When Dissent is Good: The Interplay Between Reputation and Social Networks in an Artificial Society

Elena Zinchenko, Harvard University
Daniel Joshua Steinbock, University of California, Santa Cruz
Joshua Miller, NuTech Solutions, Inc.
Reuben J. Thomas, Stanford University

ABSTRACT

This paper presents an artificial society wherein heterogeneous agents interact to form opinions about each other, and the effect of these opinions on agent behavior grows a social network. In our computational model, an opinion is a numeric score ranging on a continuum from positive to negative esteem and determines whom an agent prefers to interact with. Each agent is a node on a social network that links agents who have interacted. Each agent maintains a private list of opinions—one for every agent it has met. The list delineates which nodes an agent is adjacent to on the social network while the opinions define the connection strengths between nodes. In addition, agents "gossip" by communicating the highest scores they know of to adjacent nodes. An agent's opinion is influential in proportion to how highly the agent is regarded. As agents are influenced, and gossip passed on, an individual's reputation can reach distant agents without direct interaction. Thus reputation spreads across the self-organized social network, potentially attracting distant agents into forming direct connections, thereby transforming the network's topology. By varying how objective or subjective the agents are when evaluating each other, several typical topologies with associated dynamics emerge. The primary finding is that pure objectivity leads to multiple isolated networks whereas partial subjectivity (when differences of opinion become possible) rapidly converges to a globally connected network.

Introduction

Status and reputations are ubiquitous in human social settings: whenever people interact, they form opinions of each other, and networks of esteem emerge if none already existed (for a review on status behaviors see Ridgeway & Walker 1995 and Berger et al 1980). Where all judgements of worth reflect actual performances (judged by some shared standard, of course), a meritocratic hierarchy can be said to exist. Yet in most cases, the esteem one holds others in is at least partly a function of other factors: what social groups and categories they seem to belong to, their similarity to oneself in culture and values, and the esteem other people seem to hold them in.

The last of these is an influence process operating through social networks, but unlike most such processes, network structure and dynamics are endogenous. Whom one is influenced by is affected by whom one holds in high esteem, which in turn is affected by whom one is influenced by, and so on indefinitely. The implications of this interplay between social structure and social psychological behavior have received insufficient attention thus far, given the potential for profound and general insight into social organization (for two exceptions, see Skvoretz & Fararo 1986 and Gould 2001).

Towards this end, we offer here a simple model that illustrates the promise of growing social networks from simple simulated status processes. The model is designed to be a representation of social interaction and the emergence of stable and consistent reputation across a multitude of individually biased opinions.

Though initially intended as a representation of the academic enterprise, the model's assumptions (see below) are generic enough to allow investigation of any number of social systems where agents form impressions of each other, selectively share these impressions with others, and seek to interact with highly regarded agents.

Model Description

Agent Specification

Model actors, known as "agents", operate independently, moving around a 2D continuous non-periodic landscape with a locally limited perception of the world (vision). The following rules define the behavior each agent performs at every time step.

Interaction Rule I: Choose a random neighbor agent within my vision and assign her an opinion score with function **S** (see below). Store this score in my record of opinions. If she is among the top-ten scoring agents I've ever met, add her to my list of "friends" (displacing the lowest agent from the list).

Communication Rule C: Query my friends for a list of high-scoring agents they have heard of and incorporate their opinions into my record of opinions. In turn, share my top scorers with whoever asks.

Movement Rule M:

- i. Move one step towards the highest scoring agent I know of with probability p.
- ii. Take a step in a random direction with probability (1 – p).
- iii. Move one step away from an agent whom I've assigned a negative reputation with probability p.

Agents keep a record of scores for everybody they have ever met (via Rule **I**) or heard about via gossip (Rule **C**). We define the model's social network as linking each agent to the top-ten scoring agents they have met, referring to these as the agent's "friends". Due to differences in social interaction history, friendships are not necessarily mutual. As initial conditions, agents have no opinions, no network connections and no friends. An agent walks randomly until it has an interaction, after which its movement is biased according to Movement Rule **M**.

The agent population is heterogeneous with each agent having two endogenous attributes, one subjective (referred to as "Culture") and the other objective ("Skill"). The terms "subjective" and "objective", though controversial within the social sciences, are defined mathematically in this model by the scoring function **S** (see below).

Subjective Culture is uniformly distributed across all the agents between 1 and 100. Objective Skill is normally distributed, with a mean of 50 and a standard deviation of 9. We chose these values as they appeared to maximize the effect of the global parameter α (see below). Vision is homogenous across all agents.

The model has three exogenous global parameters: a weight, $\boldsymbol{\alpha}$, a tolerance, $\boldsymbol{\tau}$, and a movement probability, **p**. Parameters α and τ are used to calculate agent *w*'s opinion of agent *x* (*w* != *x*), according to function **S** below, and p is used in the movement rule **M**. Additionally, $\boldsymbol{\varepsilon_s}$ and $\boldsymbol{\varepsilon_c}$ introduce a small amount of noise into the value of Skill and Culture respectively.

Opinion Scoring Function S:

$$S(x) = \alpha * (skill_x + \varepsilon_s) + (1 - \alpha) * [(\tau + \varepsilon_c) - |culture_w - culture_x|]$$

The weight, α, which we label "Consensus", ranges from 0.0 to 1.0 and modulates the relative importance of the subjective trait versus the objective one in determining agent *w*'s opinion of *x*. Higher α makes agents more objective in their judgments by emphasizing the shared standard of evaluation, Skill. Skill is "objective" because it is measured on an absolute scale where higher values always score better. Culture is "subjective" because it is measured relative to agent *w*'s own Culture where smaller differences score better. When α is a fraction, both Skill and Culture contribute to the score. The tolerance, τ, determines the range of Culture values within which agent *w* evaluates other agents favorably (when Culture's contribution is non-negative). The contribution of Skill is always positive.

By the rules of our model, each agent has the explicit goal of moving toward the most appealing (highest scoring) agent it knows of. Implicitly, each is searching a score landscape for the global maximum. Due to cultural heterogeneity, different agents may perceive different score landscapes. We label α Consensus because it determines the average degree of similarity between agents' perceived landscape. If Consensus is 1.0 (pure Objectivity), then all agents are searching the exact same landscape and will always agree when scoring agent *x* (aside from ε noise). If Consensus is 0.0 (pure Subjectivity), each agent is searching a unique score landscape. Only agents with the same Culture will agree when judging another agent. In between the two extremes, agents' perceptions of the world are partially overlapping. Skill and culture combine so that a close cultural similarity can offset a low skill level, or a cultural schism can be overcome by high skill.

Assumptions of Limited Information

Implicit in the model are two simplifying assumptions. Firstly, Interaction Rule **I** assumes each agent has limit perception. If the homogenous vision parameter was too large (i.e. equal to the size of the space) there would be no sense of local interactions.

A second limiting assumption is the size of an agent's social network. Agents are allowed to have as maximum of 10. At any time t, agent *w*'s social network consists of the

10 highest-ranked with whom agent *w* has interacted (assuming *w* has interacted with at least 10 other agents).

Results

While the Consensus (α) and Tolerance (τ) parameters, as well as the movement probability (p), can all be varied, only the effects of Consensus are examined here. Due to time constraints, statistically significant results from large samples of simulated results were not obtained. Instead, the results offered below are only measures of individual network outcomes that are representative of the most likely outcomes for each α examined. As no rigorous method was used for determining this sense of representativeness, this analysis should be regarded as only exploratory, and not at all definitive.

Network topology varies considerably across values of α. Network density (Wasserman & Faust 1994) increases linearly with α, but its effects on the clustering and cohesiveness of the network show more complicated relationships (see fig. 1). Distance-based cohesion (Doreian 1974) peaks at α's of .2 and .7, with minimal values at the extremes, as its opposite measure, fragmentation, behaves conversely. The clustering coefficient (Watts 1999) peaks at .9, but has its minimum at .3 α (a very cohesive but sparsely-clustered network).

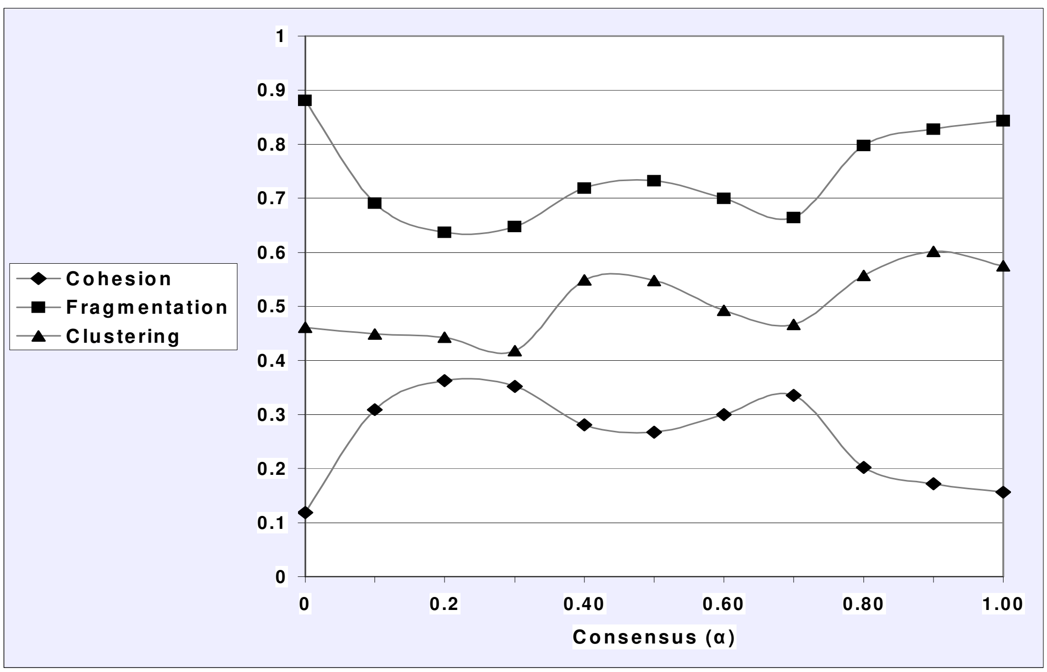

Fig. 1. Social network metrics as parameter α varies

The number of outgoing ties (out-degree) in these networks is no greater than ten for any agent, by the rules of the model, but in-degree is only limited by the size of the population. As in-degree represents the number of agents that count the agent in question as one of their ten highest esteemed colleagues (and are thus directly influenced by them), this is a good measure of both prestige and influence. When α is low, the resulting distribution of in-degree is approximately normal, if somewhat skewed to the low end. This skewing only increases with α, looking less and less Gaussian and more like a power law distribution (suggesting scale-free networks).

It is also interesting to note the qualitative visual variation in network topology. In the case of pure subjectivity ($\alpha = 0.0$), when Culture is the sole orient for forming an opinion about another agent, a very sparse, linear network emerges (fig. 2). Consensus is minimal, so agents only form connections with those of similar Culture, of which there are few.

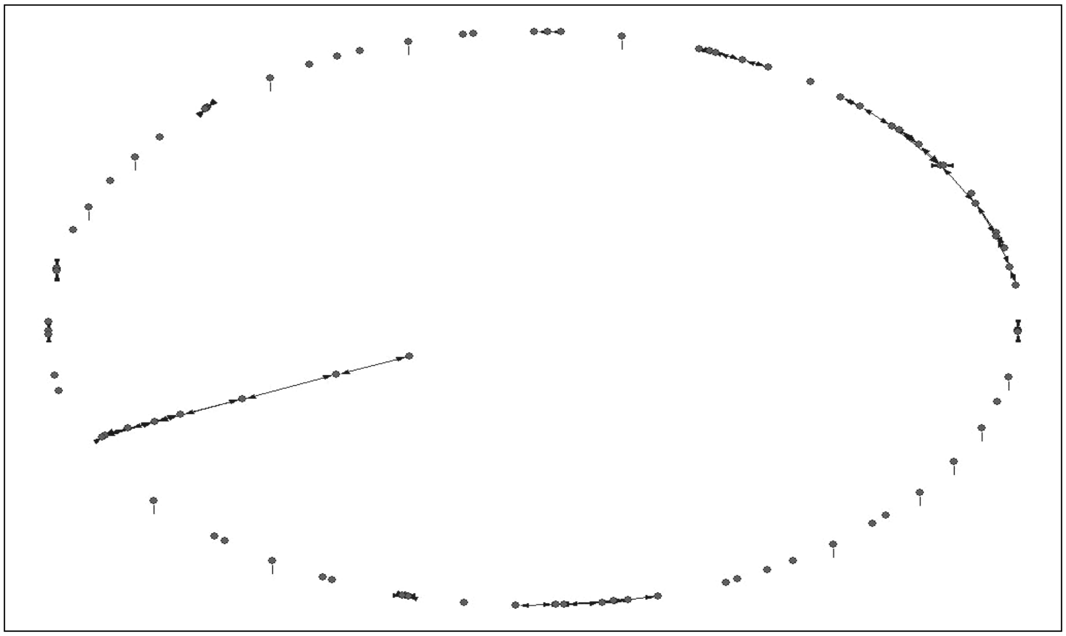

Fig. 2. Pure Subjectivity ($\alpha = 0.0$)

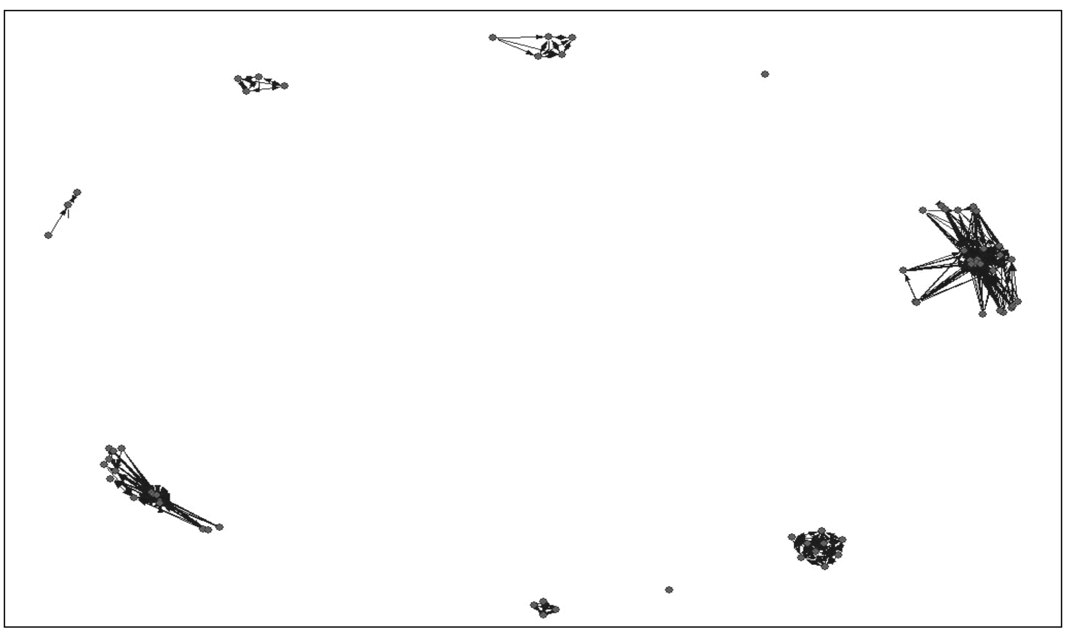

Fig. 3. Pure Objectivity ($\alpha = 1.0$)

When agents are purely objective ($\alpha = 1.0$), Consensus is maximal, but they quickly condense into many isolated networks, each strictly hierarchical according to Skill (fig. 3). Agent clusters with homogeneous Skill appear egalitarian while heterogeneity is indicated by a stratification of hierarchical levels. If two egalitarian groups merge, a two-level hierarchy is the likely result (assuming the groups' mean Skill levels differ significantly). The isolation of subnetworks is a consequence of agents getting caught in local Skill maxima. This issue is discussed in detail in the next section when we consider the social world of scientists, a system that explicitly values objectivity.

And finally, when objectivity and subjectivity are balanced ($\alpha = 0.5$), the result is rapid convergence to a globally connected network (fig. 4). The topology reflects a combination of linear Culture gradient (around the curved length of the figure) and hierarchical stellation according to Skill (centralized clusters).

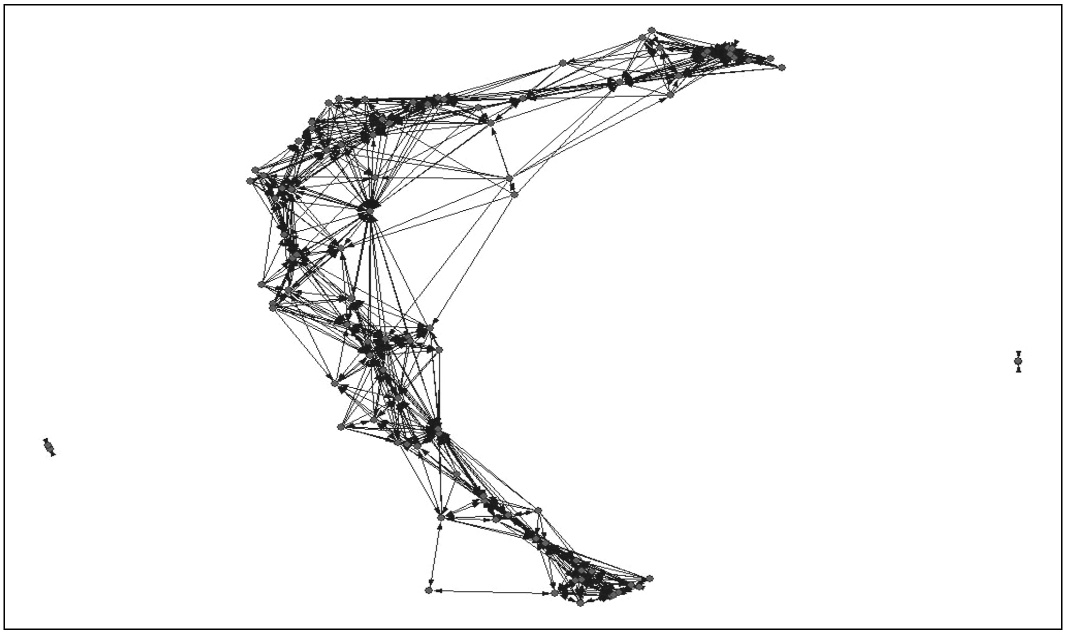

Fig. 4. Balanced Objectivity/Subjectivity ($\alpha = 0.5$)

Discussion

In human social interaction, we always form opinions about people we encounter. These opinions are mental models, constructed and refined over time out of the data of personal experience. They influence our behavior and attitudes: whom we associate with socially and professionally and whose opinion we trust.

When people we trust share their opinions with us, they alter our opinions of others. We in turn influence other people's opinions by passing on information about reputable individuals, even ones we've never interacted with directly. When many local personal opinions about individuals diffuse across the social network into the global population, the individuals obtain a status in the network (their reputation) that may be only loosely coupled with their actual performance and abilities. One's reputation cannot be found in any one place or person; it is an aggregate effect, an emergent property of processes of social estimation.

In our artificial society, status and reputation processes have simple graph-theoretic definitions. For instance, how influential an agent is can be measured by the number of incoming links to that agent's node on the network. Each link signifies that someone trusts that agent's opinion highly and the link strength signifies exactly how much trust is given. Because trust is never 100%, the opinions that influential agents communicate degrade as they traverse the network and so reputation can travel only so far. But a strong reputation will attract other agents to its source in the 2D world, creating new direct links between once-distant individuals and transforming the social network topology. Thus reputation spreads across a network of its own creation. As our results indicate, the topology that emerges out of this self-organization depends critically on Consensus (α).

Before drawing conclusions about actual human societies from the results of our simulation, it is appropriate to say a few words about the general role of artificial society models in social scientific inquiry (for more detailed analysis see Lansing 2000 or Conte 2002). While the physicist can rely on already-formulated physical laws to provide the rational assumptions on which a computational model is built, the social scientist is tasked with reducing lawless, semi-rational, social psychological behavior to mathematical formalism (e.g. Agar 2002). She creates an artificial society which simulates her own assumptions—precisely to the extent that she can represent them inside a computer. Such a model is useful if it aids the intuition in comprehending real societies and indicates promising research topics for the application traditional methods like ethnography.

In the social system of science, we tend to aspire to an idealized notion of pure objectivity when evaluating both our data and our colleagues; in practice, this is rarely, if ever, attained (Heylighen, 1997). A surprising intuition that can be drawn from the results of our abstract social simulation is that pure objectivity—when no subjective bias exists—may *not* be worth aspiring to. If agents always agree on who the best agent is, and have this consensus opinion continually reinforced by local neighbors, agents are too easily satisfied by the locally maximal "expert". In the absence of any disagreement, they have no motivation to migrate and explore other parts of the intellectual world. The result is a multiplicity of isolated, monocultural, strictly hierarchical communities. This is reflected in fig. 3 where we see a fragmented network. No information is being shared between these isolated clusters because their membership is static and each is unaware of the others' existence. At the other extreme, in the total absence of consensus, stable groups of any significant size cannot form owing to a lack of common ground (fig. 2)

Intuitively, it becomes obvious why dissenting opinions, differences in worldview and nonconformity all contribute to migration between communities. It is because such differences make possible dissatisfaction with the *status quo*. Dissatisfaction gives agents the impetus to escape local maxima of expertise and discover more competent, like-minded collaborators. This behavior very quickly establishes a globally connected social network (fig. 4), across which reputation and knowledge can potentially reach any node from any other node. The benefits of this to scientific investigation are obvious. When scientific communities don't communicate, as is the case with the traditional divisions between disciplines, time and effort are wasted on redundant research. When they do communicate, results are shared, duplication of effort is minimized, and the evolution of knowledge proceeds in a far more efficient and robust manner.

Conclusion

We have presented an artificial society that examines the effect of subjective versus objective evaluation on agents' opinion formation and the diffusion of reputation across an emergent network. The resultant networks exhibited topologies that demonstrated a consistent effect: fragmented connectivity when subjectivity or objectivity were taken to an extreme, and a globally connected network in the balanced opinion evaluation scenario.

References

Agar, Michael. 2002. "Drugmart: Heroin Epidemics as Complex Adaptive Systems". Complexity. 7,5:44-52.

Berger, Joseph, Susan Rosenholtz, & Morris Zelditch. 1980. "Status Organizing Processes." Annual Review of Sociology. 6:479-508.

Conte, Rosaria. 2002. "Agent-based modeling for understanding social intelligence." Proceedings of the National Academy of Sciences. 99,3:7189-7190.

Doreian, P. 1974. "On the Connectivity of Social Networks." Journal of Mathematical Sociology. 3:245-58.

Fararo, Thomas & John Skvoretz. 1986. "E-State Structuralism: A Theoretical Method." American Sociological Review. 51,5:591-602.

Gould, Roger. 2001. "The Origins of Status Hierarchies: A Formal Theory and Empirical Test." American Journal of Sociology. 107:1143-78.

Heylighen, F. 1997. "Objective, Subjective and Intersubjective Selectors of Knowledge." Evolution and Cognition. 3,1:63-67.

Lansing, Stephen J. 2002. "Artificial Societies and the Social Sciences." Artificial Life. 8,3:279-292.

Ridgeway, Cecilia L. & Henry A. Walker. 1995. "Status Structures." Sociological Perspectives on Social Psychology. edited by Karen Cook, Gary Alan Fine, & James House. Boston: Allyn & Bacon.

Wasserman, Stanley and Katherine Faust. 1994. Social Network Analysis: Methods and Applications. Cambridge, UK: Cambridge University Press.

Watts, Duncan. 1999. Small worlds. Princeton: Princeton University Press.