Favor-Trading in the Amazon for Intelligent People

Simon DeDeo*

February 11, 2015

This is a brief guide, for anthropologists with some quantitative experience, to our paper *Dynamical Structure of a Traditional Amazonian Social Network* [1], an article by me (Simon), Paul Hooper (an anthropologist at Emory), Ann Caldwell Hooper (Anthropology, Emory), Michael Gurven (Anthropology, UC Santa Barbara) and Hilly Kaplan (Anthropology, University of New Mexico), featuring real-world data from an indigenous community in Bolivia, and providing new insight into the origins and development of cooperation in human society. Paul and I were fellows together at the Santa Fe Institute when this paper began.

Cooperation, including altruistic exchange—where I provide a benefit to you at net short-term cost to myself—is supported, in human society, by various mechanisms that attempt to prevent cheating and to provide both group- and individual-level benefits [3]. One of the most basic mechanisms is reciprocity: you tend to cooperate with others who return the favor.

In the indigenous group studied in the paper, the Tsimane' of Bolivia, we focused on a single mode of exchange, chicha (Manoic beer). Chicha is a high-calorie alcoholic drink; it's brewed by families and shared with others in discrete social events we refer to as "chicha parties". As can be seen in the epigraph to our paper, chicha parties can be emotionally, spiritually, and socially satisfying events over and above meeting material needs. We're interested in tracking the ways in which invitations to these parties are reciprocated.

1 Two Modes of Reciprocity

We consider two forms of reciprocity: a hot mode, and a cool mode. Hot mode reciprocity is a form of tit-for-tat, where particular invitations are tracked and reciprocated—we refer to this a *conditional-action reciprocity*. Cool mode reciprocity involves the formation of stationary bonds, where individuals and families have preferred donors and recipients of favors.

Intermediate between hot and cool are the warm modes. In the warm mode families do not track individual events, but long-term averages, such that hosting of family A by family B leads to increased numbers of invitations by family B to family A. Hot and cool modes are the two extremes; we expect reality to fall somewhere in the middle, as you can verify in your own social world by introspection. If I have you over to a party, I will expect you to include me in your social

*Assistant Professor, Indiana University, School of Informatics and Computing & Cognitive Science Program; External Professor, Santa Fe Institute. sdedeo@indiana.edu; http://santafe.edu/~simon. Errors and mis-statements are the author's alone, and not the responsibility of co-authors. Part of the for Intelligent People series.

circle to a greater extent, and to receive invitations from you; but I will not consider that invitation an outstanding debt that must be paid, and I will not expect you to do so either. Indeed, we may both attempt to avoid that impression altogether, further diminishing hot mode signals in favor of cooler ones. But while we might shy away from purely tit-for-tat interactions, relationships can be built up, slowly over time, by invitations and fellow-feeling, and, conversely, if someone *never* has us over, eventually the relationship will die.

Our methods allow us to track the extent to which actual social practices are hot or cool; a warm mode will show up in our data as both a hot mode and a cool one. I bet you want to know what happens! Because science.

2 Our findings

Our findings are simple. Cool mode reciprocity is widespread in the village. Hot mode reciprocity is much rarer: it exists, on rapid timescales, among families that are not close kin. In both cases, our signals are strong (large effect size) and significant ($p < 0.01$ for all effects discussed here).

3 Our data

Our data is time resolved, which is necessary if we want to track the hot mode. Over the course of five months in the village, we (meaning my co-authors, but certainly not me) tracked, twice a week, which chicha parties people had attended. A particular line of the data set might be `Day 45: F12 hosted F9 one time`, which would record that, in the three days previous to day 45, a member of family twelve reported hosting family nine.

Another line might be `Day 45: F8 attended a party by F11 one time`, indicating that we had a record of a member of family eight claiming to have attended a party by family eleven. Yes, this means we double up and, ideally, have two data point for each host-guest pair—we find, however, that guests are shy: a little less likely to report attending a party than the host is to report inviting them. (Or, conversely, that hosts are a little boastful.)

We track any event for with at least one report (host-guest or guest-host) exists; our data then can be summarized by $h_{ij}(t)$—the number of times i hosted j at time t.

Finally, a line might be `Day 45: F8 did not attend a party by F11`, indicating that family eight was surveyed, and did not report an interaction with family eleven. We track these because not every family was interviewed on every survey day. This is the "observation window"—$o_{ij}(t)$ equal to one means that i hosting j could have been observed at time t (because either i or j was surveyed at that point).

4 Tracking the Cool Mode

The first part of the paper tracks the existence of stationary bonds—the cool mode. We look at the extent to which the rate of A hosting B correlates with the rate of B hosting A. We consider the probability of hosting, so we need to take into account the observation window. Because we're tracking stationary bonds, we dump all of the time information, and consider only the probability of a hosting event independent of time. We end up with a weighted, directed network, elegantly displayed in Fig. 1.

Paul hammered this with a massive multi-level logistic regression. The goal is to track what features of a host-guest pair influence the probability of a hosting event. This is a logistic regression, meaning that as the various parameters of interest vary over an arbitrary range, the probability moves smoothly between zero and one—see Eq. 1 in the paper.

We considered a variety of interfering effects, including boring ones, such as geographical location, that might influence hosting probabilities but that do not imply the existence of a social mechanism. We also include overall sociability of both host and guest—if a family simply has more chicha to share, it will tend to increase their hosting probabilities, again without implying the existence of a real underlying reciprocal mechanism. We also include age and age^2 as additional feed-ins.

We find a strong signal, over and above these random effects, that i hosting j predicts j hosting i—our baseline cool mode reciprocity. An increase of one percentage point in i hosting j leads to a 22% increase in j hosting i ($p < 0.001$).

Interestingly, kinship reduces this effect. While kinship itself predicts higher levels of reciprocity, this is due to the fact that kin tend to host each other at higher baseline rates. But the effect of kinship on hosting is such that when i hosts j at a certain rate, if i and j are kin, the response is lower. Kin, in other words, seem to tolerate greater asymmetry in their reciprocal exchanges, which makes sense.

5 Tracking the Hot Mode

In the hot mode, an invitation triggers (probabilistically) a return invitation. To study this, we use the timing information. We're interested, in particular, in the timescale of reciprocity—how long does a favor "hang" before being returned?

We quantify this by $R_{ij}(\Delta t)$, the conditional-action reciprocity between i and j. This is the fractional increase in the probability that j hosts i, given that i hosted j Δt steps earlier:

$$R_{ij}(\Delta t) = \frac{P(\text{“}j\text{ hosts }i\text{” at time }t+\Delta t \mid \text{“}i\text{ hosts }j\text{” at time }t)}{P(\text{“}j\text{ hosts }i\text{”})}. \tag{1}$$

In the paper, we average this quantity over all pairs, and plot $R(\Delta t)$, the overall conditional reciprocity as a function of time. There are better and worse ways to estimate R_{ij}; see Eq. 3 for a reasonable one that takes into account the observation window.

When $R(\Delta t)$ is large (greater than one) for some spacing, this indicates a hot mode reciprocity effect in play. In the absence of the hot mode, it will be precisely unity. In some cases, it will go below unity—if I host you three days later, it may depress my hosting probability ten days later. The hot mode is about the allocation of hosting over time, in response to another family.

5.1 Null Models Galore

We know stationary bonds exist. So we're really worried that these bonds will lead to false signals of a hot-mode pathway. Consider a simple example where i hosts j 5 times in 20 days, and j hosts i 7 times in the same span. If there's no hot mode, and there are no conditional reciprocity effects, these events will happen randomly and independent of each other. In that case, the data might look like

$$i: \quad \texttt{HXHXXHXXXHXHXX}$$
$$j: \quad \texttt{XHXXHXXHHXHXHH}$$

where H indicates a hosting event. Note, however, that despite the lack of correlation in the underlying process, the probability of j hosting i in the next observation period after i hosted j is measured to be $3/5$, or 60%, compared to the baseline of $7/20$. This would imply an R_{ij} of 1.7 despite the absence of any hot-mode effect.

Another effect is seasonal variation (see Fig. 2)—as the seasons change, and the village shifts from one task to another, the opportunities for parties change, and the overall level of hosting shifts. We do not want to confuse an empty village with a failure of people to return favors (or the resumption of parties later as a massive response to the few parties that were held in the downtime).

To make sure we're not confusing stationary bonds and seasonal shifts with the event-by-event return of favors, we generate null models. These preserve the stationary correlations and the overall seasonal variation.

5.2 Implementing Null Models

The simplest way to build these null models, in the modern era, is to write a bit of code that, literally, simulates the whole process. We create "fake" data that contains the stationary mechanisms but not the hot mode, and measure the same quantities as we did on the real data. If the fake data has excursions in R of similar size, we feel a bit silly—all those interesting signals could have been generated by statistical fluctuations in a pure cool-mode system.

We have three null models, but you can focus on the most rigorous one, the "full-heterogenous". This takes into account all of the cool mode activity.

Here's a bit of pseudo-code so you can see the process in action for the full-heterogenous null:

1. for each day:
2. set n to the average number of parties in the data for that week.
3. pick n families to be the hosts, based on the stationary hosting probabilities P("i is hosting").
4. for each hosting family i, pick guests j from the village based on stationary pairwise hosting probabilities, P("i hosts j"|"i is hosting").
5. for each of those pairs i and j, reproduce the original survey strategy:
   (a) report "i hosts j" if i was interviewed.
   (b) report "j is hosted by i", with probability $1 - p_{\text{omit}}$, if j was interviewed.
6. for the remainder of the pairs, report no hosting and no attendance if surveys were taken in the original data.

Following this recipe will create a fictitious data set that reproduces the overall levels of hosting, seasonal fluctuations, the effect of forgetful guests, and, most importantly, the pairwise preferences for each pair—the "stationary bonds". We're taking for granted the cool mode that Paul discovered.

The observant will notice that our null model generates an exponential distribution for party sizes for any particular host (a reasonable approximation to the data). And we will slightly underestimate the number of parties because of incomplete coverage by the observation window. More sophisticated models produce very similar results—these effects are minor. When you build a null model, it's good just to check that all the things you think you're preserving (temporal patterns, stationary probabilities) are truly preserved. Just measure them and compare to the real data!

We expect the average R to be unity in the nulls; but because of finite data and non-linearities in our estimation, fluctuations are seen over and above this in our simulated models. We thus renormalized everything, so that we plot R/R_{null} as a function of Δt.

5.3 Hot Mode and Close Kin

Fig. 3 shows that, once the nulls are taken into account, we have little evidence for an overall signal. Most of what appears to be strong hot-mode reciprocity turns out to be fluctuations that can be explained by reference to stationary bonds alone. This was somewhat disappointing—but, at the same time, intuitive. There are many modes for favor return, and we expect reciprocity to be, generically, cool in these villages with long-standing social ties.

However, then Paul suggested we look a little deeper into the hot mode signal (Fig. 4). We looked at the most rapid timescales—favor return on the order of three days or less. When we break this out between close kin ($r \geq 0.25$, or parents/children/uncles/aunts/nephew/niece) and non-close kin (everyone else), we see a strong effect. Distant kin and non-kin show a strong ($R \approx 3$, a boost of a factor of three in hosting probability) and significant ($p < 0.01$) signal. Conversely, kin show reduced reciprocity (*i.e.*, a tendency not to return favors), but this is not statistically significant compared to the null.[1]

This is a nicely interpretable result, indicating that dynamical signals of fair-play and sincerity are more important when distant kin and unrelated families interact—while replicating the stationary-mode finding that close kin can tolerate asymmetries, in this case, as they unfold over time.

6 Beyond the human

After we wrote the paper, we came across some parallel work in animal behavior [2] which looked at similar questions with different techniques for the case of grooming in chimpanzees.

References

[1] Paul L. Hooper, Simon DeDeo, Ann E. Caldwell Hooper, Michael Gurven, and Hillard S. Kaplan. Dynamical structure of a traditional amazonian social network. *Entropy*, 15(11):4932–4955, 2013.

[1] Note that Fig. 4 presents both the p value for difference from the null, and a bootstrap estimate of errors on R; it is possible to get a non-unity value of R in the null, and for the bootstrap errors to suggest that it is "real"—but this is not statistically significant. The bootstrap errors assume iid samples. In the absence of a strong positive model for the hot mode, which would allow for a Bayesian estimation, it's the best we can do

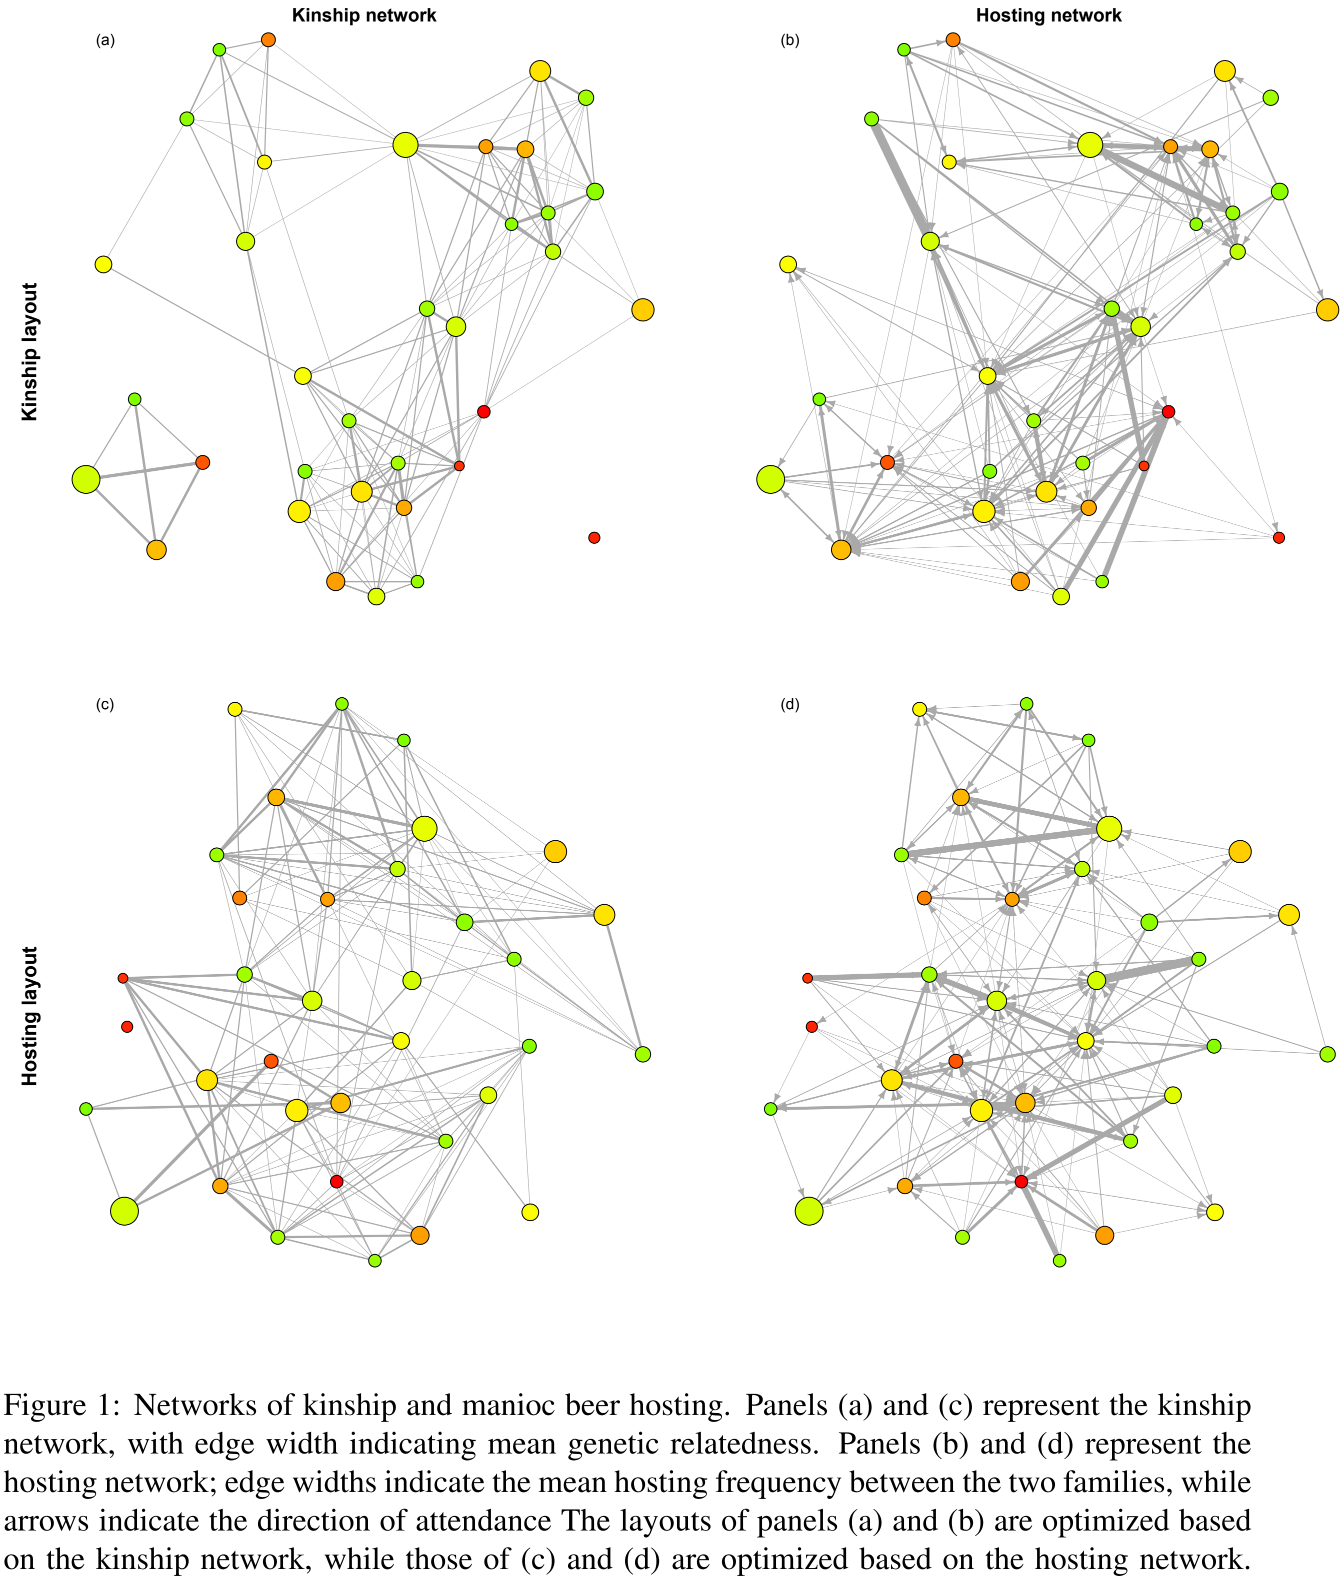

Figure 1: Networks of kinship and manioc beer hosting. Panels (a) and (c) represent the kinship network, with edge width indicating mean genetic relatedness. Panels (b) and (d) represent the hosting network; edge widths indicate the mean hosting frequency between the two families, while arrows indicate the direction of attendance The layouts of panels (a) and (b) are optimized based on the kinship network, while those of (c) and (d) are optimized based on the hosting network. Node color indicates mean age of household heads (ascending in age from green to red), while node size indicates family size. Fig. 4 in the original paper.

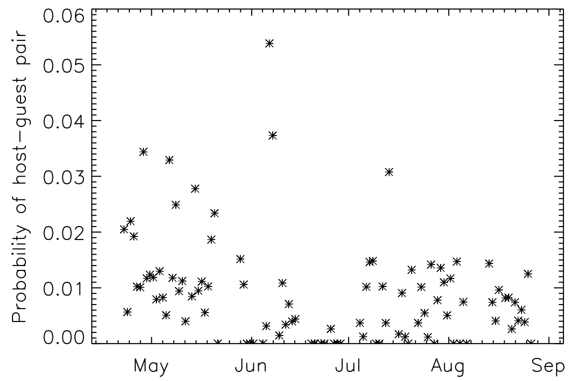

Figure 2: Seasonal effects in hosting. Shown here, for the discrete sampling points, is the probability that an observed pair is a hosting pair. A lull in late June separates two more active periods; the May period is overall the most active, followed by the July and August season. Fig. 1 in the original paper.

[2] Cristina M Gomes, Roger Mundry, and Christophe Boesch. Long-term reciprocation of grooming in wild west african chimpanzees. *Proceedings of the Royal Society B: Biological Sciences*, 276(1657):699–706, 2009.

[3] Samuel Bowles. *Microeconomics: Behavior, Institutions, and Evolution*. Princeton University Press, 2009.

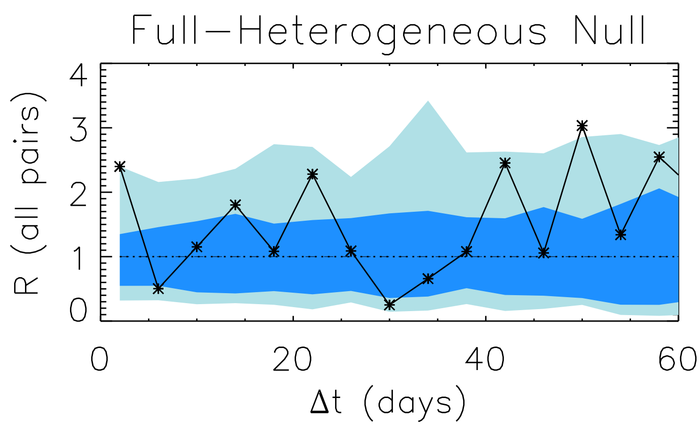

Figure 3: Average reciprocity, as a function of separation Δt, compared to the heterogenous null. Blue bands indicate the one and two-sigma ranges for null model. Nearly all the signals appear to be consistent with a world in which cool stationary-bond reciprocity explains patterns of giving and receiving—but see the next figure. Fig. 5c in the original.

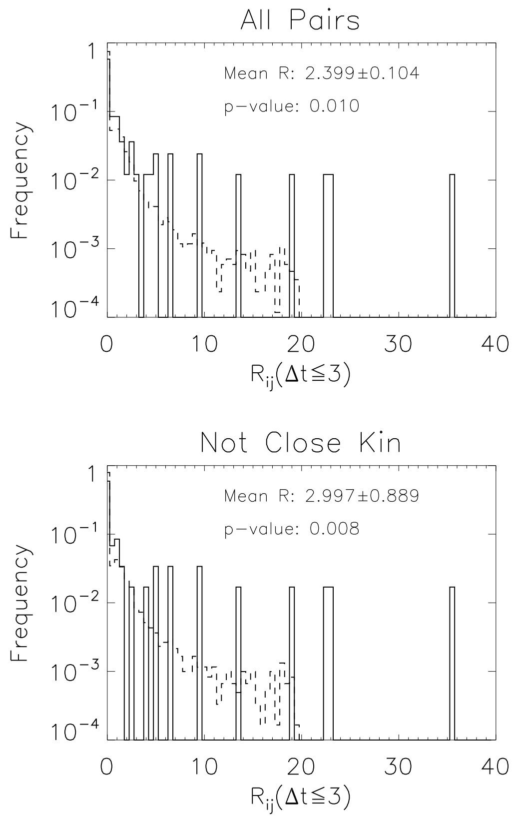

Figure 4: Rapid timescale ($t \leq 3$ days) reciprocity values. Solid lines show the distribution of observed data; dashed lines that for the full-heterogeneous null. Top left panel: all pairs; top right panel: closely related ($r \geq 0.25$) pairs; left panel: distant and non-kin ($r < 0.25$) pairs. Families that are not close kin show significant evidence of strong short-term reciprocity. Fig. 6 in the original.