Attestation and Social Structure: The Problem of the Third Party

---

Editorial note: This is a speculative theoretical essay written in the voice and idiom of Mark Granovetter. It is not authored by him and should not be cited as his work. Its purpose is to extend, in a manner consistent with his intellectual commitments, the line of argument running from "The Strength of Weak Ties" (1973) through "Economic Action and Social Structure" (1985) into the domain of digitally mediated social life.

Abstract

Social relations have always produced traces, but only recently have those traces been produced, held, and acted upon by a party who is not a member of the relation. I call such a trace an **attestation**, and I argue that the general condition of digitally mediated social life is that every dyad has become a triad whose third member is an attestor: a platform, a model, a registry, an index. From this apparently modest observation a great deal follows, and most of it is counterintuitive. I argue, first, that social structure now consists of two distinct and systematically divergent graphs—the **enacted** network of relations actors carry out, and the **attested** network of relations third parties can see—and that the divergence between them is not measurement error but a structured social product, governed by an economy of visibility. Second, I argue that the celebrated informational advantage of weak ties was never a property of weak ties as such but a rent extracted from the structural ignorance of others; where a third party observes all clusters at once, that rent is competed away, and the bridging premium migrates from the weak tie to the **illegible tie**. Third, I argue that the embeddedness which produces trust has been partially expropriated: the memory and sanctioning capacity once distributed across mutual acquaintances is now centralized in an attestor who is not itself sanctionable, which relocates malfeasance from the relational layer to the attestation layer—it becomes cheaper to fake being embedded than to be embedded. Fourth, using a threshold model of the kind I developed for collective behavior, I show that attestation decisions are strategically interdependent, so that arbitrarily small shifts in the payoff to visibility generate discontinuous restructuring of the attested graph while the enacted graph is unchanged; because resources are allocated on the attested graph, enormous reallocations occur without any underlying social event having taken place. Finally, I consider machine actors, whose capacity to generate plausible attestations at near-zero marginal cost destroys the costliness on which the informativeness of attestation has always silently depended. The argument bears on economics, computer science, epidemiology, political science, law, and the whole enterprise of computational social science, whose central inferential move—reading the enacted graph off the attested one—I argue is biased in a direction that can be specified in advance.

Keywords: social networks; attestation; legibility; embeddedness; weak ties; platforms; threshold models; digital trace data

1. Introduction

It is a curious feature of the literature on digital social networks that after three decades of prodigious empirical effort we have acquired very little theory. We have acquired instruments of extraordinary power—we can observe hundreds of millions of dyads at daily resolution, something no sociologist of my generation imagined possible—and we have used them, in the main, to test propositions formulated for a world in which no such instruments existed. That the propositions frequently survive these tests has been taken as reassuring. I want to suggest that it should be taken as suspicious. When a theory built for conditions of pervasive mutual ignorance continues to fit data generated by an apparatus of pervasive observation, one possibility is that the theory is deep. Another possibility, which I think has received insufficient attention, is that the apparatus is manufacturing the confirmation.

Consider a well-known and admirable example. Rajkumar and colleagues (2022) conducted a series of large experiments on LinkedIn's "People You May Know" system and reported that moderately weak ties are more productive of job mobility than strong ones—a striking causal vindication of an argument I made almost fifty years earlier on the basis of a small sample of professional and technical workers in Newton, Massachusetts (Granovetter 1973, 1974). I confess to having read the paper with pleasure. But pleasure is a poor epistemic guide, and on reflection the design troubles me in a way I think is instructive. The intervention was a change in an algorithm. The ties whose strength was manipulated were ties as the platform represents them. The outcome, job transition, was measured by whether an actor updated a profile field. At no point in this admirable study does anyone observe a social relation. What is observed throughout is a relation's *representation*, held by a third party, produced under conditions that third party controls, and acted upon by that third party in ways that partly constitute the outcome being measured. This is not a criticism of the authors, who were as careful as their materials permitted. It is a claim about the materials. The proposition being tested and the instrument testing it are not independent, and the entity that owns the instrument is a party to the social structure under study.

I want to take this situation seriously as a *sociological* problem rather than a methodological nuisance. My claim in this essay is that the widespread mediation of social relations by third parties who record, represent, and act upon those relations constitutes a change in social structure of the same order as the changes that preoccupied the classical theorists—the extension of markets, the growth of bureaucracy, the separation of household and workplace. It is not a change in the *content* of social relations, and I will be at pains throughout to argue that most of what we know about how relations work continues to hold. It is a change in the *ontology of the tie*: in what a tie is made of, who holds it, and who can see it. And because so much of network theory—my own very much included—rests on tacit assumptions about who can see what, a change in visibility propagates through the theory in ways that are not at all obvious in advance.

Let me state the core observation plainly, since everything else follows from it. In the social world that sociologists theorized from Simmel ([1908] 1950) through the 1990s, a relation between two persons produced knowledge of itself as a by-product, and that knowledge was held by the parties and by whoever happened to observe them—which is to say, by other members of the same social structure, who were themselves observable in turn. Reputation worked because the people who knew about your conduct were people you also knew things about. Sanction

worked because the observers were embedded in the same web as the observed. The classical mechanisms of trust and control, which I discussed at length in Granovetter (1985), depend not merely on information but on the *symmetry* of the observational relation: A can be shamed by B because B can also be shamed by A, and both can be shamed by C.

Digital mediation breaks this symmetry, and it breaks it in a specific and consequential way. When two people interact through a platform, the platform learns of the interaction. The platform is not a member of the dyad in any conventional sense: it has no reciprocal obligations, it cannot be shamed, it is not itself observable by the parties, and it does not forget. Nor is it merely a channel, on the model of the telephone company, which learned that a call occurred but had no use for the fact. The platform maintains a model of the relation, updates the model, acts on the model, and—this is the crucial point—the parties know it does. The relation is therefore always conducted before a witness whose interests are unknown, whose memory is perfect, and whose account of the relation may come to matter more, for the allocation of jobs, credit, attention, and standing, than what the parties themselves did.

I will call the third party's record of a relation an **attestation**, and I will call a relation that generates one an **attested tie**. The choice of word is deliberate. A trace is passive; a measurement is neutral; an attestation is a *claim made by someone*, and claims have authors, audiences, and consequences. To say that Facebook records a friendship is to describe a database operation. To say that Facebook *attests* to a friendship is to describe a social act, performed by a party with interests, before an audience that acts on it. It is the second description that a sociologist needs.

From this starting point I develop four connected arguments, which I state here in compressed form and defend at length below.

First, social structure has bifurcated. There is an **enacted graph**, consisting of relations as they are carried out—the interactions, obligations, expectations, and histories that sociologists have always studied—and an **attested graph**, consisting of relations as third parties represent them. These are not the same graph, they are not related by random error, and the mapping between them is itself a social product, governed by what I will call the **economy of visibility** (Section 4): actors attest to relations where attestation pays and suppress attestation where it costs. Since the payoff structure of visibility varies systematically by position, resources, and domain, the divergence between the two graphs is patterned in ways that can be predicted.

Second, the strength of weak ties is a historically contingent result, and the conditions producing it are eroding. I argued in 1973 that weak ties are valuable because they bridge otherwise disconnected clusters and thus carry non-redundant information. I now think this formulation obscured the mechanism. The bridge was valuable not because it was weak, nor even because it was a bridge, but because *no one else could see across the gap*. The bridging actor was, in economic language, an arbitrageur, and her return was a rent on the structural ignorance of everyone else. A third party with a global view of all clusters simultaneously competes that rent away. The result is not that weak ties cease to exist—there have never been more of them—but that they cease to be *advantageous*, since the information they would have carried has already been delivered by the recommender. The bridging premium does not disappear; it migrates. It migrates to relations the third party cannot see. I call this the **opacity premium**, and I regard it as the central substantive prediction of this essay.

Third, embeddedness has been partly expropriated. The mechanisms by which concrete personal relations generate trust—mutual knowledge, the shadow of the future, the capacity of third parties to sanction—have been disaggregated, with the informational component centralized in an attestor and the sanctioning component left behind, weakened, with the parties. Trust in digital settings is therefore increasingly *trust in the attestor* rather than trust in the alter, and this has a consequence that I find genuinely alarming and that follows directly from my 1985 argument about malfeasance. I argued then that fraud is not deterred by generalized morality nor by institutional arrangements but by concrete relations, and that consequently the most damaging frauds are committed by the well-embedded, since embeddedness is precisely what makes large-scale trust possible. If trust now rests on attestation, then the rational locus of malfeasance shifts: the fraudster no longer needs to become embedded, which is slow and costly, but only to become *attested as* embedded, which is fast and cheap. Sybil attacks, engagement farming, synthetic reviews, purchased followings, credential laundering, and the fabricated interlocutor are not marginal pathologies of digital life. They are its characteristic form of malfeasance, and they are predicted by the theory.

Fourth, attestation decisions are strategically interdependent, and therefore aggregate in the discontinuous ways I described in my work on threshold models of collective behavior (Granovetter 1978). Whether I make a relation visible depends on how many others in my vicinity have made theirs visible, since visibility is informative only against a background, and concealment is suspicious only against a background of disclosure. Individuals differ in their thresholds. The consequence, familiar from that earlier work, is that the equilibrium level of attestation is a discontinuous function of the underlying distribution of dispositions: a small change in the payoff to visibility—a platform redesign, a norm shift, a regulatory intervention—can produce a large and sudden restructuring of the attested graph *with no change whatever in the enacted graph*. Since credit is extended, jobs are offered, prices are set, and reputations are made on the attested graph, this means that very large reallocations of real resources can occur in the complete absence of any underlying social event. I take this to be among the more disquieting implications of the framework, and among the most testable.

I close with the problem of machine actors, which I regard not as a separate topic but as the limiting case of the argument. The informativeness of an attestation, like the informativeness of any signal, has always rested on its costliness (Spence 1973; Zahavi 1975; Gambetta 2009). Maintaining a hundred relationships that would each vouch for you was expensive; that expense is what made the vouching credible. Generative systems reduce the marginal cost of producing a plausible attestation to approximately zero. When the cost of the signal falls to zero, the signal stops carrying information—not gradually, but at a threshold. What replaces it, I suggest, is a return to precisely the costly, slow, non-scalable, and structurally embedded forms of verification that digital infrastructure was supposed to have made obsolete. The end state of maximal legibility is a renewed premium on the unrecorded encounter.

Before proceeding I should say what this essay is not. It is not an argument that technology determines social structure; I will argue in the next section that such claims are a variant of the oversocialized conception I have criticized elsewhere, with the platform substituted for internalized norms. Nor is it an argument that nothing has changed and that classical network theory applies without modification, which is the corresponding undersocialized error. It is an argument that a specific and identifiable feature of digital infrastructure—the presence of a non-reciprocating third party who holds the record of relations and acts upon it—alters the operation

of mechanisms we already understand, in directions we can derive. That is all theory can reasonably be asked to do.

2. Under- and Oversocialized Conceptions of Digital Life

Wrong (1961) observed that sociology oscillates between two errors: an oversocialized conception, in which behavior follows mechanically from internalized norms, and an undersocialized conception, in which behavior follows mechanically from atomized self-interest. I have argued (Granovetter 1985, 2017) that these apparent opposites share a common structure: both dispose of ongoing social relations, one by absorbing them into the actor's interior, the other by denying their relevance. Both permit the analyst to skip the messy middle level where actual social life occurs.

The literature on digital networks reproduces this pair with remarkable fidelity, and I think identifying the reproduction is worth some effort, because the correct position is neither a compromise between the two nor a bit of each.

The **oversocialized** position in digital scholarship holds that platforms determine behavior. Architecture is destiny; the algorithm shapes the self; affordances constrain action; the recommender system produces polarization, addiction, loneliness, or radicalization as an output of its optimization objective. Actors in these accounts have interior lives but no relations: they confront the platform individually, are individually shaped by it, and aggregate into publics by summation. This is precisely the structure of the oversocialized argument, with the recommender in the position formerly occupied by internalized culture. It shares the older error's central weakness, which is that it cannot explain variation. If the architecture determines the outcome, why do outcomes differ so markedly across populations facing identical architectures? The empirical record on political polarization is instructive here: the effects of platform exposure are real but modest, heterogeneous, and heavily conditioned by offline social position (Bakshy, Messing, and Adamic 2015; Bail et al. 2018; Bail 2021). The architecture is a constraint on relations, not a substitute for them.

The **undersocialized** position holds that digital ties are simply ties, that the platform is a channel, and that the accumulated results of network analysis therefore transfer intact. On this view the interesting questions are of scale and measurement: we now have more data about the same objects. This position has the enormous advantage of allowing existing work to continue without disturbance, which may explain some of its popularity. Its weakness is that it treats the medium as transparent, when the whole burden of a generation of work on infrastructure and code is that media are constitutive rather than neutral (Lessig 1999; Star 1999; Gillespie 2014). The telephone was, roughly, a transparent medium: it transmitted the relation without modeling it, without acting on the model, and—crucially—without the parties adjusting their behavior in anticipation of how the model would be used. A platform does all three. To treat these as equivalent is to make the same move that treats a market as a neutral meeting-place for pre-existing preferences, and I have spent a good deal of my career objecting to that move.

The correct position, as usual, is at the level of concrete relations, but the relations in question now include one that classical network theory did not represent: the relation between each actor and the attestor. This is the key analytical step. Most treatments of platforms place them *outside* the network, as environment, context, or infrastructure. I propose instead to place the attestor

inside the network, as a node—an extremely peculiar node, with properties no human node has, but a node nonetheless, connected to every actor and every tie. This is a thoroughly relational move in Emirbayer's (1997) sense, and it has the formalist virtue that its consequences can be worked out rather than merely asserted (Erikson 2013).

Once one draws the graph this way, several things become visible at once. The attestor has degree equal to the size of the population, which no human actor ever approaches; it therefore occupies, trivially and permanently, every structural hole in the system (Burt 1992). It has ties that are entirely non-reciprocal in information: it observes without being observed, which is the pure limiting case of the asymmetric dependence Emerson (1962) analyzed. Its ties do not decay, since it does not forget, whereas all human ties require maintenance and decay without it—a fact whose importance for network structure I do not think has been adequately appreciated (see Burt 2002 on tie decay). And it is not subject to the constraint that generates the forbidden triad: it has no time budget, no cognitive limit, and no need for balance among its relations.

That last point deserves emphasis, because the forbidden triad is the engine of the weak-tie argument and it rests on assumptions about human limitation. I argued in 1973 that if A has strong ties to both B and C, then a tie between B and C is highly likely, for three reasons: time (strong ties consume time, so strongly-tied alters are frequently co-present), homophily (A's strong ties are similar to A and therefore to each other), and balance (in Heider's sense, the configuration is psychologically unstable). Every one of these reasons is a fact about *human* actors under conditions of scarce time and bounded cognition. An attestor faces none of them. It is strongly tied, in the relevant informational sense, to everyone, and the triads it participates in are never forbidden. Nor, for the same reason, does it occupy any structurally equivalent position with any other node in the system (Lorrain and White 1971), which is one way of saying that it is not comparable to the actors it observes. This is not a curiosity. It means that the structural results derived from the forbidden triad—that bridges must be weak, that local bridges are the channels of novel information, that removing weak ties fragments the network while removing strong ties does not—hold only for the subgraph of human nodes, and cease to describe the system once the attestor is included. In the full graph there are no structural holes, no bridges, and no clusters, because one node connects everything to everything. The entire architecture of advantage that network theory identified is an artifact of a graph from which the most connected node has been deleted.

I want to be careful here, because this observation can be overread. It does not follow that structural position no longer matters to human actors. It follows that structural position matters *conditionally on what the attestor does with its position*, which is an empirical question with a determinate structure, and that is precisely the theory I now want to develop.

3. The Argument in Brief

For the reader who wants the skeleton before the flesh, the argument runs as follows.

A relation may be *enacted* (carried out between parties) and it may be *attested* (recorded and represented by a third party). These are logically independent. Four cases exist: relations both enacted and attested; relations enacted but unattested; relations attested but not enacted; and relations neither, which need not detain us. The three substantive cases define the field of study.

The mapping from enactment to attestation is not technical but social. Actors have some control over whether and how relations are attested, and they exercise it strategically, because attestation has payoffs and costs that vary by relation, position, and domain. This produces the economy of visibility (Section 4).

Third parties—employers, lenders, insurers, states, matchmakers, recommenders, and social scientists—act on the attested graph, because it is the only graph they can see. Their actions have real consequences. Therefore the attested graph is *causally efficacious independent of the enacted graph*, which is the sense in which representation becomes constitutive rather than merely descriptive (compare Espeland and Sauder 2007; Fourcade and Healy 2017).

Because the attestor sees all clusters at once, informational arbitrage across clusters is no longer available to human bridges; the value of bridging migrates to relations that are enacted but unattested (Section 5).

Because the attestor holds the record of the relation, the informational basis of trust is centralized in a party that cannot itself be sanctioned; trust becomes trust in attestation; and malfeasance relocates to the production of false attestations (Sections 6 and 7).

Because attestation decisions are interdependent, they aggregate discontinuously, so that the attested graph exhibits sudden reorganizations unrelated to any change in the enacted graph, and resource flows follow the attested graph (Section 8).

Because generative systems make plausible attestation costless, the informational value of attestation collapses, and verification returns to costly, embedded, non-scalable forms (Section 9). ## 4. Two Graphs: Definitions and Elementary Properties

I will now be more precise, though I want to enter the usual caveat that formalization in sociology is useful chiefly for exposing what one has assumed, and is dangerous when mistaken for demonstration.

Let N be a population of human actors. The **enacted graph** $G_E = (N, E_E)$ contains an edge (i, j) when i and j stand in a social relation of some specified type—they interact, they have obligations, they would recognize a claim from the other. This is the graph sociology has always meant when it said "network." It has never been fully observable and never will be. Every survey instrument, every ethnography, every name generator is an attempt to estimate it (Marsden 1990).

Let A be a set of **attestors**: parties who are not members of the relations they record. An attestor $a \in A$ maintains an **attested graph** $G_a = (N_a, E_a)$, where $N_a \subseteq N$ is the population it can see and E_a its representation of relations among them. Attestors include platforms, but also credit bureaus, professional registries, citation indexes, immigration authorities, insurers, and—I include us deliberately—social scientists holding trace data.

Define the **attestation function** $\alpha_a : E_E \to [0,1]$, giving the probability (or, if one prefers, the intensity) with which an enacted relation appears in a's representation, and the **fabrication function** ϕ_a, giving the rate at which edges appear in E_a with no enacted counterpart. The naive assumption underlying nearly all analysis of digital trace data is that α_a is a constant and ϕ_a is zero: that the attested graph is a uniform random sample of the enacted graph, plus noise. I will

argue that α_a is a systematic function of properties of the relation, of the actors, and of the domain, and that ϕ_a is not only nonzero but *strategically produced*. If either claim holds, inference from G_a to G_E is biased, and biased in a direction we can specify—which is better than unbiased, in the sense that a known bias is a resource, but worse than the literature assumes.

Four elementary properties of this setup deserve statement.

(i) Attestation is not observation. An attestor does not see the relation; it sees events it treats as evidence of a relation, and it applies a classification. Two people who message daily are attested as strongly tied; two people who see each other daily at the same workplace and never message are not attested at all. The attestor's category system—friend, connection, follower, contact, match, co-author—is not a description of a social form but an *imposition* of one, with all the consequences that follow from the installation of a classification into working infrastructure (Bowker and Star 1999), and actors then arrange their conduct with reference to the imposed category. This is Hacking's (1986) looping effect, applied to relations rather than to persons, and I regard it as one of the most powerful mechanisms in the entire domain. The categories of the attestor become the categories of social life, not because anyone believes in them, but because they are what pays.

(ii) Attestation is asymmetric in control. In most systems, an edge appears in E_a only if at least one party takes an action, but neither party can unilaterally remove the underlying evidence. Attestation is thus easy to create and hard to destroy—the reverse of the enacted tie, which is hard to create and decays on its own. This asymmetry alone predicts that G_a will be systematically *denser* than G_E and will exhibit far less turnover, which is exactly what one finds when platform network data are compared to contemporaneous survey data.

(iii) Attestation is stratified. The capacity to control one's attestations—to have relations that are not recorded, to have records corrected, to segment audiences—is unequally distributed and follows familiar lines of advantage. The wealthy conduct consequential business in rooms with no logs; the poor are attested continuously by welfare agencies, employers' scheduling software, parole systems, and gig platforms. Fourcade and Healy (2013, 2017) have described the resulting stratification order with more care than I can here, and it belongs to a longer history of the political consequences of quantification (Desrosières 1998; Porter 1995). My addition is structural: it is not only that some people are more surveilled than others, but that *the shape of the graph the attestor holds differs systematically by class position*, which means that any algorithm operating on that graph will make systematically different kinds of errors about different populations, without any discriminatory intent and indeed without any way of detecting the problem from inside the data.

(iv) The attestor is a party to the relation it attests. This is the point I have already labored and will not labor further, except to note its formal consequence: any model that treats G_a as a measurement of G_E has misspecified the system, because G_a enters the dynamics of G_E. The attestor recommends, ranks, prices, and matches on the basis of G_a, and those actions change E_E. The measurement is inside the causal loop. Physicists were forced to take this seriously a century ago and it changed their science considerably; I do not think social scientists have yet absorbed the analogous point, though Merton (1948) saw it clearly and MacKenzie (2006) has shown in detail how it operates in financial markets.

The economy of visibility

Why should α vary? Because attestation has consequences, and actors anticipate them.

Consider an actor i with an enacted relation to j. Attesting the relation—accepting the connection request, tagging the photograph, listing the reference, appearing together in a recorded meeting—produces a stream of benefits and costs. The benefits are chiefly of three kinds: *access*, since attested relations are the ones third parties will route opportunities through; *credibility*, since an attested relation to a high-status alter transfers status, which is the classical mechanism Podolny (2005) analyzed; and *convenience*, since the infrastructure works better on attested relations.

The costs are of four kinds, and it is the costs that have been undertheorized. First, *audience collapse*: an attested relation is visible to audiences before whom the actor maintains different presentations, in exactly Goffman's (1959) sense, and the attestation forcibly merges front and back regions (Marwick and boyd 2011). Second, *inferential exposure*: an attestation reveals not only the relation but everything a competent inference engine can derive from it, including facts the actor has deliberately withheld—one's ties disclose one's health, politics, sexuality, and solvency more reliably than one's statements do. Third, *obligation*: an attested relation is a claim others may invoke, and there are many relations one wishes to maintain precisely on the condition that no third party can prove they exist. Fourth, *strategic disclosure*: in any bargaining situation, revealing one's alternatives and alliances reduces one's position.

We may then write the actor's decision as a comparison of an expected benefit B_{ij} and cost C_{ij}, and define the **visibility differential** $v_{ij} = B_{ij} - C_{ij}$, with attestation occurring when $v_{ij} > 0$. Nothing is gained by pretending this is more than a bookkeeping device. What matters is that v_{ij} is not randomly distributed across relations. The following propositions state its structure.

Proposition 1 (Instrumental bias). *Attestation rates are higher for relations whose value is realized through third parties than for relations whose value is realized within the dyad.* Professional acquaintance, patronage, and credentialing relations are attested at high rates because their value consists precisely in being knowable by others. Kinship, intimacy, confidence, and dependency are attested at low rates, and the deepest relations often at the lowest rates of all, since their value is intra-dyadic and their exposure is costly—indeed, part of the work of maintaining an intimate relation consists in marking it off from the accounting practices that govern other ties (Zelizer 2005). It follows that the attested graph systematically over-represents instrumental relations relative to affective ones, and that any theory built from trace data will systematically overstate the instrumental character of social life. I note without pleasure that this is a bias in favor of the economist's picture of the world, generated not by ideology but by the sampling properties of the instrument.

Proposition 2 (Asymmetry of illicit relations). *Relations that violate a norm salient to any audience with access to the attestation are suppressed at rates increasing in the severity of the sanction and the breadth of the audience.* This yields the familiar result that the attested graph is far cleaner than the world—corruption networks, sexual networks, dissident networks, and criminal networks are precisely the networks that do not appear—but it yields a less familiar corollary. Since suppression is costly and imperfect, the *residue* of illicit ties in the attested graph is not a random sample of illicit ties but a sample biased toward the incompetent and the

unprotected. Enforcement systems trained on that residue will therefore learn to detect incompetence rather than wrongdoing, and will apply that detection with great confidence and increasing scale.

Proposition 3 (Status gradient in opacity). *The ratio of enacted to attested ties increases with position in the status order.* High-status actors have more relations that pay without being recorded, more capacity to conduct relations in unattested channels, and more to lose from exposure. Low-status actors are attested involuntarily and comprehensively. Consequently the attested graph is not merely a biased sample of the enacted graph; it is a sample whose *bias varies monotonically with the outcome variable of most sociological interest*, which is the worst case for inference.

Proposition 4 (Reactive densification). *Where third parties reward degree, actors attest relations they have not enacted; where third parties reward tie strength, actors enact interaction they would not otherwise have had.* This is Campbell's law and Goodhart's law in network form, and it is the mechanism by which the attested graph becomes a poor measure precisely to the extent that it becomes an important one. Espeland and Sauder (2007) documented the process in detail for law school rankings, and a considerable literature has traced the same reactivity through audit regimes of every description (Strathern 1997; Power 1997; Muller 2018). The network case is the same process operating on relations rather than attributes, and it is far less visible because relations are harder to audit than test scores.

Taken together, these four propositions imply that the attested graph is denser, more instrumental, more respectable, more persistent, and more class-skewed than the enacted graph, and that the discrepancy grows with the stakes attached to attestation. Every one of these implications is testable by comparing platform data with contemporaneous survey or ethnographic data on the same population, and my strong impression is that where such comparisons have been made, they come out in this direction (Kossinets and Watts 2006; Tufekci 2014; Ruths and Pfeffer 2014; Lazer et al. 2014).

5. The Obsolescence of the Bridge

I now turn to the argument I expect will be most contested, since it concerns the undoing of a result I am associated with.

Recall the structure of the 1973 argument. Novel information reaches an actor from outside her cluster. Only ties that bridge clusters can carry it. By the forbidden-triad argument, bridges must be weak. Therefore weak ties carry novel information, and actors rich in weak ties enjoy informational advantage—in job search, in innovation diffusion, in the capacity of communities to mobilize. Burt (1992, 2004) restated the mechanism in terms of structural holes and showed convincingly that the advantage is a return to *brokerage position* rather than to tie weakness as such.

I want to press one step further than Burt did, and ask what kind of return this is. It is, plainly, an arbitrage return. The broker profits from a price differential across two markets that cannot see each other. Everything we know about arbitrage in economics tells us the return is a function not of the differential itself but of the *exclusivity of access* to it, and that returns are competed away as access spreads. Burt (2004) saw this and noted that brokerage returns decay as a hole is filled. What has not been drawn out, so far as I know, is the implication for the digital case, which is

not that individual holes get filled but that a single actor comes to occupy all of them simultaneously and permanently.

Let me put the point as directly as I can. **The strength of weak ties was never a theory about ties. It was a theory about ignorance.** The weak tie was valuable because of a fact about everybody else: that no one had a view of the whole. Remove that fact and the theorem does not survive it.

Formally, let the value to i of a bridging tie to cluster K be

$$W_i(K) = \Delta(K) \cdot (1 - c(K)) \cdot \tau,$$

where $\Delta(K)$ is the informational differential between i's cluster and K, $c(K)$ is the probability that a third party has already delivered the relevant information to i, and τ captures timeliness. Classical network theory studied the case $c \approx 0$ and could safely suppress the term; the small-world results that so impressed us all (Watts and Strogatz 1998) described how short the paths were, not who else could already see along them. Modern recommendation infrastructure drives c toward one for exactly the kinds of information that recommenders handle well: job openings, potential collaborators, market prices, news, available partners, and the existence of similar others. As $c \to 1$, $W_i(K) \to 0$ *regardless of* Δ and regardless of how many bridges i maintains.

Proposition 5 (Erosion of the bridging premium). *The return to brokerage in a domain declines monotonically with the coverage and accuracy of third-party intermediation in that domain.* The prediction is sharply domain-specific and thus falsifiable: brokerage returns should have fallen most where matching infrastructure is best developed (routine white-collar labor markets, commodity trade, mass-market dating, standardized credit) and should be intact or rising where it is poorly developed (elite labor markets, illiquid and idiosyncratic assets, political coalition-building, scientific problem selection, corruption). This is a comparative-static claim across domains, not a claim that networks no longer matter, and I would regard evidence of uniform decline as disconfirming.

Proposition 6 (Redundancy inflation). *Third-party intermediation increases the number of weak ties while decreasing the average non-redundancy of each.* The recommender's function is to identify likely ties, and it identifies them chiefly by structural and attributional similarity—friends of friends, people like you—which is to say that it operationalizes homophily and organizational focus as a prediction rule (Feld 1981; McPherson, Smith-Lovin, and Cook 2001). It therefore generates ties that close triads rather than bridging holes. Every such tie increases measured degree and decreases the actual diversity of the neighborhood. Actors thus experience their networks as expanding at the very moment those networks are becoming informationally poorer, which is a fair description of the subjective texture of professional networking sites and, I suspect, part of the explanation for the peculiar sense of exhausted futility they induce. Aral and Van Alstyne (2011) showed that bandwidth and diversity trade off against each other in ordinary networks; intermediation resolves the trade-off in the wrong direction, purchasing bandwidth with diversity at every step.

Proposition 7 (Migration of the premium: the opacity premium). *As third-party coverage rises, the informational advantage shifts from ties that are weak to ties that are unattested.* If everything the attestor sees is available to everyone the attestor serves, then non-redundant information is by definition information the attestor does not have. The valuable relation is the

one that leaves no record: the conversation over dinner, the call from a landline, the encounter at a conference bar, the note passed inside an institution whose logs are not integrated. I want to insist that this is not a nostalgic or romantic claim. It is a straightforward implication of treating brokerage as arbitrage: the rent accrues wherever ignorance remains, and the only remaining ignorance is the attestor's.

Proposition 8 (Stratification of opacity). *Because the capacity to sustain unattested relations is distributed by Proposition 3, the opacity premium accrues disproportionately to those already advantaged, and does so through a mechanism invisible in attested data.* This is, I think, the most consequential single claim in the essay. The classical weak-tie mechanism was in a modest way *democratizing*: acquaintanceship is cheap, and the poor have acquaintances too, so the mechanism distributed some information downward, a finding that motivated a generation of policy interest in "bridging social capital" (Bourdieu 1986; Putnam 2000; Chetty et al. 2022). The opacity premium is anti-democratizing, because unattested relations are expensive to maintain: they require private space, discretionary time, physical co-presence, institutional protection, and counterparties with something to lose. Where the weak tie was a poor person's asset, the illegible tie is a rich person's asset. And—here is the sting—because the mechanism operates in the unattested layer, every study conducted on trace data will find that networks matter less than they used to, and will be wrong. What will have happened is not that networks stopped mattering but that the part of the network that matters moved out of view.

I regard Proposition 8 as the essay's central hypothesis and the one on which I would most like to be proven wrong. Chetty and colleagues (2022) showed with extraordinary data that cross-class friendship predicts mobility. Their data are attested data. If Proposition 8 is right, the true cross-class structure among the advantaged is substantially denser than their instrument records, and the mobility effect they estimate is a lower bound on the advantage conferred by an upper-class network and an upper bound on what an intervention operating in the attested layer could achieve. ## 6. Third-Party Embeddedness: The Expropriation of Trust

I argued in 1985 that the problem of order in economic life is solved neither by generalized morality nor by institutional arrangements designed to prevent malfeasance, but by concrete personal relations and the structures they compose. Two components did the work in that argument, and it is worth separating them now because digital mediation separates them in practice.

The first is **relational**: I trust you because we have a history, because I know what you are like, and because our relation has a future I do not want to spoil. The second is **structural**: I trust you because we have mutual acquaintances who will learn of your conduct, whose good opinion you require, and who can impose costs on you. The second component is what allows trust to extend beyond dyads one has personally tested, and it is therefore what makes complex economic life possible at all. Coleman (1988) formalized the point as network closure; Gambetta (1993) showed with unpleasant vividness what happens to markets when it fails; Uzzi (1997) traced its operation and its limits in the garment trade.

Now observe what digital mediation does to the structural component. It does not remove it. It *disaggregates* it, taking the informational function—who knows what about whose conduct—and centralizing it in the attestor, while leaving the sanctioning function distributed, weakened, and largely unexercised.

The result is a form of embeddedness with a peculiar topology, which I will call **third-party embeddedness**: the parties are connected to each other through the attestor rather than through mutual acquaintances who are themselves connected to each other. Topologically this is a star, not a closed triad. And a star has properties a closed triad does not.

Proposition 9 (Non-reciprocal closure). *Third-party embeddedness supplies information without supplying mutual vulnerability.* In classical closure, C's capacity to sanction A is limited by A's capacity to sanction C, and this mutual limitation is what makes closure a stable social arrangement rather than a tyranny. The attestor has no counterparty risk. It can revoke standing, reduce reach, deny access, or simply cease to attest, and the actor has no reciprocal instrument. Every substantive account of how reputation systems restrain conduct assumes the closed-triad case and thereby misses the fact that the modern reputational third party is not a neighbor but a sovereign.

Proposition 10 (Displacement of the object of trust). *As attestation coverage rises, the object of trust shifts from the alter to the attestor.* One does not, in an unfamiliar city, trust the driver; one trusts the rating system, the tracked route, and the recourse mechanism. This is efficient, and it enables transactions that would otherwise be impossible—I do not want to be misread as objecting to it. But it has a structural consequence: trust becomes *fragile in a new way*, because it is now concentrated. Distributed relational trust degrades gracefully—when one relation fails, the rest hold. Centralized attestational trust fails catastrophically, because every relation in the system depends on the same node. This is the network-theoretic point about scale-free topologies and their tolerance of random failure but vulnerability to targeted attack (Barabási and Albert 1999; Albert, Jeong, and Barabási 2000) applied to trust rather than to connectivity, and it predicts that trust in digitally mediated systems will appear unusually robust for long periods and then fail all at once.

Proposition 11 (Erosion of relational skill). *Where attestation substitutes for relational judgment, the capacity for relational judgment atrophies, increasing dependence on attestation.* This is not a lament about kids today. It is a claim about the distribution of investment. If assessing a counterparty through slow relational work yields no advantage over consulting a score, no one will do the slow work, and the tacit skills involved—which are real, learned, and largely unarticulated—will not be transmitted. The system thus generates increasing returns to itself, and populations lose the ability to operate the fallback precisely as the fallback becomes necessary.

Proposition 12 (Trust without history). *Third-party embeddedness permits trust between parties with no relational history and no mutual acquaintances, and therefore extends the range of cooperation enormously while lowering its average depth.* The gain is real: the classical mechanism could not support a stranger economy at planetary scale. The cost is that the cooperation so produced is *thin*—it survives only within the scope of the attestor's enforcement. Cooperation of this kind does not generalize, does not accumulate into obligation, and does not produce the kind of relations from which further cooperation grows. It is the interactional analogue of a spot market—and markets, as White (1981) insisted, are made of durable mutual observation among a small set of producers, not of anonymous encounters. Those who have hoped that digital infrastructure would build community have, I think, been misled by a confusion between the volume of cooperation and its depth; the theory predicts that the first

should rise sharply while the second falls, which is consistent with what has been observed regarding the coexistence of unprecedented connectivity with rising reports of isolation.

7. Malfeasance Moves to the Attestation Layer

The 1985 essay's least comfortable claim was about fraud. I argued that since trust is produced by concrete relations, and since fraud on a large scale requires trust, the most damaging malfeasance is committed by the well-embedded. Trust and betrayal are not opposites; the second is parasitic on the first. Every large financial fraud in the succeeding forty years has, so far as I can tell, confirmed this.

Apply the same logic to attestation. If trust is now produced by attestation, then the rational fraudster does not invest in becoming embedded. He invests in becoming *attested as* embedded. And this is a far better investment, because embeddedness is slow, costly, irreversible, and requires that one actually behave well for a considerable period, whereas attestation is fast, cheap, reversible, and requires only that one produce the right traces.

Proposition 13 (Relocation of malfeasance). *Malfeasance concentrates at the layer where trust is produced. Where trust is produced by relations, fraud takes the form of betrayed relations; where trust is produced by attestation, fraud takes the form of fabricated attestation.* The characteristic crimes of the attested order are therefore not confidence tricks in the classical sense but *counterfeits of social structure*: purchased followings, review farms, sockpuppet networks, mutual-citation rings, fabricated employment histories, credential laundering through obscure institutions—an inversion of the illegitimacy discount Zuckerman (1999) identified, in which one now manufactures categorical fit rather than suffering for its absence—synthetic identities assembled from fragments of real ones, and—the limiting case—wholly generated interlocutors. The computer science literature has studied these under the heading of Sybil attacks (Douceur 2002) and has generally treated them as a security problem. They are not a security problem. They are the predictable form that opportunism takes given the structure of the trust regime, and no amount of detection engineering will remove them, since the incentive is structural.

Proposition 14 (The cost differential). *The rate of attestation fraud in a domain is increasing in the ratio of the cost of becoming embedded to the cost of appearing embedded.* This ratio has been falling continuously since the invention of the platform and has fallen precipitously with generative systems. The proposition explains variation: attestation fraud should be rare where attestations require costly, verified, adversarially audited acts (a securities filing, a marriage record, a criminal conviction) and pervasive where they require only cheap ones (a follow, a review, a profile claim, a citation). It also implies that hardening verification does not reduce fraud but *raises its price and concentrates it*, moving it from many small actors to a few well-capitalized ones—which is what one observes in mature review markets, and which is exactly the pattern I would predict for any regulatory scheme premised on identity verification.

Proposition 15 (Detection asymmetry). *An attestor can detect counterfeit attestations only by their structural signature, and can therefore detect only counterfeits that are structurally naive.* Detection systems look for anomalous subgraph patterns—implausible density, temporal synchrony, degree distributions that no organic process would generate. But the structural properties of organic networks are published, well characterized, and easy to simulate. A counterfeiter who reads the same literature as the detector produces subgraphs that are

indistinguishable in every statistic the detector computes. The consequence is a peculiar and I think under-appreciated equilibrium: **detection selects for sophistication**, so that the observed rate of counterfeiting falls while the true rate rises, and the attestor's confidence in its own graph increases exactly as that graph becomes less trustworthy. I would call this the central pathology of the attested order. It is the same structure as the residue problem in Proposition 2, and it recurs so often in this framework that I am inclined to treat it as a general principle: *any system that acts on representations of social structure while adversaries can shape those representations will grow more confident as it grows more wrong.*

8. Thresholds, Interdependence, and Discontinuous Restructuring

I turn now to aggregation, which is the problem I have always thought sociology exists to solve. The interesting question is never how individuals behave but how individual behaviors compose into collective outcomes, and the persistent finding of my work on this—beginning with the threshold models of collective behavior (Granovetter 1978; Granovetter and Soong 1983, 1986)—is that the composition is *not* a summing, that intuitions built on averages fail badly, and that populations with nearly identical distributions of dispositions produce wildly different outcomes.

Attestation is a textbook case of interdependent decision, for two reasons that push in opposite directions.

The first is that **the benefit of attesting increases with the number of others who attest**. An attested professional profile is worthless if no one else has one and indispensable if everyone does; a review is informative only within a corpus; a network is valuable in proportion to its coverage. This is a standard positive externality.

The second, and more interesting, is that **the cost of not attesting also increases with the number of others who attest**. When attestation is rare, absence carries no information. When it is universal, absence is a signal, and specifically a signal of something to hide. This is a phenomenon for which I do not think we have a good name; I will call it the **suspicion of the unattested**. It converts privacy from a default into a claim, and a claim that requires justification.

Let r be the proportion of an actor's relevant reference set that has attested a given class of relation. Let actor i attest when $r \geq \tau_i$, where the threshold τ_i summarizes i's private costs and benefits. Let F be the cumulative distribution of τ in the population. Then the dynamics are $r_{t+1} = F(r_t)$, equilibria are fixed points $r^* = F(r^*)$, and an equilibrium is stable when $F'(r^*) < 1$. This is exactly the apparatus of the 1978 paper and I make no apology for reusing it; the point of a mechanism is that it recurs.

Three results follow immediately, and all three matter.

Proposition 16 (Discontinuity). *Arbitrarily small changes in the distribution of attestation thresholds—or in the payoff structure that generates them—can produce arbitrarily large changes in the equilibrium level of attestation.* A platform alters a default from opt-in to opt-out; an employer begins consulting a registry; a norm shifts among a professional cohort. If the perturbation moves F across a tangency with the diagonal, the stable equilibrium jumps, and the attested graph reorganizes wholesale. Nothing has happened to anyone's relations. What has changed is which relations are visible.

Proposition 17 (Decoupling of representation from reality). *Because third parties allocate resources on the attested graph, discontinuous restructuring of the attested graph produces discontinuous reallocation of real resources in the absence of any change in the enacted graph.* This is the proposition I would most like to see tested, and it is testable: it predicts that credit, hiring, insurance, visibility, and price outcomes should exhibit jumps timed to changes in attestation regimes rather than to changes in underlying conduct, and that the jumps should be concentrated among actors near the margin of the attestation decision. It also implies that a great deal of what we interpret as social change—rising polarization, collapsing institutional trust, shifting patterns of association—may in part be change in the attestation regime rather than in social structure. I do not claim it is *all* artifact. I claim that we currently have no way of telling, because we measure social structure with the very instrument whose properties are changing, and this seems to me an intolerable situation for an empirical discipline.

Proposition 18 (Hysteresis). *Attestation equilibria are path-dependent and asymmetric: the level of visibility at which a population begins attesting is higher than the level at which it stops.* Since attestations are hard to destroy (Property ii), and since the suspicion of the unattested makes withdrawal costly once disclosure is general, a population that has moved to a high-attestation equilibrium cannot return to the previous one even if every member prefers it. This is a genuine Pareto-inferior trap of the sort that has occupied Schelling (1978) and the collective-action literature, and it is worth recalling from Salganik, Dodds, and Watts (2006) that outcomes generated by cumulative interdependent choice are both highly unequal and, in the relevant sense, unpredictable from the properties of the objects being chosen, and it explains an otherwise puzzling fact: that widespread and sincere preference for privacy coexists with near-universal disclosure, without anyone having been coerced and without any inconsistency in anyone's preferences. The usual explanation, that people say they value privacy but do not act on it, is a claim about individual hypocrisy where a structural explanation is available and better.

A further observation connects this section to Section 5. The opacity premium of Proposition 7 is bounded by the suspicion cost of Proposition 18: an unattested relation is valuable because it is invisible, but invisibility is itself increasingly read as evidence. Whether opacity pays therefore depends on whether the actor can bear the suspicion cost. Elite actors can—their standing is established by other means, and their opacity reads as discretion. Marginal actors cannot—their opacity reads as concealment and is priced accordingly. The two mechanisms thus reinforce each other in the same direction: **the returns to invisibility accrue to those whose visibility is not required**, which is a compact statement of the stratification order the framework predicts.

9. Machine Actors and the Collapse of Costly Attestation

I have deferred until now the question that will occur to any reader in the middle 2020s: what happens when some of the parties producing attestations are not people?

I want to resist the temptation to treat this as a new topic. It is the limiting case of a mechanism already in the argument. Every theory of signaling since Spence (1973)—and in biology since Zahavi (1975)—rests on the proposition that a signal is informative only if it is costly to produce falsely. Attestations have been informative for the same reason. A hundred people vouching for you was informative because acquiring a hundred willing vouchers took years of conduct that was hard to fake. The information was never in the vouchers. It was in the cost.

Generative systems drive the marginal cost of producing a plausible attestation toward zero: plausible text, plausible interaction histories, plausible faces, plausible references, plausible interlocutors with plausible relations to other plausible interlocutors, at whatever scale is required. The relevant fact is not that the outputs are convincing to people. It is that they are convincing to *the attestor's classifier*, which is a much lower bar and, being a fixed function, an exploitable one.

Proposition 19 (Signal collapse). *The informational content of an attestation class falls to zero at the point where the marginal cost of fabrication falls below the marginal benefit of fabrication for the marginal actor—and the transition is a threshold, not a gradient.* Signal collapse is a cascade of the type described in Section 8: as fabrication spreads, the pooling equilibrium breaks, honest actors' attestations lose value, and the incentive to invest in honest attestation disappears, which accelerates the collapse. Markets for lemons (Akerlof 1970) unravel this way, and reputation systems are markets for lemons in which the good being traded is social structure itself.

Proposition 20 (Return of the costly channel). *Following signal collapse, verification reverts to channels whose costliness is guaranteed by physics or by institutional structure rather than by convention.* Namely: physical co-presence, which cannot be manufactured at scale; long duration, which cannot be compressed; institutional staking, in which a party with something durable to lose puts it at risk on your behalf; and closed networks of mutual acquaintance, which are costly to enter precisely because they are small. Every one of these is a mechanism from the pre-digital repertoire. The theory therefore predicts, not a collapse into anomie, but a *return of embeddedness in its classical form* at the top of the system, coexisting with a fully attested and increasingly uninformative layer below it. Introductions will matter more, not less. Rooms will matter more, not less. And the sociologist will note, with the melancholy appropriate to the observation, that a technology adopted in the name of open access will have ended by re-establishing the salon.

Proposition 21 (Artificial alters). *Where machine actors occupy positions in the enacted graph—as interlocutors, advisers, or intermediaries—they alter structure not by adding nodes but by absorbing the functions for which human ties were previously required.* If an actor's need for advice, information, emotional regulation, and routine sociability can be met by a non-human alter, then the ties that met those needs are not replaced but simply not formed, and the effect on the graph is subtractive, dispersed, and nearly impossible to observe—one cannot sample from the population of ties that failed to come into existence. I raise the point mainly to note that it is the one place where the framework predicts an effect that is in principle unobservable in attested data by construction, and where I therefore think ethnography and long-panel survey work are not a supplement to computational methods but the only available instrument. The same difficulty attends every attempt to identify social influence from observational network data (Christakis and Fowler 2007; Centola 2010), with the added indignity that here the counterfactual tie is not merely unidentified but nonexistent. ## 10. Implications Beyond Sociology

I have never thought sociology's proper ambition was to describe a residual domain left over after economics and psychology have taken their share. The interesting propositions in our field are propositions about how relations among units generate outcomes, and that problem is not the

property of any discipline. I therefore want to set out, briefly and without pretending to competence in each, what the framework implies elsewhere.

Economics

Information economics has been built on asymmetry between the parties to a transaction—the seller knows the car, the worker knows her ability, the insured knows his health. The framework here describes a different and, I would argue, now more consequential asymmetry: between the parties jointly and a third party who observes both and trades on the observation. The attestor is not an intermediary in the classical sense, because it does not bear inventory or risk; it holds the *relational* information from which the value of a match is derived. Standard results do not straightforwardly transfer. Consider the market for matching: a monopolist attestor with complete relational information can extract essentially the whole surplus from a match, since it can price at each party's reservation value, which it knows better than the party does. Consider signaling: Spence's separating equilibrium requires that the cost of the signal differ by type, and attestation fraud makes the cost differential collapse. Consider transaction cost economics, whose treatment of governance structures as efficient adaptations I criticized at length in 1985: the current literature explaining platforms as efficient solutions to matching and trust problems repeats the same functionalist error, inferring the cause of an institution from the benefits it happens to confer, and ignoring the distributional struggle by which the institution's form was actually determined. Whether an attestor exists, what it records, and who may act on the record are outcomes of conflict, not of optimization, and any account that omits the conflict will systematically mistake the winner's interest for the system's function.

Computer science and machine learning

Nearly all of graph learning proceeds by treating an observed graph as the object of study. The framework says the observed graph is an attested graph produced by a social process with the properties in Propositions 1–4, and therefore that: (a) missingness is not at random but correlates with the sensitive attributes practitioners are asked to protect; (b) link prediction trained on attested edges learns the *attestation function* at least as much as the social process, so that a model which perfectly predicts new edges may have learned nothing about friendship and everything about disclosure; (c) recommender systems that close triads accelerate Proposition 6, degrading the informational value of the network they are optimizing; and (d) adversarial robustness cannot be achieved by better detection, for the reason given in Proposition 15. Point (a) is the structural root of what the legal literature discusses as disparate impact without an intent requirement (Barocas and Selbst 2016), and it is not remediable by fairness constraints imposed on the model, because the distortion is in the world the model is shown rather than in the model's treatment of it (see also Pasquale 2015). I would add a design implication that follows directly and that I have not seen drawn: since the value of a recommendation to the recipient depends on the recommendation's exclusivity, a recommender that serves the same non-redundant information to everyone destroys the value it is measuring. Recommender evaluation metrics are, in this precise sense, self-defeating, and one might reasonably design systems that *deliberately withhold* coverage in order to preserve informational differentials.

Epidemiology and public health

Contact networks are attested networks. During the recent pandemic, digital contact tracing was deployed on the assumption that attested proximity approximates infectious contact. Propositions 1 through 3 predict specifically who will be missed: those whose contacts are unattested by choice (the affluent, the illicit) and those unattested by exclusion (the undocumented, the unbanked, the phoneless). Since these groups differ systematically in exposure, the resulting estimates of network structure are biased in the region of the parameter space that matters most for control. The general lesson for any field using digital traces to reconstruct interaction—epidemiology, transportation, criminology—is that the sampling frame is a social product with a determinate shape, and the shape can be modeled rather than lamented. It is worth recalling that the most instructive study of a sexual contact network was conducted by asking adolescents, in person, whom they had been involved with (Bearman, Moody, and Stovel 2004), and that the resulting structure—a long, sparse, spanning chain—is precisely the sort of structure an attestation-based instrument would have destroyed.

Political science

Mobilization theory has long held that recruitment runs through networks. If attested networks diverge from enacted ones, then movements recruit differently depending on which network organizers and observers can see. States and parties act on the attested graph, which means repression, targeting, and canvassing all operate on a representation whose bias is knowable. The theory predicts that successful clandestine mobilization will increasingly resemble pre-digital forms—cell structure, physical meeting, deliberate non-attestation—not from tactical nostalgia but because the opacity premium is very large where the third party is a state. It also suggests caution about the vast literature inferring polarization from platform data: Propositions 4 and 16 imply that a measured increase in the homophily of attested ties is fully consistent with no change in the homophily of enacted ties, if the payoff to attesting politically expressive relations has risen. I do not assert that polarization is an artifact. I assert that the dominant measurement strategy cannot distinguish the two, and that the theory says exactly what evidence would.

Law and regulation

Regulation of digital systems has focused on data protection—rights over records held about individuals. The framework suggests this misses the object. The consequential entity is not the record about a person but the record about a *relation*, which by construction concerns at least two people with divergent interests and which no individual-rights framework can allocate. When I disclose my relation to you, I disclose you. Consent, the central device of data protection law, is structurally incapable of handling this, since the disclosing party is not the exposed party. A relational rather than individual conception of informational interests seems to me the necessary direction, and I note that the older law of confidences, which attached duties to relationships rather than to data, was closer to the right structure than anything since.

Signaling theory and biology

The account in Section 9 is a straightforward extension of the handicap principle (Zahavi 1975): the informativeness of a display depends on its cost, and cheap displays are ignored. What the social case adds is that here the *receiver* is often an automated classifier with a fixed and learnable decision rule, which is a much weaker receiver than natural selection typically

produces. Biologists have models of signaling arms races in which receivers coevolve; the interesting feature of attested social systems is that receivers update slowly, in discrete releases, by decisions of an organization with its own interests. The dynamics of deception under a slow, corporate receiver are, so far as I know, unmodeled, and would repay attention from anyone equipped to do it.

The philosophy and practice of social measurement

Finally, and closest to home, the framework is an argument about our own instruments. Sociologists using platform data are attestors of a particular kind, subject to the same properties. Every result derived from trace data is a result about the attested graph, and its transfer to the enacted graph requires a model of the attestation function. We do not have such models. We could build them: the requirement is paired observation of the same population by attested and unattested means, which is expensive, unglamorous, and exactly what the field has stopped doing as trace data became abundant. Salganik (2018) has set out with admirable clarity how designed and found data might be combined; my point is that in the network case the combination is not an improvement but a precondition, since the found data are a representation produced by an interested party. I regard the substitution of available data for designed data as the most serious methodological error of the past two decades, and I say this as someone whose own most-cited result rests on fifty-four interviews.

11. Objections

"This is Scott's legibility, or Zuboff's surveillance capitalism, in network dress." The affinities are real and I have signposted them (Scott 1998; Zuboff 2019; and behind both, Foucault). But those are accounts of *power over subjects*: a center makes a population knowable in order to govern it. My argument concerns *the structure of relations*, not the condition of individuals, and it generates propositions those accounts cannot. Legibility theory does not predict that brokerage returns decline with intermediary coverage; it does not predict that the premium migrates to unattested ties; it does not predict discontinuous reallocation from threshold effects; and it does not predict that malfeasance relocates to the attestation layer. It also lacks a mechanism for divergence, since it tends to assume the state's map is simply cruder than the terrain, whereas my claim is that the map is *strategically produced by the mapped*, which is a different and more interesting situation.

"Attestation is not new. Parish registers, guild rolls, credit ledgers, and the confessional all attested relations." Entirely correct, and I would resist any argument that requires a sharp break. The framework is meant to be historical and comparative: attestation regimes vary in coverage, granularity, latency, integration, and reciprocity, and the propositions are stated as comparative statics in those variables, not as claims about a digital rupture. The Florentine catasto was an attestation regime, and Padgett and Ansell's (1993) Medici are a case of an elite exploiting the gap between the attested and enacted graphs—which is precisely the opacity premium, four centuries early. What is distinctive now is the conjunction of near-total coverage, near-zero latency, cross-domain integration, and total non-reciprocity, and it is the conjunction, not any single element, that does the theoretical work.

"You have exaggerated the divergence between the graphs; people mostly connect online with people they know." In the aggregate, yes, and I do not dispute the correlational finding.

But my claim is not that the graphs are unrelated; it is that they diverge *systematically*, and the propositions specify where. A correlation of 0.7 between attested and enacted ties is entirely compatible with the residual 0.3 being exactly the ties that carry the consequential information—indeed Proposition 7 says it must be. Averages are the wrong statistic here, as they usually are when structure is the object.

"The theory is unfalsifiable, since the crucial variable is by definition unobserved." This objection has force, and I want to meet it rather than deflect it. The unattested tie is unobservable *to the attestor*, not to the sociologist, who has other instruments: interviews, observation, diaries, archives, records held by attestors with different coverage. The whole empirical strategy of the framework is triangulation between attestors with non-overlapping blind spots, which is a well-understood design. Section 12 sets out specific tests, several of which would embarrass the theory badly if they came out wrong.

"You have made the platform an actor, which is reification." The charge would be fair if I had attributed intentions to software. I have attributed to the attestor only what I would attribute to a firm or an agency: a position in a structure, a set of interests, a capacity to act, and a record. Whether the attestor's conduct issues from a boardroom, a policy document, or an objective function seems to me immaterial to its structural consequences, and I would note that sociology has always been willing to treat organizations as actors when they behave like them.

"If the bridging premium is gone, why do networks still so evidently pay?" Because Proposition 5 is domain-specific and the domains where networks visibly pay are precisely the domains where third-party coverage is worst: elite hiring, venture finance, political appointment, deal origination, and academic recognition. The theory does not predict that networks stop paying. It predicts that the *kind* of network that pays changes, and that the change runs against equality. If a well-designed study found brokerage returns declining uniformly across domains regardless of intermediary coverage, or found them declining fastest in elite settings, the theory would be in serious trouble.

12. An Empirical Agenda

A theory that cannot be embarrassed is not worth having. I list the tests I would most like to see, in rough order of feasibility.

1. **Paired-instrument studies.** Administer a conventional name generator to a population for which platform data are available, and estimate α as a function of relation type, tie strength, alter status, and ego status. Propositions 1–3 make directional predictions about every one of these coefficients. If α turns out to be roughly constant, the framework is wrong at its foundation.

2. **Coverage-exploiting comparative designs.** Brokerage returns should decline with intermediary coverage. Coverage varies across occupations, countries, and time in ways that are plausibly exogenous to network structure—the staggered entry of matching platforms into national labor markets is close to a natural experiment. Proposition 5 predicts a negative interaction between coverage and the return to structural holes; a null or positive interaction would be strong disconfirmation.

3. **Regime-change event studies.** Proposition 17 predicts jumps in resource allocation timed to changes in attestation regimes—a platform default change, a registry's opening, a disclosure mandate—concentrated among actors near the attestation margin, and *absent* corresponding change in independently measured conduct. This is the sharpest available test of the decoupling claim, and the required designs are standard.

4. **Withdrawal studies.** Proposition 18 predicts hysteresis. Populations induced to reduce attestation should show slower and smaller reversion than the original adoption, and should incur measurable suspicion costs. Existing platform-deactivation experiments could be re-analyzed for this asymmetry.

5. **Detection-sophistication studies.** Proposition 15 predicts that the structural sophistication of detected counterfeit subgraphs increases over time while detected volume falls. Platform trust-and-safety archives contain the data; the field would benefit enormously from access to them, and I am aware of how unlikely that is.

6. **Ethnography of the unattested.** The opacity premium is a claim about what happens in rooms with no logs, and the only instrument for rooms with no logs is a person in the room. That the framework's central proposition can be tested only by the oldest method in the discipline is a result I find satisfying, and I would encourage anyone tempted to treat computational and qualitative methods as rivals to notice that here they are strict complements: the computational work identifies where the gap should be, and only the fieldwork can look into it.

13. Conclusion

I have argued that the elementary unit of digitally mediated social life is not the dyad but the triad, and that the third member of that triad is a party of a sort no previous social structure contained: one that observes without being observed, remembers without forgetting, participates without reciprocating, and acts on its record with consequences the parties cannot escape. From this I have derived a bifurcation of social structure into enacted and attested graphs; an economy of visibility that governs the mapping between them; the erosion of the bridging premium and its migration to illegible ties; the expropriation of the structural component of trust and the relocation of malfeasance to the attestation layer; the discontinuous aggregation of attestation decisions and the resulting decoupling of resource flows from social reality; and the collapse of attestation as a signal once its production becomes costless.

If I am right, several things follow that are worth stating plainly. The first is that the most consequential social structure of our period is, by construction, the part we cannot see, and that this is not a temporary limitation of instruments but a stable equilibrium property of the system—as measurement improves, the valuable relations move. The second is that the apparent triumph of computational social science may be substantially a triumph of measuring the measurable, and that the confidence of our findings is growing faster than their validity, for reasons the framework specifies. The third, and the one I would emphasize to anyone concerned with inequality, is that a technology universally described as connective and democratizing has, on this analysis, produced a new and unusually well-hidden form of advantage: the capacity to have relations that no one records. It is an old form of advantage in a new register. Those who could formerly afford walls can now afford silence.

I will end where I have ended before. Sociology's contribution is not to insist that social relations matter, which everyone concedes and no one acts on, but to specify how relations of particular kinds, in particular configurations, produce outcomes that no one intended and that no summary of individual dispositions could predict. The apparatus we have built to observe ourselves is now among those configurations. It is a peculiar situation for a discipline to be in, and it seems to me the most interesting problem we have.

References

Akerlof, George A. 1970. "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism." *Quarterly Journal of Economics* 84(3):488–500.

Albert, Réka, Hawoong Jeong, and Albert-László Barabási. 2000. "Error and Attack Tolerance of Complex Networks." *Nature* 406:378–382.

Aral, Sinan, and Marshall Van Alstyne. 2011. "The Diversity–Bandwidth Trade-off." *American Journal of Sociology* 117(1):90–171.

Bail, Christopher A. 2021. *Breaking the Social Media Prism: How to Make Our Platforms Less Polarizing*. Princeton, NJ: Princeton University Press.

Bail, Christopher A., et al. 2018. "Exposure to Opposing Views on Social Media Can Increase Political Polarization." *Proceedings of the National Academy of Sciences* 115(37):9216–9221.

Bakshy, Eytan, Solomon Messing, and Lada A. Adamic. 2015. "Exposure to Ideologically Diverse News and Opinion on Facebook." *Science* 348(6239):1130–1132.

Barabási, Albert-László, and Réka Albert. 1999. "Emergence of Scaling in Random Networks." *Science* 286(5439):509–512.

Barocas, Solon, and Andrew D. Selbst. 2016. "Big Data's Disparate Impact." *California Law Review* 104(3):671–732.

Bearman, Peter S., James Moody, and Katherine Stovel. 2004. "Chains of Affection: The Structure of Adolescent Romantic and Sexual Networks." *American Journal of Sociology* 110(1):44–91.

Bourdieu, Pierre. 1986. "The Forms of Capital." Pp. 241–258 in *Handbook of Theory and Research for the Sociology of Education*, edited by J. G. Richardson. New York: Greenwood.

Bowker, Geoffrey C., and Susan Leigh Star. 1999. *Sorting Things Out: Classification and Its Consequences*. Cambridge, MA: MIT Press.

Burt, Ronald S. 1992. *Structural Holes: The Social Structure of Competition*. Cambridge, MA: Harvard University Press.

Burt, Ronald S. 2002. "Bridge Decay." *Social Networks* 24(4):333–363.

Burt, Ronald S. 2004. "Structural Holes and Good Ideas." *American Journal of Sociology* 110(2):349–399.

Campbell, Donald T. 1979. “Assessing the Impact of Planned Social Change.” *Evaluation and Program Planning* 2(1):67–90.

Centola, Damon. 2010. “The Spread of Behavior in an Online Social Network Experiment.” *Science* 329(5996):1194–1197.

Chetty, Raj, et al. 2022. “Social Capital I: Measurement and Associations with Economic Mobility.” *Nature* 608:108–121.

Christakis, Nicholas A., and James H. Fowler. 2007. “The Spread of Obesity in a Large Social Network over 32 Years.” *New England Journal of Medicine* 357(4):370–379.

Coleman, James S. 1988. “Social Capital in the Creation of Human Capital.” *American Journal of Sociology* 94(Supplement):S95–S120.

Desrosières, Alain. 1998. *The Politics of Large Numbers: A History of Statistical Reasoning*. Cambridge, MA: Harvard University Press.

Douceur, John R. 2002. “The Sybil Attack.” Pp. 251–260 in *Peer-to-Peer Systems (IPTPS 2002)*. Berlin: Springer.

Emerson, Richard M. 1962. “Power-Dependence Relations.” *American Sociological Review* 27(1):31–41.

Emirbayer, Mustafa. 1997. “Manifesto for a Relational Sociology.” *American Journal of Sociology* 103(2):281–317.

Erikson, Emily. 2013. “Formalist and Relationalist Theory in Social Network Analysis.” *Sociological Theory* 31(3):219–242.

Espeland, Wendy Nelson, and Michael Sauder. 2007. “Rankings and Reactivity: How Public Measures Recreate Social Worlds.” *American Journal of Sociology* 113(1):1–40.

Feld, Scott L. 1981. “The Focused Organization of Social Ties.” *American Journal of Sociology* 86(5):1015–1035.

Fourcade, Marion, and Kieran Healy. 2013. “Classification Situations: Life-Chances in the Neoliberal Era.” *Accounting, Organizations and Society* 38(8):559–572.

Fourcade, Marion, and Kieran Healy. 2017. “Seeing Like a Market.” *Socio-Economic Review* 15(1):9–29.

Gambetta, Diego. 1993. *The Sicilian Mafia: The Business of Private Protection*. Cambridge, MA: Harvard University Press.

Gambetta, Diego. 2009. *Codes of the Underworld: How Criminals Communicate*. Princeton, NJ: Princeton University Press.

Gillespie, Tarleton. 2014. “The Relevance of Algorithms.” Pp. 167–194 in *Media Technologies*, edited by T. Gillespie, P. Boczkowski, and K. Foot. Cambridge, MA: MIT Press.

Goffman, Erving. 1959. *The Presentation of Self in Everyday Life*. New York: Doubleday.

Granovetter, Mark S. 1973. “The Strength of Weak Ties.” *American Journal of Sociology* 78(6):1360–1380.

Granovetter, Mark S. 1974. *Getting a Job: A Study of Contacts and Careers*. Cambridge, MA: Harvard University Press.

Granovetter, Mark S. 1978. “Threshold Models of Collective Behavior.” *American Journal of Sociology* 83(6):1420–1443.

Granovetter, Mark S. 1985. “Economic Action and Social Structure: The Problem of Embeddedness.” *American Journal of Sociology* 91(3):481–510.

Granovetter, Mark S. 2005. “The Impact of Social Structure on Economic Outcomes.” *Journal of Economic Perspectives* 19(1):33–50.

Granovetter, Mark S. 2017. *Society and Economy: Framework and Principles*. Cambridge, MA: Harvard University Press.

Granovetter, Mark, and Roland Soong. 1983. “Threshold Models of Diffusion and Collective Behavior.” *Journal of Mathematical Sociology* 9(3):165–179.

Granovetter, Mark, and Roland Soong. 1986. “Threshold Models of Interpersonal Effects in Consumer Demand.” *Journal of Economic Behavior and Organization* 7(1):83–99.

Hacking, Ian. 1986. “Making Up People.” Pp. 222–236 in *Reconstructing Individualism*, edited by T. Heller, M. Sosna, and D. Wellbery. Stanford, CA: Stanford University Press.

Healy, Kieran. 2015. “The Performativity of Networks.” *European Journal of Sociology* 56(2):175–205.

Heider, Fritz. 1958. *The Psychology of Interpersonal Relations*. New York: Wiley.

Kossinets, Gueorgi, and Duncan J. Watts. 2006. “Empirical Analysis of an Evolving Social Network.” *Science* 311(5757):88–90.

Lazer, David, Ryan Kennedy, Gary King, and Alessandro Vespignani. 2014. “The Parable of Google Flu: Traps in Big Data Analysis.” *Science* 343(6176):1203–1205.

Lessig, Lawrence. 1999. *Code and Other Laws of Cyberspace*. New York: Basic Books.

Lorrain, François, and Harrison C. White. 1971. “Structural Equivalence of Individuals in Social Networks.” *Journal of Mathematical Sociology* 1(1):49–80.

MacKenzie, Donald. 2006. *An Engine, Not a Camera: How Financial Models Shape Markets*. Cambridge, MA: MIT Press.

Marsden, Peter V. 1990. “Network Data and Measurement.” *Annual Review of Sociology* 16:435–463.

Marwick, Alice E., and danah boyd. 2011. “I Tweet Honestly, I Tweet Passionately: Twitter Users, Context Collapse, and the Imagined Audience.” *New Media and Society* 13(1):114–133.

McPherson, Miller, Lynn Smith-Lovin, and James M. Cook. 2001. “Birds of a Feather: Homophily in Social Networks.” *Annual Review of Sociology* 27:415–444.

Merton, Robert K. 1948. “The Self-Fulfilling Prophecy.” *Antioch Review* 8(2):193–210.

Muller, Jerry Z. 2018. *The Tyranny of Metrics*. Princeton, NJ: Princeton University Press.

Padgett, John F., and Christopher K. Ansell. 1993. “Robust Action and the Rise of the Medici, 1400–1434.” *American Journal of Sociology* 98(6):1259–1319.

Pasquale, Frank. 2015. *The Black Box Society: The Secret Algorithms That Control Money and Information*. Cambridge, MA: Harvard University Press.

Podolny, Joel M. 2005. *Status Signals: A Sociological Study of Market Competition*. Princeton, NJ: Princeton University Press.

Porter, Theodore M. 1995. *Trust in Numbers: The Pursuit of Objectivity in Science and Public Life*. Princeton, NJ: Princeton University Press.

Power, Michael. 1997. *The Audit Society: Rituals of Verification*. Oxford: Oxford University Press.

Putnam, Robert D. 2000. *Bowling Alone: The Collapse and Revival of American Community*. New York: Simon and Schuster.

Rajkumar, Karthik, Guillaume Saint-Jacques, Iavor Bojinov, Erik Brynjolfsson, and Sinan Aral. 2022. “A Causal Test of the Strength of Weak Ties.” *Science* 377(6612):1304–1310.

Ruths, Derek, and Jürgen Pfeffer. 2014. “Social Media for Large Studies of Behavior.” *Science* 346(6213):1063–1064.

Salganik, Matthew J. 2018. *Bit by Bit: Social Research in the Digital Age*. Princeton, NJ: Princeton University Press.

Salganik, Matthew J., Peter Sheridan Dodds, and Duncan J. Watts. 2006. “Experimental Study of Inequality and Unpredictability in an Artificial Cultural Market.” *Science* 311(5762):854–856.

Schelling, Thomas C. 1978. *Micromotives and Macrobehavior*. New York: Norton.

Scott, James C. 1998. *Seeing Like a State: How Certain Schemes to Improve the Human Condition Have Failed*. New Haven, CT: Yale University Press.

Simmel, Georg. [1908] 1950. *The Sociology of Georg Simmel*, translated and edited by Kurt H. Wolff. Glencoe, IL: Free Press.

Spence, Michael. 1973. “Job Market Signaling.” *Quarterly Journal of Economics* 87(3):355–374.

Star, Susan Leigh. 1999. “The Ethnography of Infrastructure.” *American Behavioral Scientist* 43(3):377–391.

Strathern, Marilyn. 1997. “‘Improving Ratings’: Audit in the British University System.” *European Review* 5(3):305–321.

Tufekci, Zeynep. 2014. “Big Questions for Social Media Big Data: Representativeness, Validity and Other Methodological Pitfalls.” *Proceedings of the International AAAI Conference on Web and Social Media* 8:505–514.

Uzzi, Brian. 1997. “Social Structure and Competition in Interfirm Networks: The Paradox of Embeddedness.” *Administrative Science Quarterly* 42(1):35–67.

Watts, Duncan J., and Steven H. Strogatz. 1998. “Collective Dynamics of ‘Small-World’ Networks.” *Nature* 393(6684):440–442.

White, Harrison C. 1981. “Where Do Markets Come From?” *American Journal of Sociology* 87(3):517–547.

Wrong, Dennis H. 1961. “The Oversocialized Conception of Man in Modern Sociology.” *American Sociological Review* 26(2):183–193.

Zahavi, Amotz. 1975. “Mate Selection—A Selection for a Handicap.” *Journal of Theoretical Biology* 53(1):205–214.

Zelizer, Viviana A. 2005. *The Purchase of Intimacy*. Princeton, NJ: Princeton University Press.

Zuboff, Shoshana. 2019. *The Age of Surveillance Capitalism*. New York: PublicAffairs.

Zuckerman, Ezra W. 1999. “The Categorical Imperative: Securities Analysts and the Illegitimacy Discount.” *American Journal of Sociology* 104(5):1398–1438.