rules of any particular games).[31]

This is an intriguing suggestion, and it would give substance to Garfinkel's talk of generalized *trust* (beyond talk of expectations): members of the social order are bound by commitments to fundamental constitutive practices, we could say that we trust when we rely on those commitments, and the result is a stable social order. But we must be clear: the fact that the stability of the social order depends on the mutual or reciprocal commitment to fundamental constitutive practices doesn't yet show—as Anne Warfield Rawls suggests—that persons have a general moral obligation to participate in and sustain a *reciprocal* social order, one in which all participate as equals, share in the benefits of the social order, are committed at least in some minimal way to the well-being of others, etc. That seems to be a further claim about the particular shape of the social order and a claim about the substantive obligations binding persons—beyond the (more minimal) obligation to sustain order itself.

One question here is whether we have candidates for such fundamental rules. We would need to identify candidates to defend the claim above, that is, to show that mutual intelligibility, joint activity, primary goods, and personhood are all at stake. But assuming we could identify such rules, generalized trust of this sort would be necessary for but likely not sufficient—by itself—for generating the essential social goods listed above, mutual intelligibility, joint activity, and primary goods. Moreover, with Garfinkel, we act in specific social settings on the basis of constitutive expectancies, but we should not talk of trust with respect to these expectancies—with respect to those rules of the game—because participation in those is a choice, we can expect commitment to those rules (we expect persons to play ticktacktoe according to the usual rules), but we cannot claim to be wronged

when someone choses otherwise. So the ticktacktoe example reveals the role of constitutive expectancies but not generalized trust. As a result, Garfinkel's form of generalized trust would not replace the one described in the previous section, according to which persons rely on (presumed) background moral commitments—Garfinkel's work concerns deeper practices that sustain intelligibility and joint activity rather than substantive (but general) moral obligations binding persons.

Concluding Summary

In sum, then, this paper has distinguished between four ways to think about generalized trust. On the first, generalized trust is a willingness to engage in trust relationships with particular content with an out-group member, generalized trust is itself not a trust relationship. On the second, generalized trust is a (general) expectation or belief about the behavior of others; we should reject this conception for the reasons outlined above. But the third conception offered a related alternative: with generalized trust, A trusts B to act in accordance with some (specified) general or background moral obligation, where A can assume that B is committed to acting in that way because of the character of the obligation (where the prototypical one might be an obligation to not willingly harm others, but the obligations here could be much stronger—as in the example from Fukuyama outlined above). Here B could be an in-group or an out-group member. Fourth, with Garfinkel, we might talk of generalized trust in terms of the fundamental constitutive practices that make a social order possible (assuming we could identify candidate fundamental constitutive practices). Setting the second aside, the remaining three forms of generalized trust (the first, third, and fourth) are not in competition, they each sustain the social order in different ways, and for that reason conceptual and empirical work on generalized trust should be clear about which is at issue.

[1] Peter Nannested, "What Have We Learned About Generalized Trust, If Anything?," *Annual Review of Political Science* 11 (2007): 413-36, 422. Compare Jan Delhey, Ken Newton, and Christian Welzel, "How General Is Trust In 'Most People'? Solving the Radius of Trust Problem," *American Sociological Review* 76:5 (2011): 786-807, 787.

[2] Nannested, "What Have We Learned," 415-416.

[3] Markus Freitag and Richard Traunmüller, "Spheres of Trust: An Empirical Analysis Of the Foundations of Particularized and Generalized Trust," *European Journal of Political Research* 48:6 (2009): 782-803, 784, and see their further references.

[4] To emphasize one point here: Freitag and Traunmüller distinguish between the two kinds of trust empirically (or, put another way, they show that the common distinction is supported by empirical findings). They use data from surveys of German citizens and show that these two spheres of trust have different foundations (they depend on different experiences and dispositions): trust in close relations (particularized trust) affects trust in more distant others (what they call generalized trust), and so particularized trust might be a necessary foundation for generalized trust; but the reverse is not true—higher levels of generalized trust do not affect levels of particularized trust. The material that follows does not challenge Freitag and Traunmüller's empirical findings but, from the perspective developed here, their findings should be re-framed in terms of particularized trust in two different spheres, without reference to generalized trust.

[5] Nannested, "What Have We Learned," 428-431.

[6] Margaret Urban Walker's account of trust is an exception in the philosophical literature even though she does not use the term 'generalized trust.' She defines trust in terms of A's expectation that B will comply with moral standards or social norms, where A and B are unknown others, persons in general—which is best taken to be an account of generalized trust. See for example her formulation at p. 84, where she defines "default trust" as the "unreflective and often nonspecific *expectation* that strangers or unknown others may be relied upon to behave in an acceptable and unthreatening manner" (my emphasis). The second section in the main text argues that we should reject expectation-accounts of this sort, though Walker's account could be re-cast in terms consistent with the account developed in section three, making reference to generalized moral obligations (where she talks of reliance in the passage just quoted). See Walker, *Moral Repair: Reconstructing Moral Relations After Wrongdoing* (Cambridge: Cambridge University Press, 2006). Accounts of trust in Walker's terms—as

expectations with respect to social norms—have been present in the sociology literature for some time; for example see Bernard Barber, *The Logic and Limits of Trust* (New Brunswick, NJ: Rutgers University Press, 1983).

[7] Nannestad, "What Have We Learned," 417.

[8] See the references in Nannestad, "What Have We Learned," 419.

[9] Delhey, Newton, and Welzel, "How General Is Trust In 'Most People'?" See, Francis Fukuyama, *Trust: The Social Virtues and the Creation of Prosperity* (New York: The Free Press, 1995).

[10] Jan Delhey and Ken Newton, "Predicting Cross-national Levels of Social Trust: Global Pattern or Nordic Exceptionalism," *European Sociological Review* 21:4 (2005): 311-327, 311. Compare, Freitag and Traunmüller, "Spheres of Trust," 782-783, who refer back to Delhey and Newton.

[11] Delhey and Newton, "Predicting Cross-national Levels of Social Trust," p. 311. They also describe their conception as being "close to" Russell Hardin's and to one other account—without further explanation—even though there are substantial differences between those two and Gambetta. This reflects their general lack of concern with the conceptual questions here.

[12] Diego Gambetta, "Can We Trust Trust?," in *Trust: Making and Breaking Cooperative Relations*, ed. D. Gambetta (New York: Basil Blackwell, 1988), 217.

[13] Guido Moellering, *Trust: Reason, Routine, Reflexivity* (Amsterdam: Elsevier, 2006).

[14] The main text focuses on Gambetta because Gambetta makes the role of expectations most explicit, and so provides a clear contrast with my alternative account as developed in section three. Mayer, Davis, and Schoorman reformulate Gambetta's account using the term 'expectation' instead of "subjective probability," and they also make explicit mention of the trusting party's vulnerability. They define trust as a willingness to be vulnerable given a set of expectations about the trusted party—expectations that the risks of a trusting action won't materialize, and by extension, that the risks are justified by the potential benefit. This account reduces trust to a willingness to take a calculated risk, so—anticipating the critical point that follows in the main text—there is no space to talk about the possibility of betrayal: if A takes a calculated risk in some interaction with B, given A's expectations about B's behavior, person A cannot claim to be wronged or betrayed if B acts contrary to those expectations. This is therefore inadequate as an account of trust. See Roger C. Mayer, James H. Davis, and F. David Schoorman, "An Integrative Model of Organizational Trust," *Academy of Management Review* 20:3 (1995): 709-734. The broad acceptance of Mayer, Davis and Schoorman in the management literature is documented in Denise M. Rousseau et al., "Not So Different After

All: A Cross-discipline View of Trust," *Academy of Management Review* 23:3 (1998): 393-404, and also in Roderick M. Kramer and Roy J. Lewicki, "Repairing and Enhancing Trust: Approaches to Reducing Organizational Trust Deficits," *The Academy of Management Annals* 4 (2010): 245-277. Nannestad, in the (2008) paper discussed in the introduction, also refers to Russell Hardin's account. According to Hardin, A trusts B to do x when A believes that B is motivated to act on A's interests, in this case by doing x. This is another formulation of the expectations account, not a distinct alternative we need to consider here. See Hardin, *Trust* (Cambridge, UK: Polity, 2006).

[15] Annette Baier, "Trust and Antitrust," *Ethics* 96:1 (1986): 231-260.

[16] Marc A. Cohen and John Dienhart, "Moral and Amoral Conceptions of Trust, with an Application in Organizational Ethics," *Journal of Business Ethics* 112:1 (2013): 1-13, and Marc A. Cohen, "Genuine, Non-calculative Trust with Calculative Antecedents: Reconsidering Williamson on Trust," *Journal of Trust Research* 4:1 (2014): 44-56.

[17] Margaret Gilbert, *Joint Commitment: How We Make the Social World* (Oxford: Oxford University Press, 2013).

[18] To anticipate the account of generalized trust below: in a case of this sort, with an explicit promise, there is also a background moral obligation to keep one's promises, so we might say that A trusting B to do x (particularized trust) depends on a context in which generalized trust, in the form of keeping promises, is assumed. Generalized trust is more fundamental in this sense though we cannot explain particularized trust relationships in terms of, or reduce particularized trust relationships to, generalized trust.

[19] With respect to process, my work defined trust in terms of invitations: when A trusts B to do x, A invites B to acknowledge and accept an obligation to do x. When—or if—B accepts the invitation, B makes a commitment and takes on that obligation. In that way trust creates an obligation. We might go beyond this point and note that the invitation could move in the opposite direction: B might invite A to rely on him or her, meaning that B invites A to trust him or her—and in that way B could initiate the trust relationship. The prototypical case is the one noted in the main text, involving an explicit promise. Trust-acceptance can create obligations when there is no pre-existing relationship, and it can also create new obligations in the context of existing relationships. And we could extend the point: we might see trust-invitations in both directions as seeking out a moral relationship (defined by obligations). Moreover, there could be cases in which a particularized trust relationship is forced onto B, perhaps in some emergency in which A must rely on B and, if B refuses to accept, A would be

greatly harmed.

[20] See Philip J. Nickel, " Trust and Obligation-ascription," *Ethical Theory and Moral Practice* 10:3 (2007): 309-319. Quotes in the main text are from p. 310 and p. 312 respectively. This same limitation applies to Richard Holton's account in "Deciding to Trust, Coming to Believe," *Australasian Journal of Philosophy* 72:1 (1994): 63-76, and also to Colin O'Neil, "Lying, Trust, and Gratitude," *Philosophy & Public Affairs* 40:4 (2013): 301-333. This critical point is developed in Cohen and Dienhart, "Moral and Amoral Conceptions of Trust," 4. Katherine Hawley's commitment conception of trust is an exception to the critical point in the main text: her account is motivated by different questions (distinguishing trust from reliance, and explaining both trust and distrust), and some of the details differ, but her account very much tracks the one applied here. See Hawley, "Trust, Distrust, and Commitment," *Nous* 48:1 (2014): 1-20.

[21] Fukuyama, *Trust*, 188.

[22] Here, moreover, we can see why Baier's own account of trust is problematic (even though the account of trust developed here makes use of her claim that trust involves the possibility of betrayal). According to Baier, A trusts B when A makes him- or her-self vulnerable by giving to B discretionary control of some object, relying on or depending on B's "goodwill" ("Trust and Antitrust," 235). Baier doesn't distinguish between particularized and generalized trust, but this cannot serve as a conceptual account of particularized trust: using one of the examples in the main text, a general obligation to act with goodwill could not account for the particular obligations involved when A trusts B to pick up his mail, Baier's reference to goodwill could not account for the possibility of betrayal in that trust relationship. For this reason, Baier is better understood as offering a conceptualization of generalized trust, but even as such her account is problematic, because it is unclear that (or why, or even when) A is entitled to B's goodwill—especially if goodwill goes beyond an obligation to keep promises and not harm others. In order to think about generalized trust in Baier's terms, she would have to establish that persons are bound by that particular obligation.

[23] The account here focuses on Garfinkel's (1963) essay because the experimental demonstration in that essay—discussed in the main text below—makes Garfinkel's point especially clear, and because the chapter on trust in the (1967) book is largely drawn from that earlier essay. See, Garfinkel, "A Conception Of, and Experiments With, 'Trust' As a Condition of Stable Concerted Action," in *Motivation and Social Interaction*, ed. O.J. Harvey (New York: The Ronald Press Company, 1963). And, Garfinkel, *Studies In Ethnomethodology* (Englewood Cliffs,

NJ: Prentice Hall, 1967).

[24] Garfinkel, "A Conception Of...'Trust'," 190.

[25] Garfinkel, "A Conception Of...'Trust'," 187.

[26] Anne Warfield Rawls uses the term "second-order" in this context in her paper, "An Essay on Two Conceptions of Social Order," *Journal of Classical Sociology 9*:4 (2009): 500-520, 509.

[27] Garfinkel, "A Conception Of...'Trust'," 193.

[28] Garfinkel, "A Conception Of...'Trust'," 190.

[29] Another example: at the beginning of the chapter on trust in the (1967) book, Garfinkel makes direct reference to the "moral order," but then immediately defines this moral order in non-moral terms, in the sense that it is not at all connected to philosophical or even ordinary conceptions of moral obligation; see p. 35.

[30] See the candidates for fundamental basic rules in Garfinkel, "A Conception Of...'Trust'," 210-214; the items are from sociologist Alfred Schutz's phenomenology of ordinary life.

[31] Anne Warfield Rawls, "An Essay... on Social Order," and Anne Warfield Rawls and Gary David, "Accountably Other: Trust, Reciprocity and Exclusion In a Context of Situated Practice," *Human Studies* 28:4 (2006): 469-497.