



## Stanford Social Web (ca. 1999)





## LADA ADAMIC

CURRENT  
PAST  
ALWAYS

ME  
YOU  
THEM

WORK  
PLAY

*ladamic@umich.edu*



## THEM

TJ  
Stacy  
Eytan  
Orkut  
Rajan  
Ryan  
Wendy  
Sergio  
Amy  
Josh  
Annie  
Martin  
Jen  
Amnon

# awkwardness of specifying friendship

The screenshot shows an email client window titled "2005to2007 - umich". The search bar contains "robert ackland" and shows "5 messages". The message list is filtered by "Sender".

Subject	From	Date	
my blogosphere paper	Robert Ackland	11/3/05 5:25 PM	• • •
SNAC participant mailing list	Hank Green	11/21/05 12:04...	• • •
↻ "Divided They Blog" - availability of data?	Robert Ackland	11/13/07 7:18 PM	• • •
↻ Re: "Divided They Blog" - availability of data?	Robert Ackland	11/26/07 8:06 PM	• • •
Re: "Divided They Blog" - availability of data?	Robert Ackland	11/29/07 1:42 ...	• • •

On the right, a sidebar titled "Selected Nodes(19)" lists names:

- Ackland, Robert
- Amaral, Luis
- Aral, Sinan
- Barabasi, Albert-Laszlo
- Borner, Katy
- Burt, Ronald
- Contractor, Noshir
- Han, Jiawei
- Hargittai, Eszter
- Leskovec, Jure
- Macy, Michael
- Monge, Peter

# Are online social networks accurate?

---

- ▶ OSNs provide a goldmine of data
- ▶ Social networking features enhance functionality (e.g. recommendations)
- ▶ How can one elicit honest ratings?
  - ▶ when users review items
  - ▶ when users rate other users' contributions
  - ▶ when users rate other users





# Papers

---

- ▶ Lauterbach, Truong, Shah, Adamic, 'Surfing a web of trust: Reputation and reciprocity on CouchSurfing.com', SIN 2009
- ▶ Teng, Lauterbach, Adamic, I rate you. You rate me. Should we do so publicly?, WOSN 2010



## related work

---

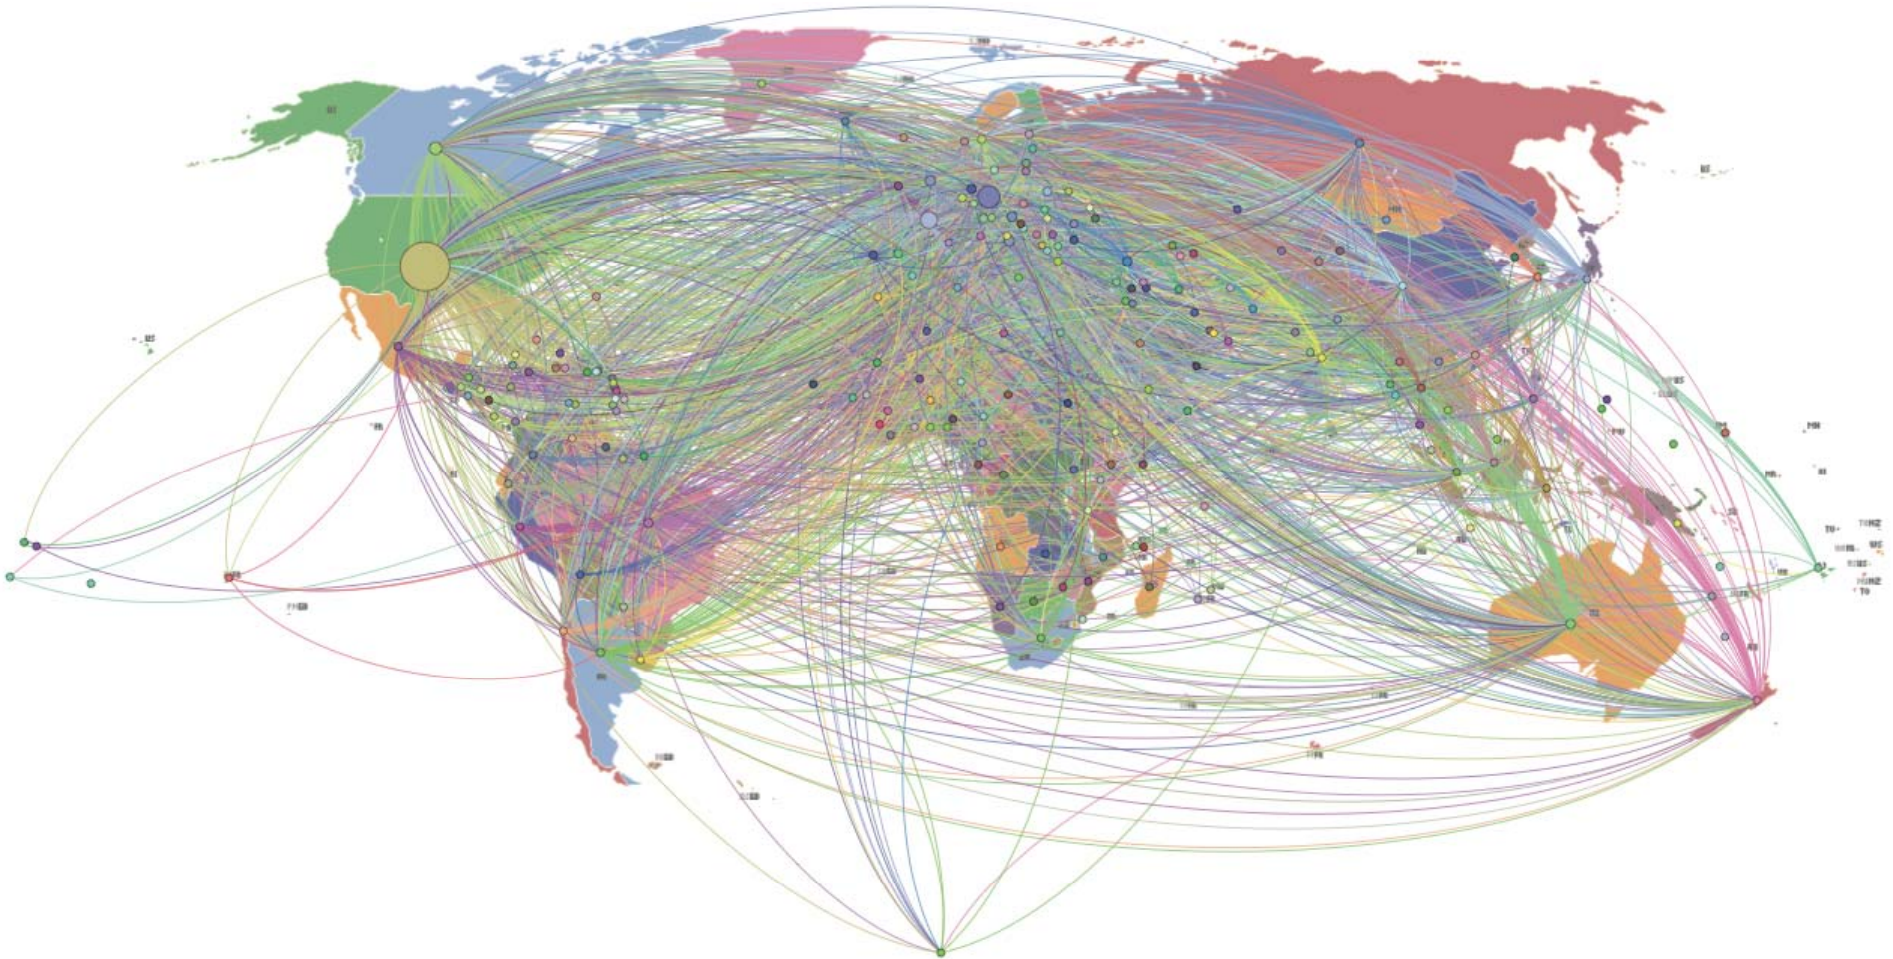
- ▶ Predicting trust & ratings
  - ▶ propagation of trust and distrust (Guha et al WWW'04)
  - ▶ signed social networks (Leskovec et al WWW'10, CHI'10)
  - ▶ trust based on shared attributes
  - ▶ trust on CouchSurfing.org (Bialski and Batorski, Sunbelt 2007)
- ▶ Relationship between online ratings and offline relationships
  - ▶ strength of ties (Gilbert and Karahalios CHI'10)
  - ▶ online ties and social capital (Ellison et al, JCMC 2007)
- ▶ Anonymity and ratings
  - ▶ e.g. teaching evaluations, peer review



# Why online reputation systems matter

---

hospitality exchange on CouchSurfing.org






# Direct reciprocity


---

- ▶ 12-18% of stays are directly reciprocated



# Generalized reciprocity

- ▶ Generalized reciprocity: large strongly connected component (150,000 individuals: 1/3 of active users)



## How can this be?

---

- ▶ *It's the exact same [] as if you had just met, casually met somebody while you were traveling and they said, "You know, you could stay on my couch if you want". [P9]*
- ▶ *We arrived in Brussels at 9 o'clock in the morning and the guy that we were staying with, our host, had to come to his work and he just handed his house keys right over to us and said, ``This is where I live. Go here, go here and my wife and I won't be home until after 5."*

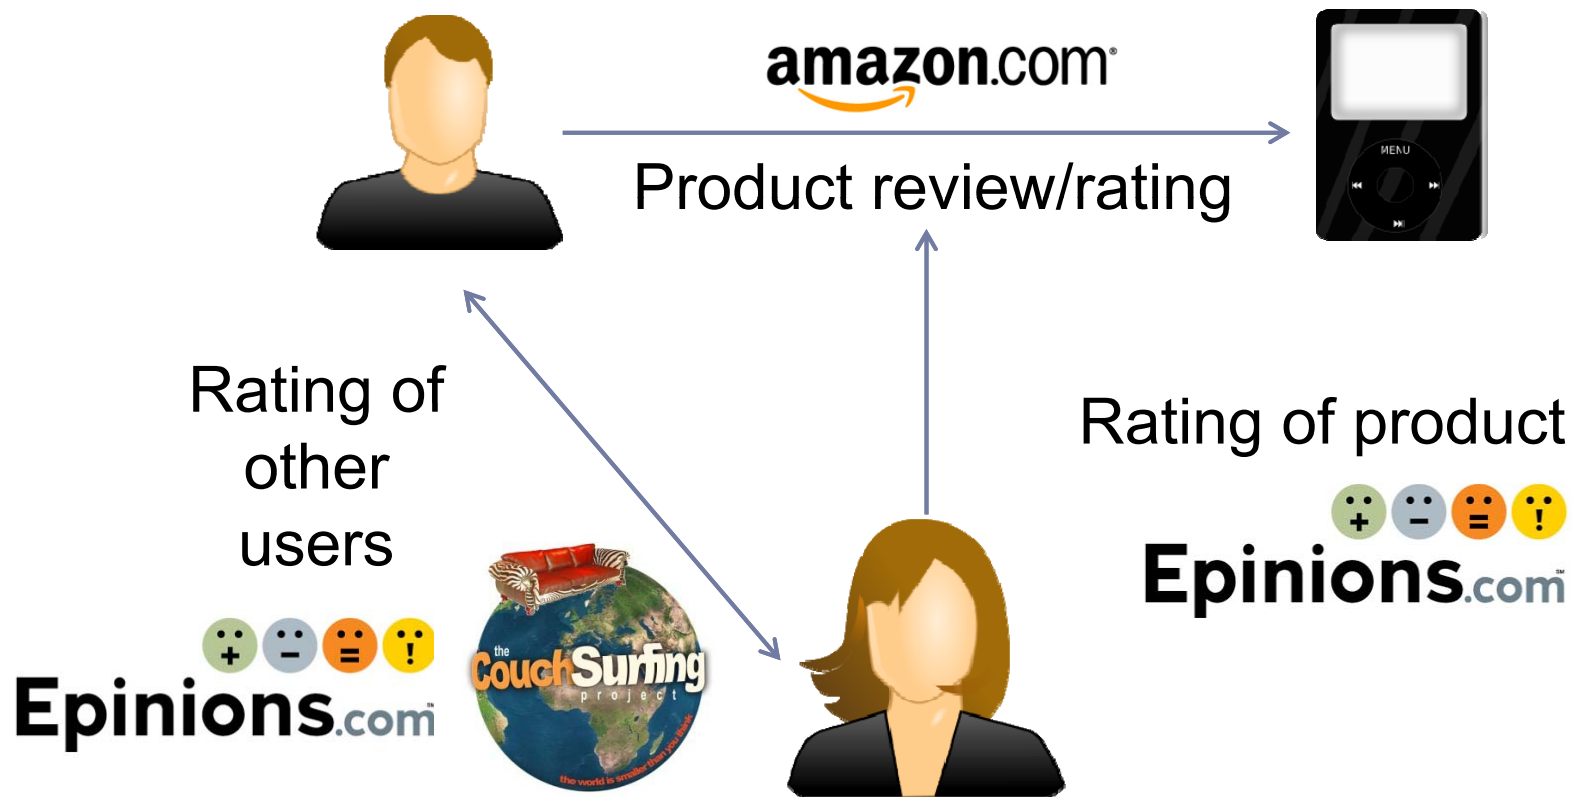


# Can online ratings really be trusted?



# Datasets: Who is rating what/whom?

---





# Amazon

---

- ▶ Amazon.com provides a platform which allows users to review products
- ▶ Users can decide to use a pen name or real name to review products
- ▶ 15 thousand reviews from top 1500 reviewers (about one half using pen names)



# Effect of anonymity in absence of reciprocity

**amazon.com**

attribute	pen name	REAL NAME™	statistically significant
product rating # stars	4.19	4.21	no
# reviews	498	551	yes
length of review (words)	364	377	yes
# of fan voters	28.6	37.1	yes

**PUBLIC**

**ANONYMOUS**

**RECIPROCAL**

# Epinions

---


- ▶ Epinions.com allows users to share product reviews.
- ▶ Users can write reviews, rate other users' reviews, and specify which users they “trust” or “distrust”
  - ▶ ~800K user-to-user ratings (trust or not)
  - ▶ ~100K users rating 3 million articles



# Anonymity when there is potential for reciprocity



- ▶ Anonymous ratings of others' reviews are lower (3.84) on average than identified ratings (4.71)
- ▶ For the same user, anonymous ratings still average lower (4.01) than identified ones (4.76)



PUBLIC

ANONYMOUS

RECIPROCAL

## Evidence for reciprocity?

---

- ▶ We average multiple user-to-article ratings into user-to-user ratings.
- ▶ Rating from A to B is correlated with rating from B to A ( $\rho = 0.48$ )
- ▶ # of ratings from A to B and B to A also correlated ( $\rho = 0.49$ )
- ▶ anonymously given ratings between users have much lower correlation ( $\rho = 0.14$ )

PUBLIC	ANONYMOUS	RECIPROCAL
--------	-----------	------------





- ▶ Epinions allows users to “trust” others publicly, but “distrust” privately
- ▶ non-trivial fraction (14.7%) are “distrust” ratings.
- ▶ For pairs of users who mutually rated one another (35% of public trust, and 6% of private distrust were reciprocated):

A->B, B->A	observed	expected
trust, trust	97.1%	72.8%
trust,distrust	1.1%	25.1%
distrust,distrust	1.8%	2.2%



# CouchSurfing

---

- ▶ Can trust and friendship be quantified?
- ▶ What factors into a rating?
- ▶ What are trust and friendship anyway?



# CouchSurfing

---



- data: 600K users, 3 million edges
  - ~ 500 survey respondents
  - 18 interviews
- ▶ Users can do the following for other users:
- ▶ specify friendship level (e.g. acquaintance, friend, best friend)
  - ▶ specify how much they trust them (e.g. "somewhat", "highly")
  - ▶ vouch for them
  - ▶ leave positive, neutral, or negative references

**PUBLIC**  
**PRIVATE**


**PUBLIC**





# Scarcity of public, negative ratings when identified

- ▶ Users leave a positive reference for 87.7 % of those they host and for 90.1% of those who host them
- ▶ Neutral/missing references are confounded in data
- ▶ The ratio of positive to negative references is 2500:1!



**PUBLIC**


**ANONYMOUS**

**RECIPROCAL**

# Why only positive references?

- ▶ 55 % say they always leave references

- ▶ For those who don't:



percentage of respondents	reason for not leaving a reference
51.3%	too busy
31.7%	neutral experience and didn't want to state it
12.1%	negative experience and didn't want to state it



# Lack of negative references

---

- ▶ Concern for others' reputation:
  - ▶ *the few times when I had a neutral experience, I believe it was because of personal character differences, and not because I had complaints against the person in question. Somebody else might have a positive experience, why write them a negative/ neutral one and prevent people from considering the person? [S83]*



## Is there reciprocity in references?

---

- ▶ *generally, I prefer the host or the surfer to leave a reference first, so I can kind of... And I do gauge. I mean I gauge on the way that their reference was. If their reference was very detailed and inclusive of our expense then I will meter that, and if it was a little more scoped. [laughter] then I will narrow mine down too. [P9]*
- ▶ *I usually don't write references to those I hosted and didn't leave a reference on my profile [S399]*



# Reciprocity deterring negative ratings?

---

- ▶ fear of reciprocal action

- ▶ *But the big problem is that if you leave a bad reference, what happens then. What will that person say about you. You leave a bad reference and he can do the same. And its not true. [S37]*

- ▶ results in lack of information about negative experiences

- ▶ *I chose not to leave a reference because I just felt uncomfortable [] then I actually ended up speaking to the country ambassador [] and she told me that several other surfers have had the same experience with this particular host, and I just chose on a personal level not to leave a reference for him. Today, I regret that. I wish I had left a*

- ▶ *negative reference so no one else would have been put in that situation that we were [P9]*

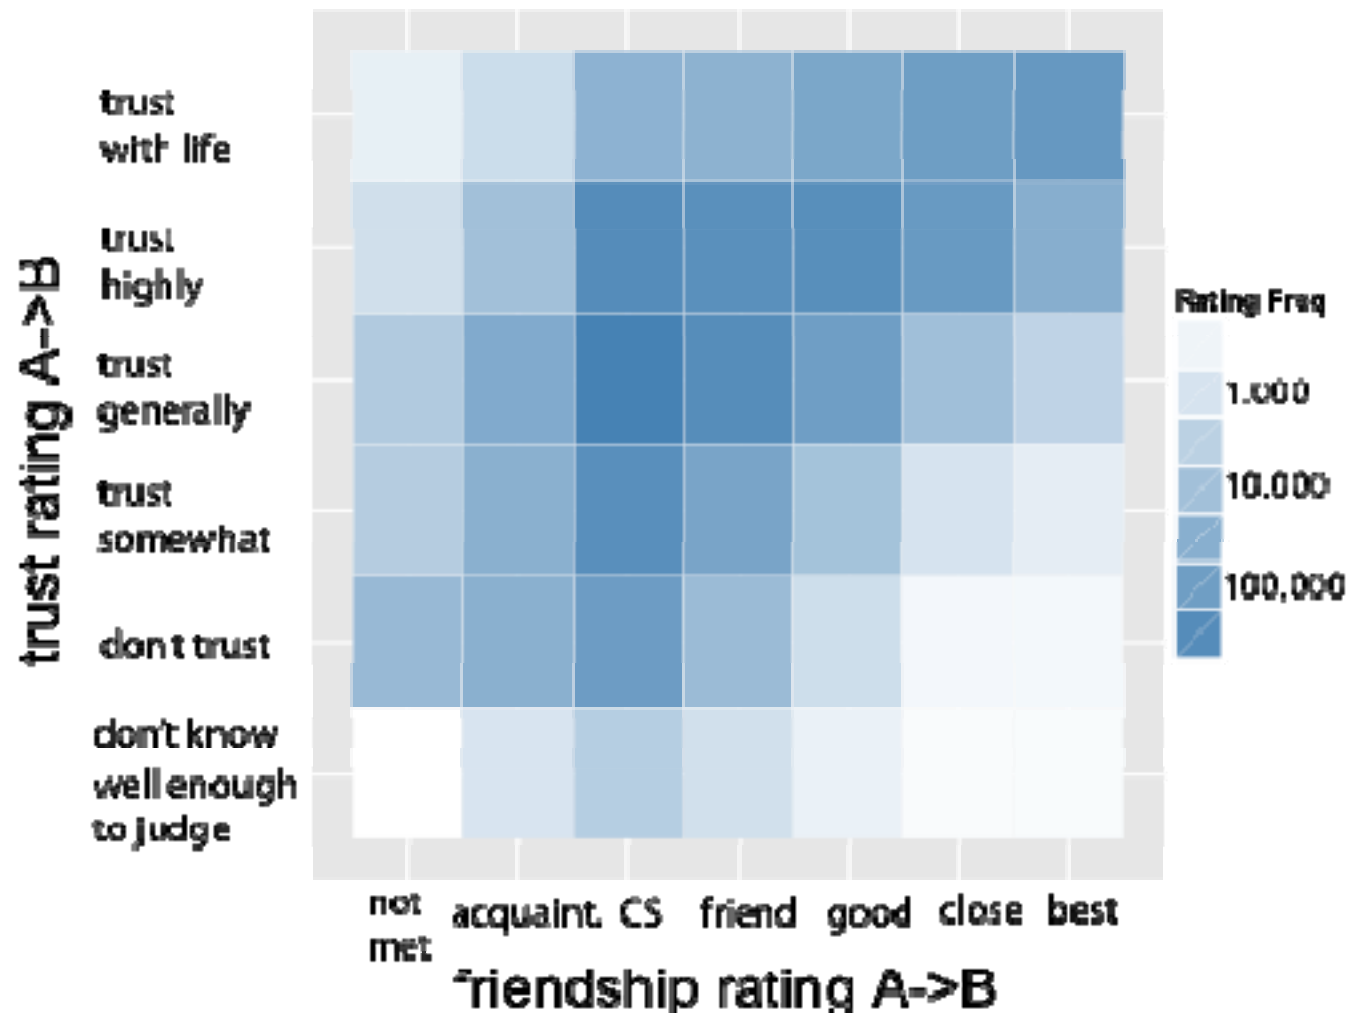
# Can trust and friendship be quantified?

---

- ▶ How easy is it to quantify friendship and trust?
  - ▶ *One “level” is never enough to point to the correct tone of a human relation. [S291]*
- ▶ Can friendship be interpreted as trust?
  - ▶ *I think close friends you trust, but I don’t think everyone you trust is a close friend. [P12]*
  - ▶ *Friendship includes trust. You can trust someone, but still without that person being a friend. I guess friendship is a more elusive concept and therefore more difficult to judge. [S312]*



# Is the friendship/trust asymmetry reflected in the data?





## Time and friendship/trust

---

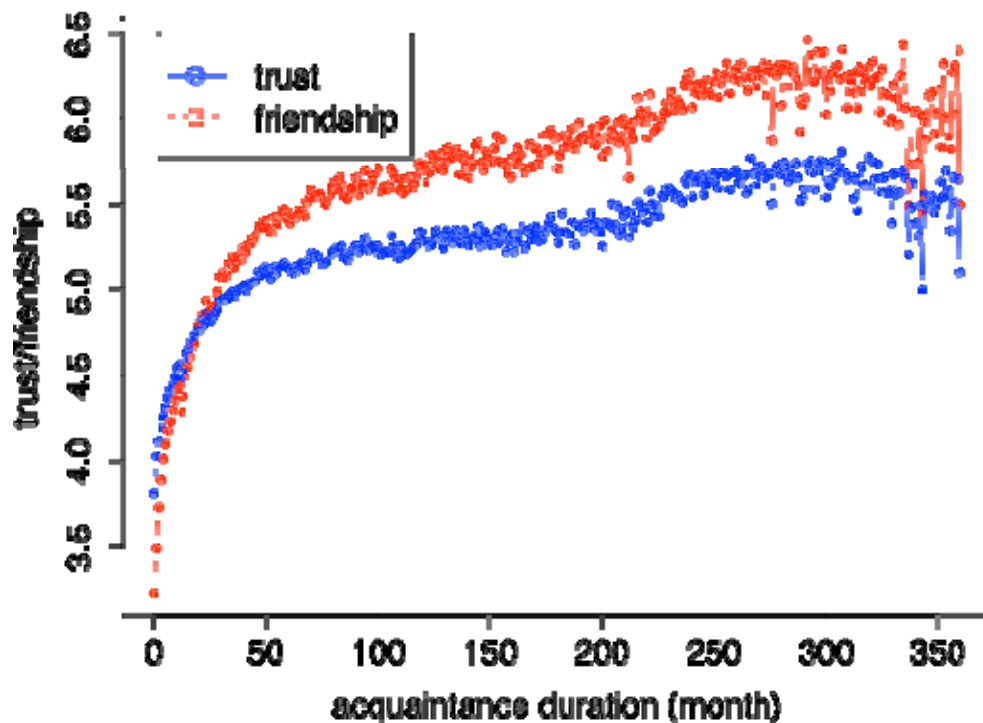
- ▶ *trust takes, you know, in a lot of cases years to build, whereas friendship and that sort of thing can sort of happen instantly [P8]*
- ▶ *I have a gut feeling about who I could trust, but not so much about who actually counts as my friend. [S10]*
- ▶ *In general I only \*really\* trust my closest, real life friends. That takes years to earn. [S256]*



# Which one takes time?

---

- ▶ both do



*In part, it's a question of hours that we stay together. If we still stay together the level of trust increases... Not always, but increases. And sometimes you understand some limitations so maybe your level of trust can arrive to a certain level and that's it, and doesn't increase over that level.  
[P10]*

# What else does it take?

---

	correlation w/ friendship	correlation w/ trust
how well you know other	0.719	0.670
log (how long...)	0.592	0.365
days traveled	0.404	0.250
same country	0.212	0.075
abs. age difference	-0.106	-0.068
days surfed	0.268	0.203
days hosted	0.279	0.193
same gender	0.078	0.045



## node vs. edge attributes


---

- ▶ To what extent are trust and friendship node vs. edge properties?
- ▶ Friendship is something between two people (edge attribute?)
- ▶ Does trust depend more on the trustworthiness of the individual (node attribute?)
  - ▶ median variance per node (normalized by variance over all ratings of individuals w/ same # of friends)
  - ▶ 0.60 (friendship) vs. 0.50 (trust)




# Reciprocity in friendship & trust

- *Public* friendship ratings are more highly correlated ( $\rho = 0.73$ ) than *private* trust ratings ( $\rho = 0.39$ )



(a) alignment of trust ratings between user pairs



(b) alignment of friendship ratings between user pairs

We omit trust rating of 2 (I don't know the person)

**PUBLIC**

**ANONYMOUS**

**RECIPROCAL**


**PUBLIC**

**ANONYMOUS**


**RECIPROCAL**

# Reciprocity friendship & trust

- ▶ reminder: friendship = public, trust = private
  - ▶ *It can be difficult to select a friendship level if I am unsure of how the other person may react or if I think they may see our friendship as being at a different level. [S114]*
  - ▶ *Cause sometimes you don't want to be unpolite[] ... you want to have the person the same friendship level. [S175]*
  - ▶ *the trust level is anonymous, and I tend to trust people more easily. that's why friendship level is more difficult: everyone can see it [S276]*



(a) alignment of trust ratings between user pairs



(b) alignment of friendship ratings between user pairs


## reactions to misaligned ratings

---

- ▶ Only 41% of users ever recalled noticing a misaligned friendship rating. Those who did typically did not attach much importance. But for some:

- ▶ *I once said one girl was a “good friend” - however, she added me as an acquaintance. It actually made me feel quite bad to hear that she didn’t even consider me as her friend [S491]*

- ▶ *Not a big deal but yes it feels not great. Because you see that the feelings about the friendships is not really mutual. [S31]*



# How useful are numerical ratings of human relationships?

- How important are the following when choosing whether to host or be hosted by someone?

	Very Important	Important	Neutral	Unimportant	Very Unimportant
how many vouches they have	8.6%	29.3%	37.1%	15.8%	6.9%
whether they are verified	9.1%	23.0%	36.4%	18.9%	11.6%
number of references received	23.0%	57.0%	13.7%	4.4%	1.1%
text of references received	47.6%	40.8%	8.2%	2.1%	0.8%
number of friends	4.4%	23.6%	37.7%	25.1%	8.0%
friends' friendship level	3.2%	19.0%	35.6%	25.0%	13.5%



# Why are textual references more useful?

---

- ▶ Many include information about the individual that signal to others the person's personality and interest
- ▶ It is possible to leave a **neutral** reference while using a seemingly **positive** tone.
  - ▶ *I've gone pretty keen on what certain references mean, and you can tell when a reference is just like a simpatico-nice; you-were-a-nice-person-reference: "[She] was great. She was very hospitable. She's a great host." That can mean in a sense you might be kind of boring. [P9]*



# Are truthful ratings reliable?


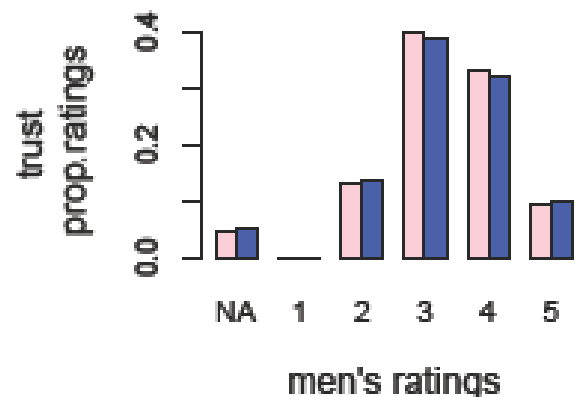


---

- ▶ Even if one were able to elicit truthful ratings, would there still be biases?




# Gender effects for trust & friendship

- ▶ Men rate both men and women about equally on trust and friendship
- ▶ Women rate other women more highly on both




# Age

- ▶ Trust is very slightly higher the smaller the age difference between rater and ratee ( $\rho = -0.06$ )
- ▶ Trust depends on age of ratee – typical CouchSurfing demographic preferred?




# Lest you think it is just CouchSurfing



## The U-bend

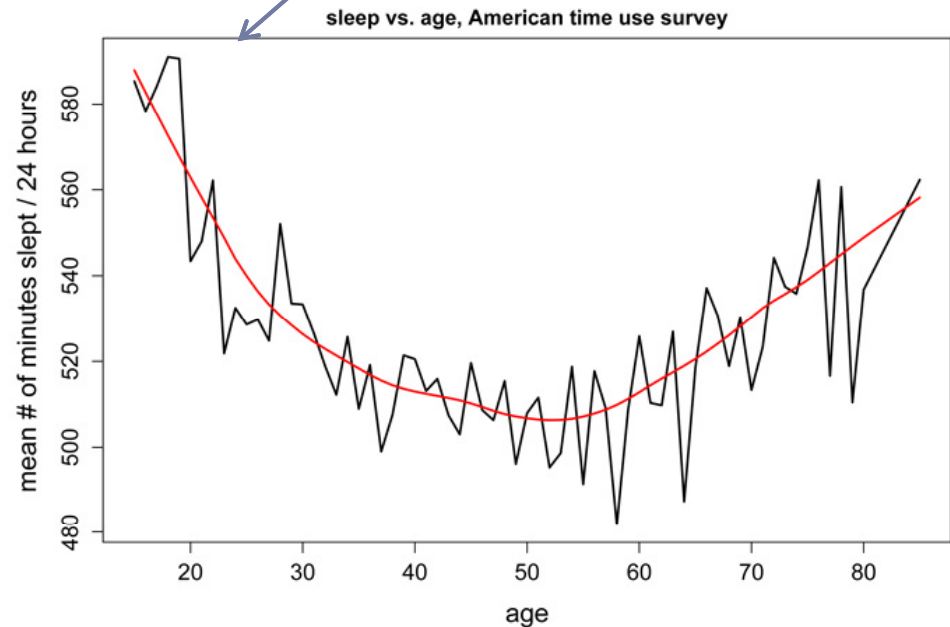
Self-reported well-being, on a scale of 1-10



Source: PNAS paper: "A snapshot of the age distribution of psychological well-being in the United States" by Arthur Stone

how you would rate yourself

how much sleep you are getting



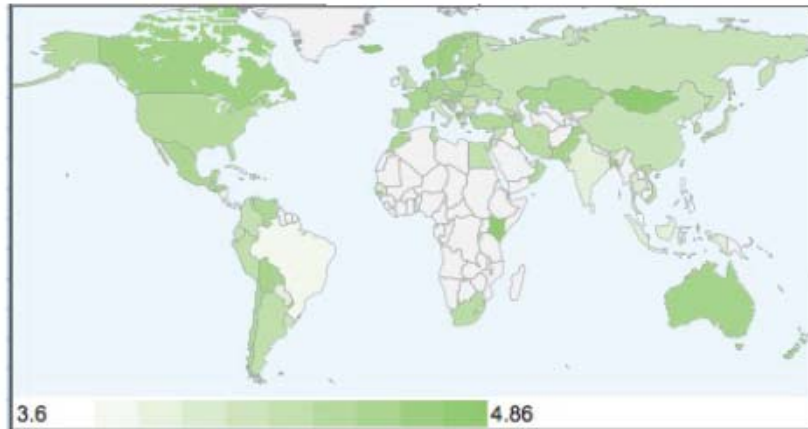
# Geography

---

- ▶ Closer friends tend to be geographically proximate
  - ▶ Friendship for one's countrymen (4.19) is higher than foreigners (3.65)
  - ▶ Trust for one's countrymen is higher than for foreigners (4.33 vs 4.16)
- ▶ Countries with similar cultural background tend to be trusting of one another (e.g. Austria and Germany)
- ▶ Sharing a border does not always correspond to greater trust (e.g. Canadians did not rate US contacts more highly)

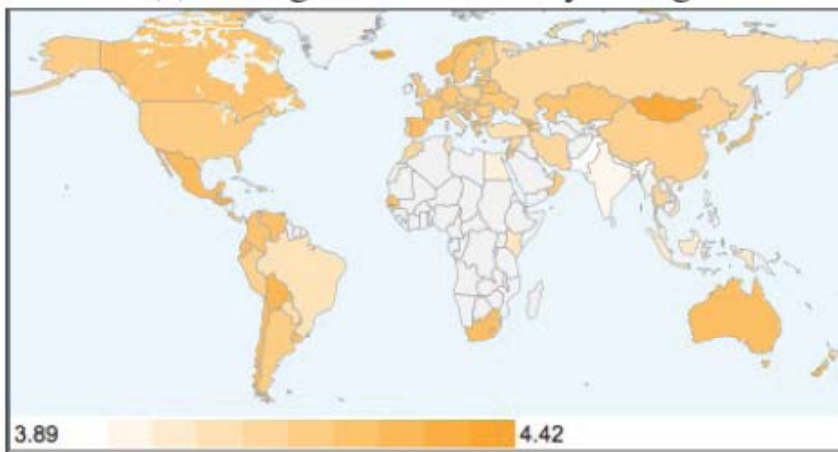


# Geography

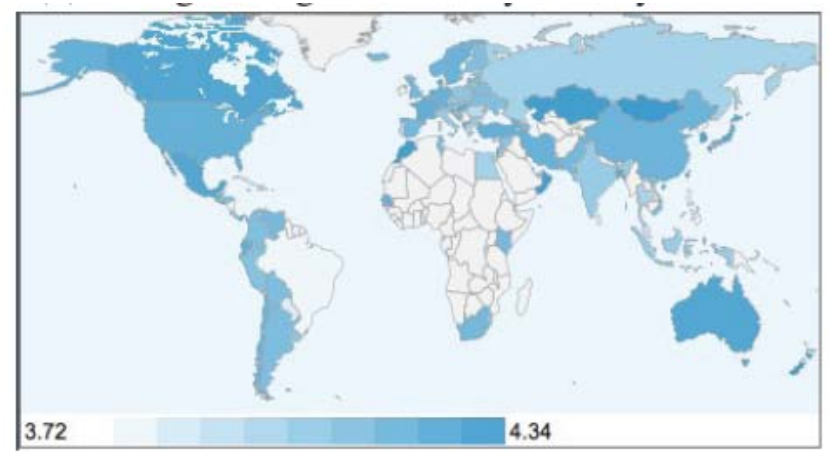


(a) average within-country ratings

- ▶ Within country ratings correlate with across-country ratings
- ▶ Language or culture?



(b) average ratings received by country residents



(c) average ratings given by country residents

# Vouching: eliciting more honest ratings?



“Respecting the significance of vouching is essential to the integrity of the network... It is very important that you ONLY vouch for people that you ... know well enough to believe that he or she is trustworthy.”

- ▶ You can only vouch for others if you have at least 3 vouches yourself
- ▶ Vouching forms a small “web of trust” in the network
  - ▶ 6.8% of users have been vouched at least once
  - ▶ 1.8% can vouch for others






# Who are users vouching?

---

- ▶ A high number of vouches are between “CouchSurfing friends”



Friendship degree:

1= Haven't met yet

2= Acquaintance

3= CouchSurfing friend

4= Friend

5= Good friend

6= Close friend

7= Best friend

## Reciprocity in vouches on CouchSurfing

---

- ▶ If A vouched for B, 70% of the time B also vouched for A
- ▶ Mean **private** trust score for *reciprocated* **public** vouches was higher (4.47) than unreciprocated ones (4.19)
  - ➔ lack of rating could signal lower trust



# Can one predict vouches?

---


- ▶ Logistic regression model (10-fold cross-validation)
- ▶ 71% accuracy in predicting whether a random edge is vouched
- ▶ Most predictive attributes were friendship degree, rating of experience, how they met



# Predicting vouches - network measures

---

- ▶ Two-step indirect measure for propagating vouches:



Indirect vouch score for A→D:  
 $= 1/n(B) + 1/n(C)$



# Predicting vouches - global measures

---

- ▶ Results from logistic regression for each variable alone:

Variable	Predictive accuracy:
Friendship degree	67.7%
Jaccard coefficient	55.8%
2-step vouch propagation	54.2%
PageRank	50.6%

- ▶ Global measures are poor predictors of trust & friendship



# Conclusion

---

- ▶ Social networks online are a product of true feelings and the difficulty in transforming them into hyperlinks:
  - ▶ Public, identified ratings tend to be positive when there is potential for reciprocity
  - ▶ Even truthful ratings can be biased based on demographics
  - ▶ Trust and friendship are localized on the network – global metrics are of little use in predicting them.



# Thanks

---

- For more info: <http://netsi.org>

