

Strengthening Weak Identities Through Inter-Domain Trust Transfer

Giridhari Venkatadri, **Oana Goga**, Changtao Zhong, Bimal
Viswanath, Nishanth Sastry, Krishna Gummadi




Max
Planck
Institute
for
Software Systems

KING'S
College
LONDON

Trustworthiness of online identities



Trustworthiness of online identities




Trustworthiness of an identity: The **likelihood** that the **identity will respect the terms of service (ToS)** of its domain in the future

Online identity-infrastructure



Online identity-infrastructure




strong
identity-infrastructure



weak
identity-infrastructure

Accountability	✓	✗
Anonymity	✗	✓
Adoption	✗	✓
Resistance to fake identity attacks	✓	✗

Online identity-infrastructure



Trusted certificate

strong
identity-infrastructure




Trusted certificate

weak
identity-infrastructure

Up to 40% of newly
created identities on
Twitter are malicious!!

- Accountability
- Anonymity
- Adoption
- Resistance to fake
identity attacks




Online identity-infrastructure



Challenge: How to reason about the trustworthiness of weak identities?

are of newly created identities on Twitter are malicious!!

- Accountability
- Anonymity
- Adoption
- Resistance to fake identity attacks



Trustworthiness of weak identities


Current techniques: Based on the **past activity** of each identity **within the domain**

Limitation: Domains need to observe the behavior of weak identities over time (**time lag**)

- Malicious users can still exploit new identities to misbehave
- Honest users must wait to acquire access to resources (e.g., Reddit posting quotas)


Key idea

Strengthen weak identities through
inter-domain trust-transfer



Key idea

Strengthen weak identities through
inter-domain trust-transfer



- ▶ Use the weak identities of users on other domains as external trust certificates

Why would this work?

1. Many honest users maintain weak identities on multiple domains
2. Users already interconnect their identities (e.g., social login)
3. Malicious attackers would incur additional costs
4. More established domains could provide good trust references for newer domains

This talk

1. Potential for inter-domain trust transfer
2. Inter-domain trust transfer framework
3. Leverage inter-domain trust transfer for identity curation

This talk

1. Potential for inter-domain trust transfer
2. Inter-domain trust transfer framework
3. Leverage inter-domain trust transfer for identity curation


Potential for inter-domain trust transfer

Can **activity signals** from **Facebook** and **Twitter** help **Pinterest** reason about trustworthiness better?

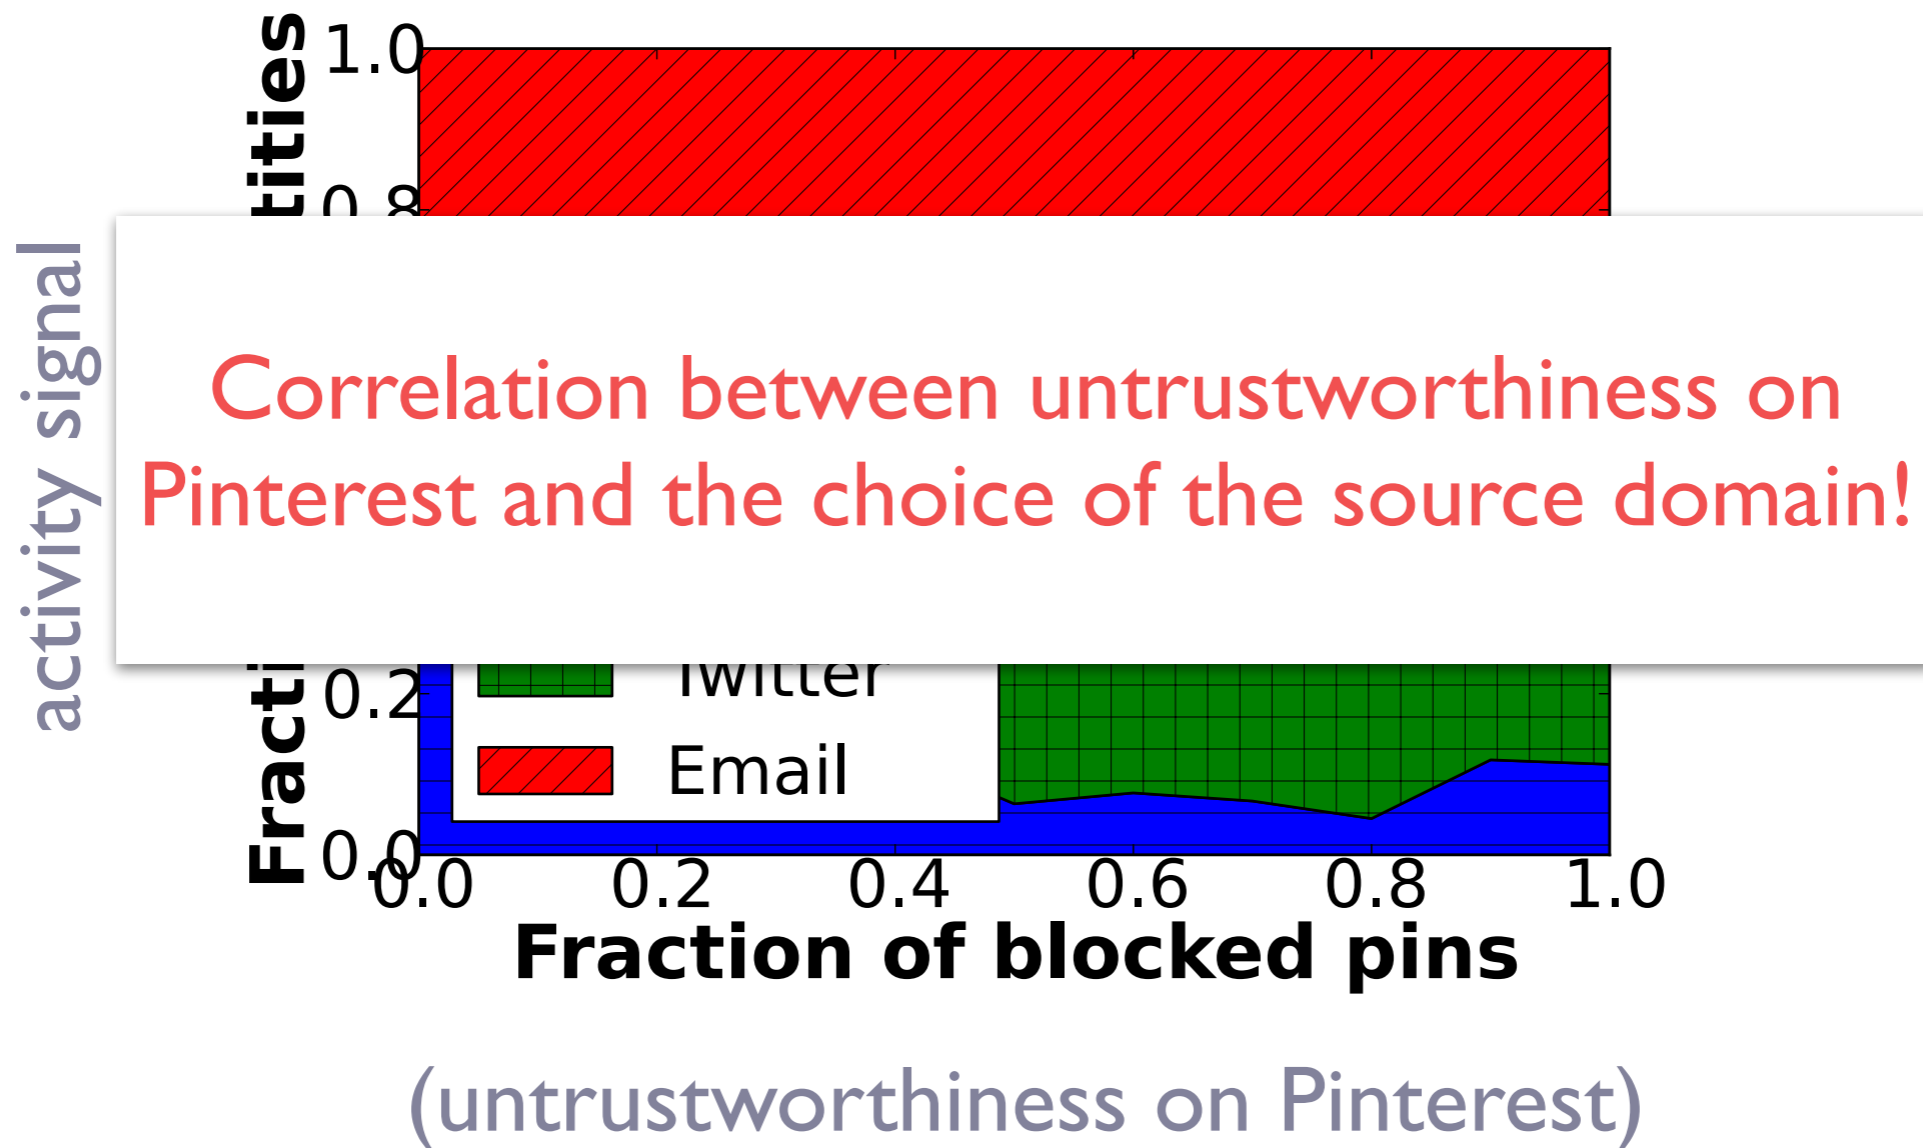
Dataset

- 1.7M random identities on Pinterest, and their matching identities on Facebook and Twitter
- **Activity signals** computed based on **public data** on Twitter and Facebook (e.g., account age, # followers, suspension)
- Diverse set of **untrustworthy identities** on Pinterest


Source domain and trustworthiness



Source domain and trustworthiness

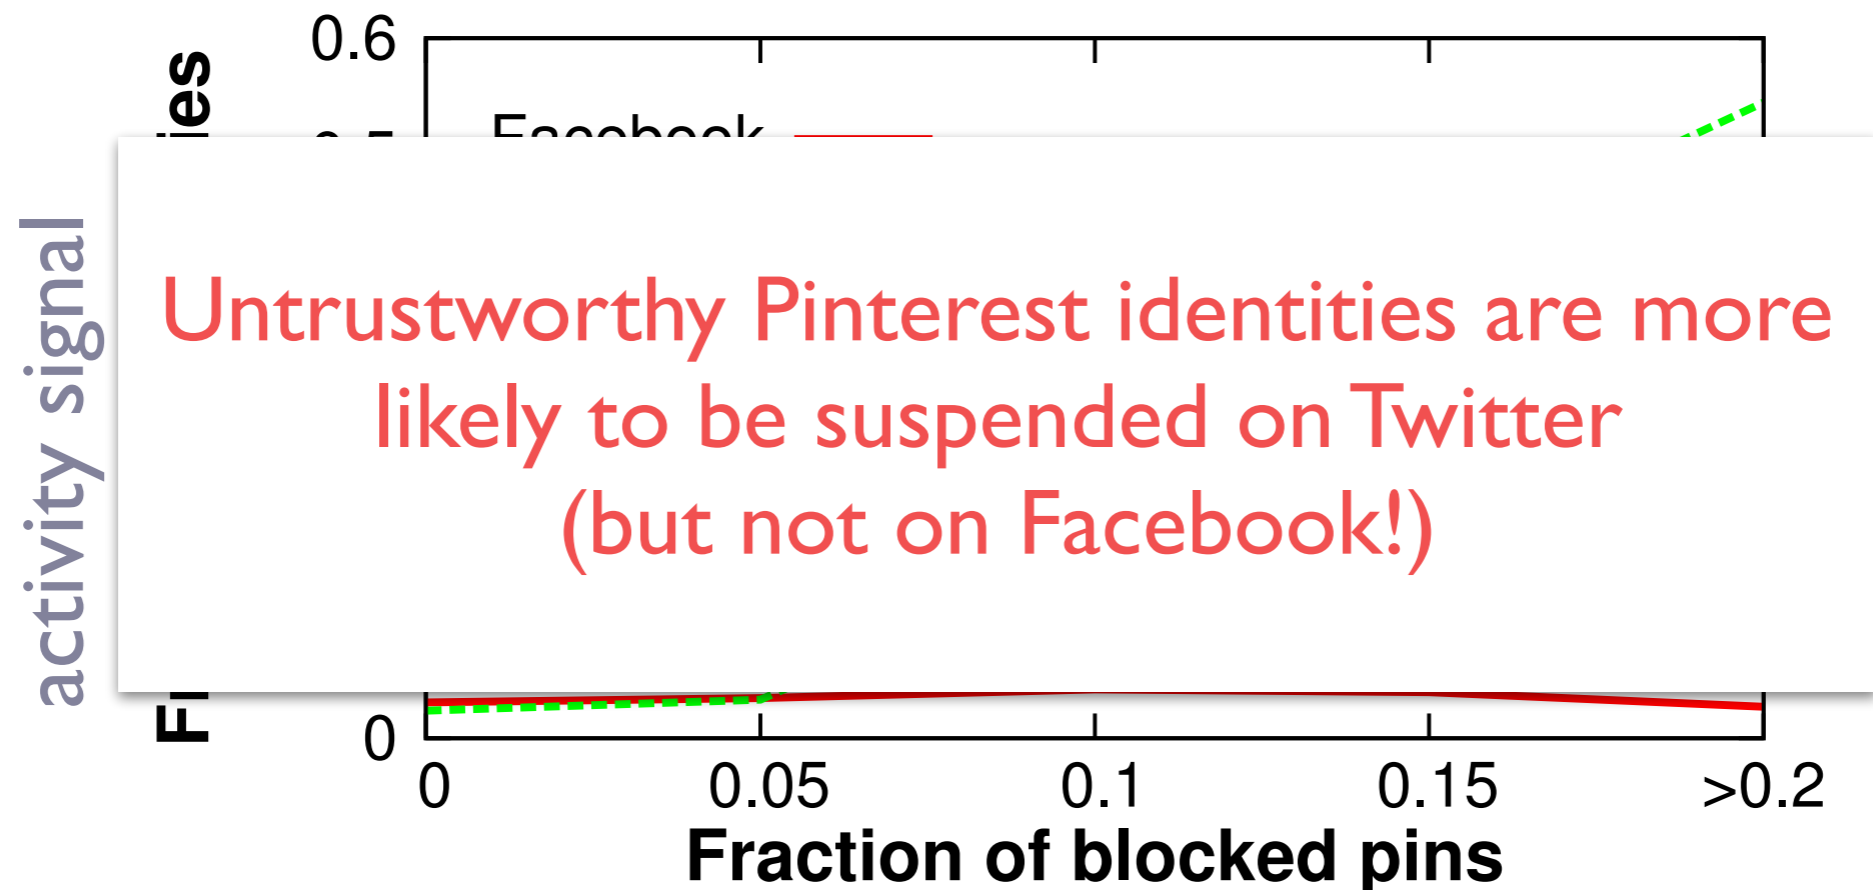


Suspension signal and trustworthiness



(untrustworthiness on Pinterest)

Suspension signal and trustworthiness



(untrustworthiness on Pinterest)


This talk

1. Potential for inter-domain trust transfer
2. Inter-domain trust transfer framework
3. Leverage inter-domain trust transfer for identity curation

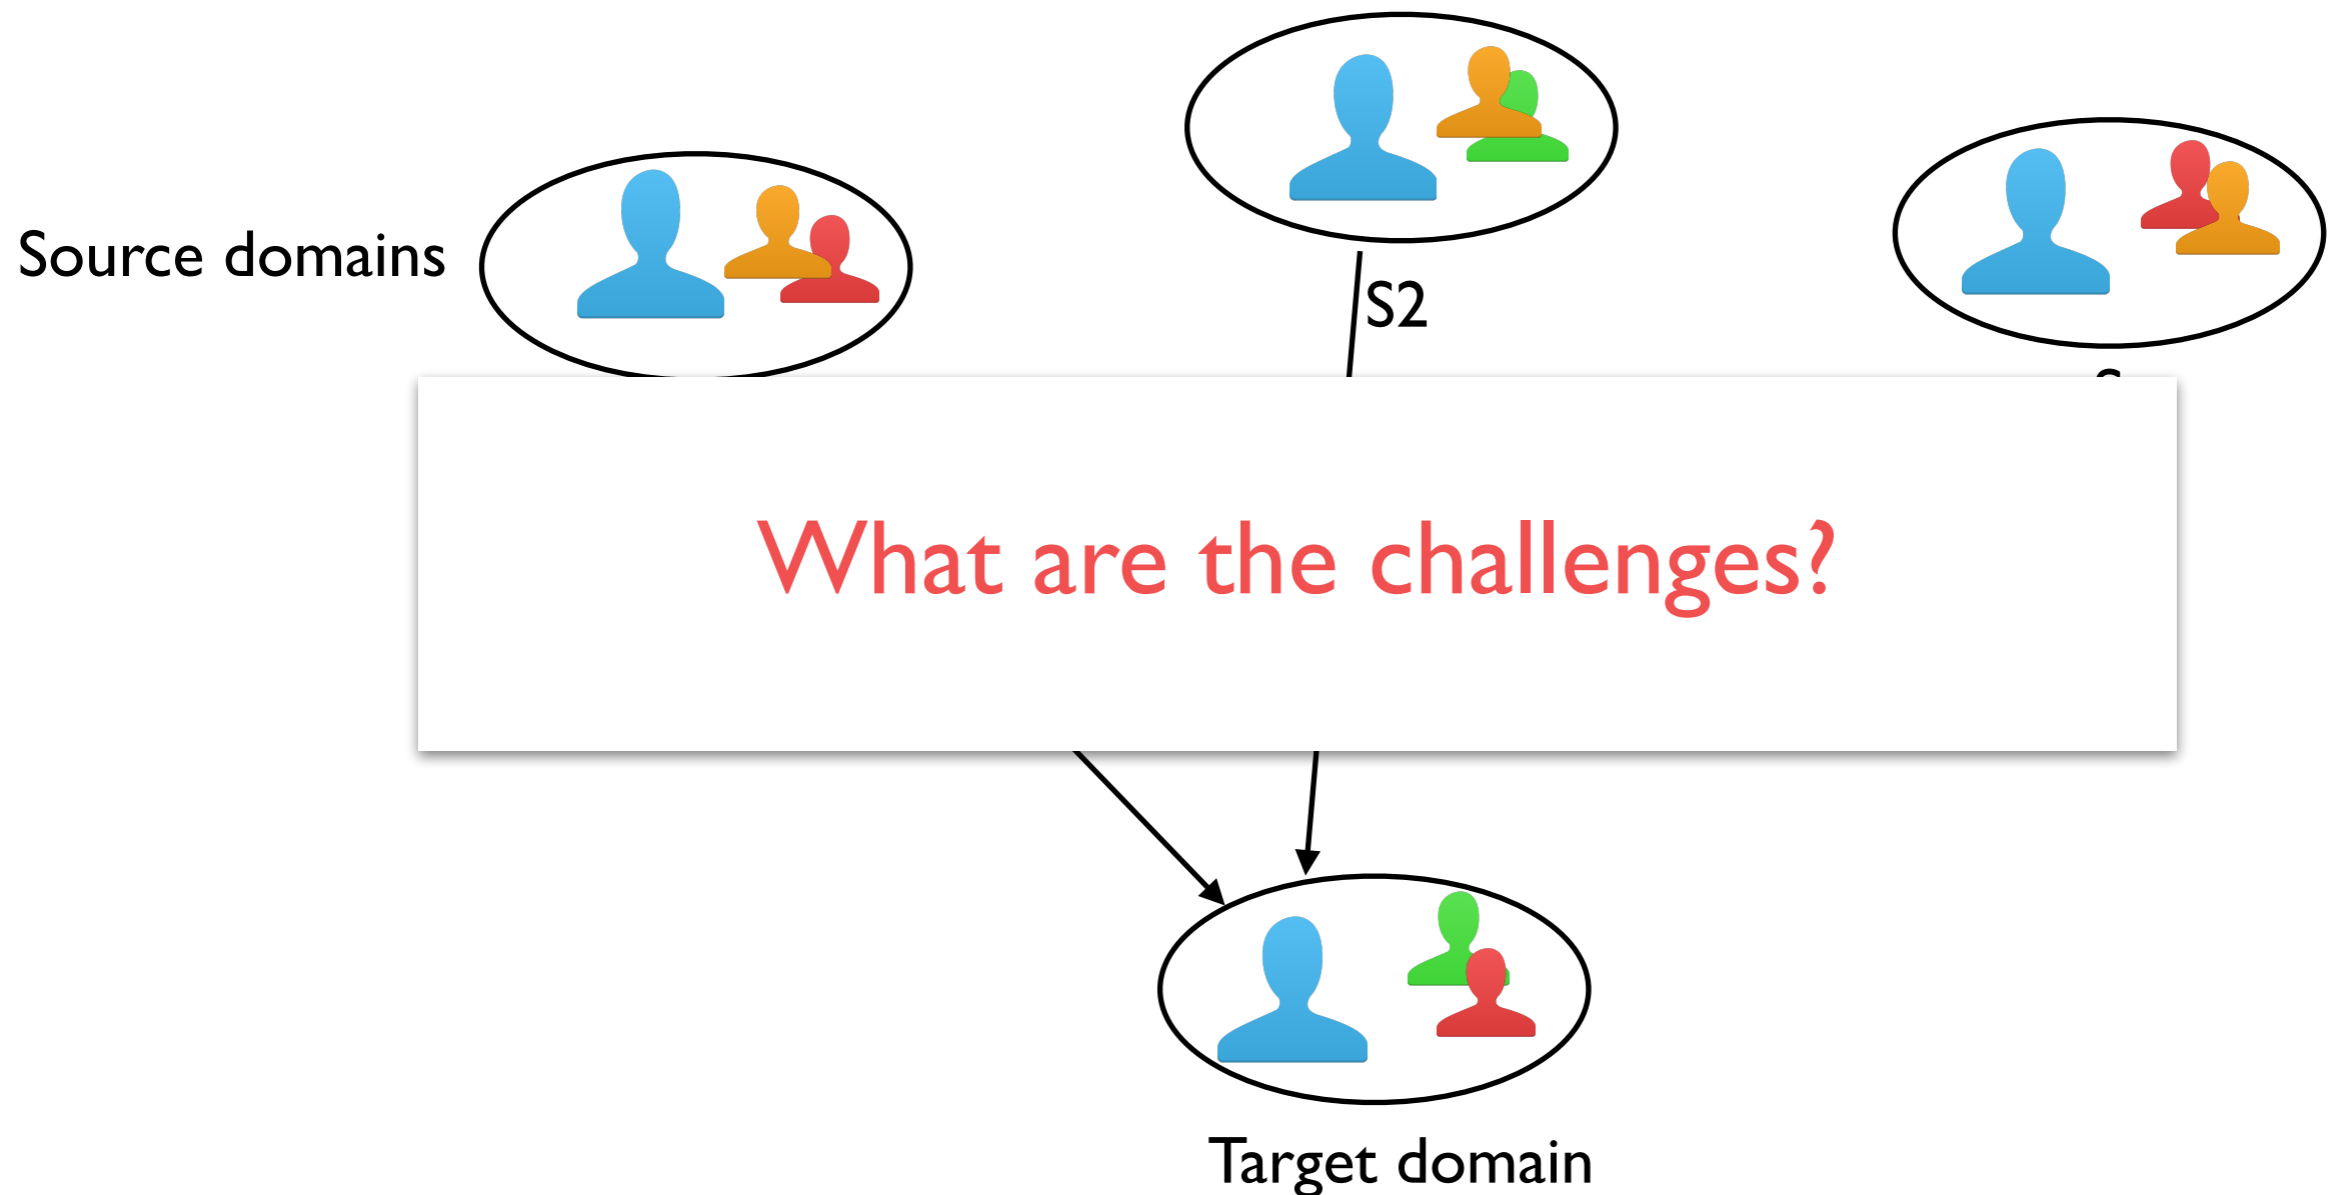
This talk

1. Potential for inter-domain trust transfer
2. Inter-domain trust transfer framework
3. Leverage inter-domain trust transfer for identity curation


Inter-domain trust transfer framework




Inter-domain trust transfer framework



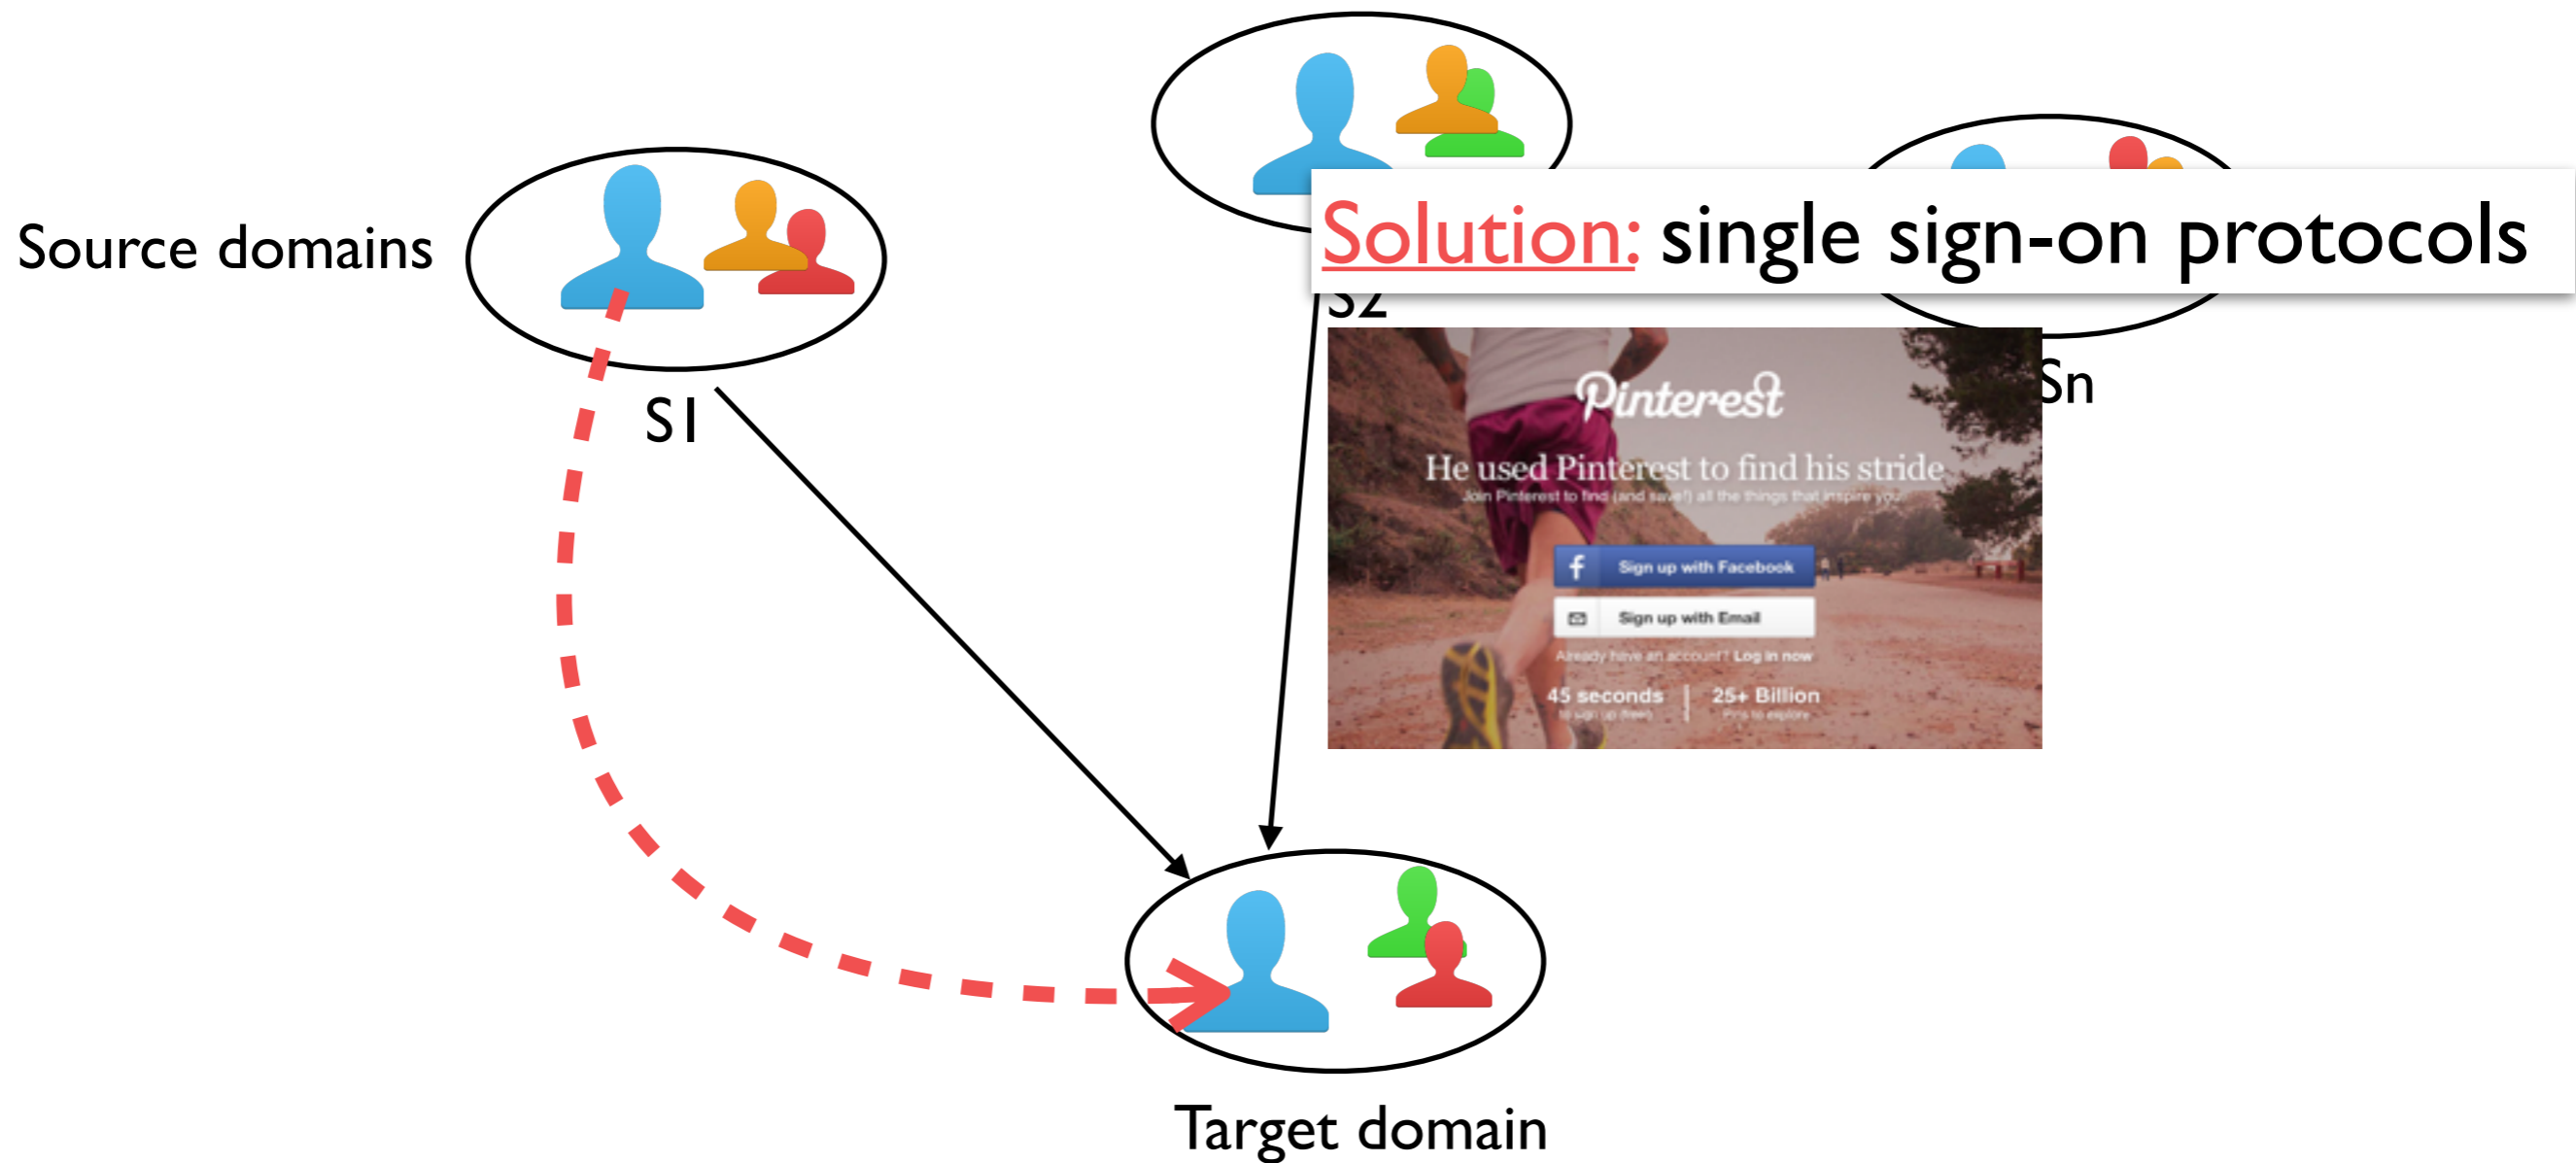
I. How to link the matching identities of a user?



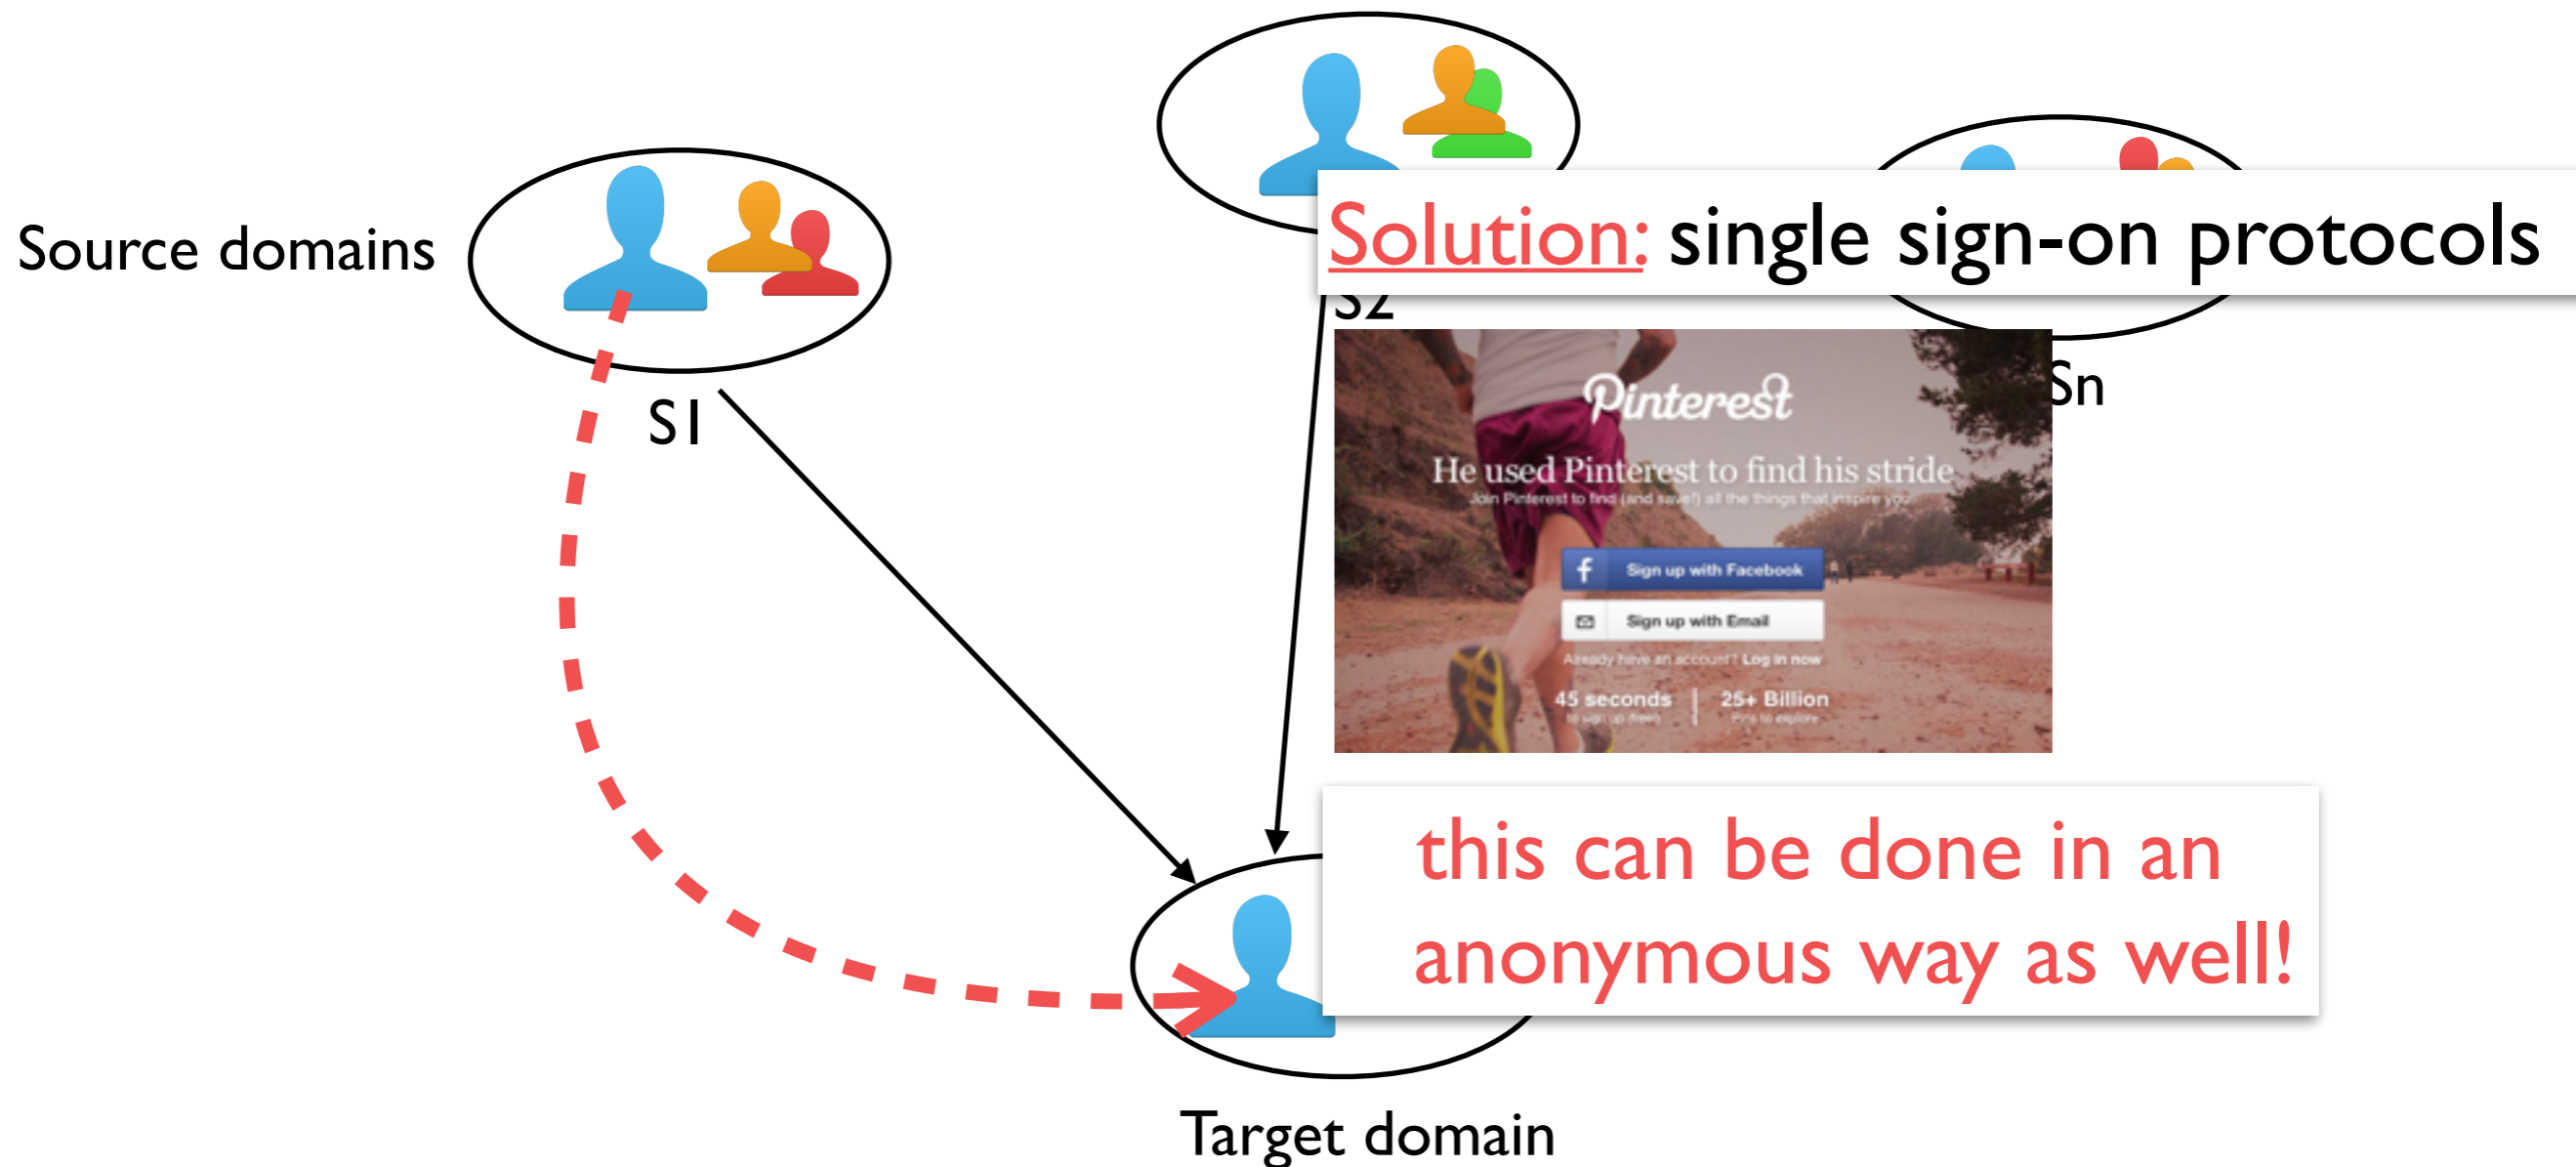
I. How to link the matching identities of a user?



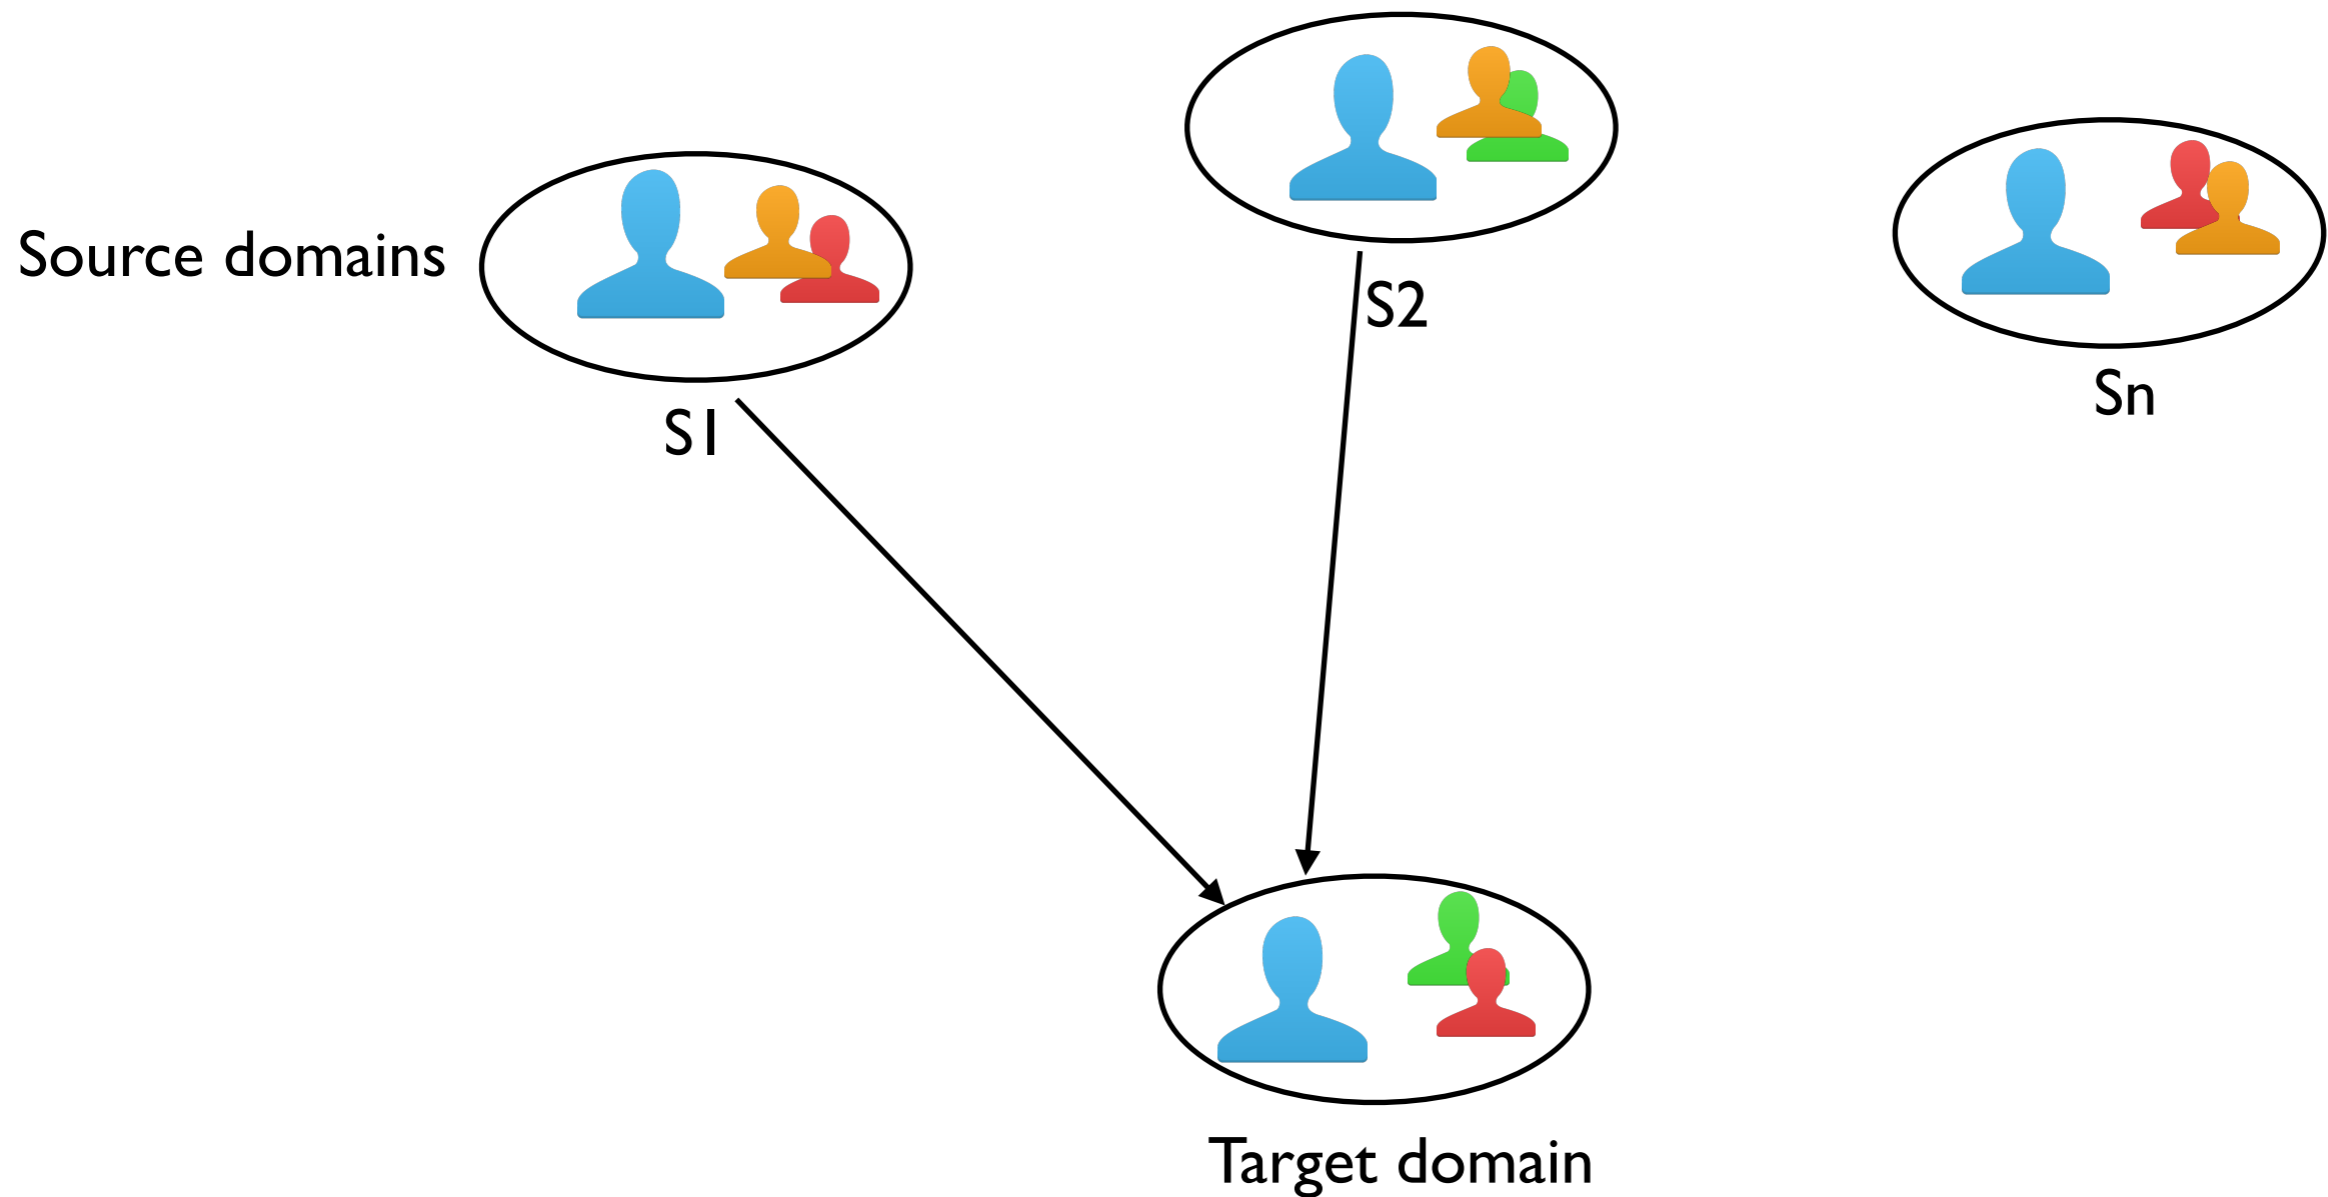
I. How to link the matching identities of a user?




I. How to link the matching identities of a user?



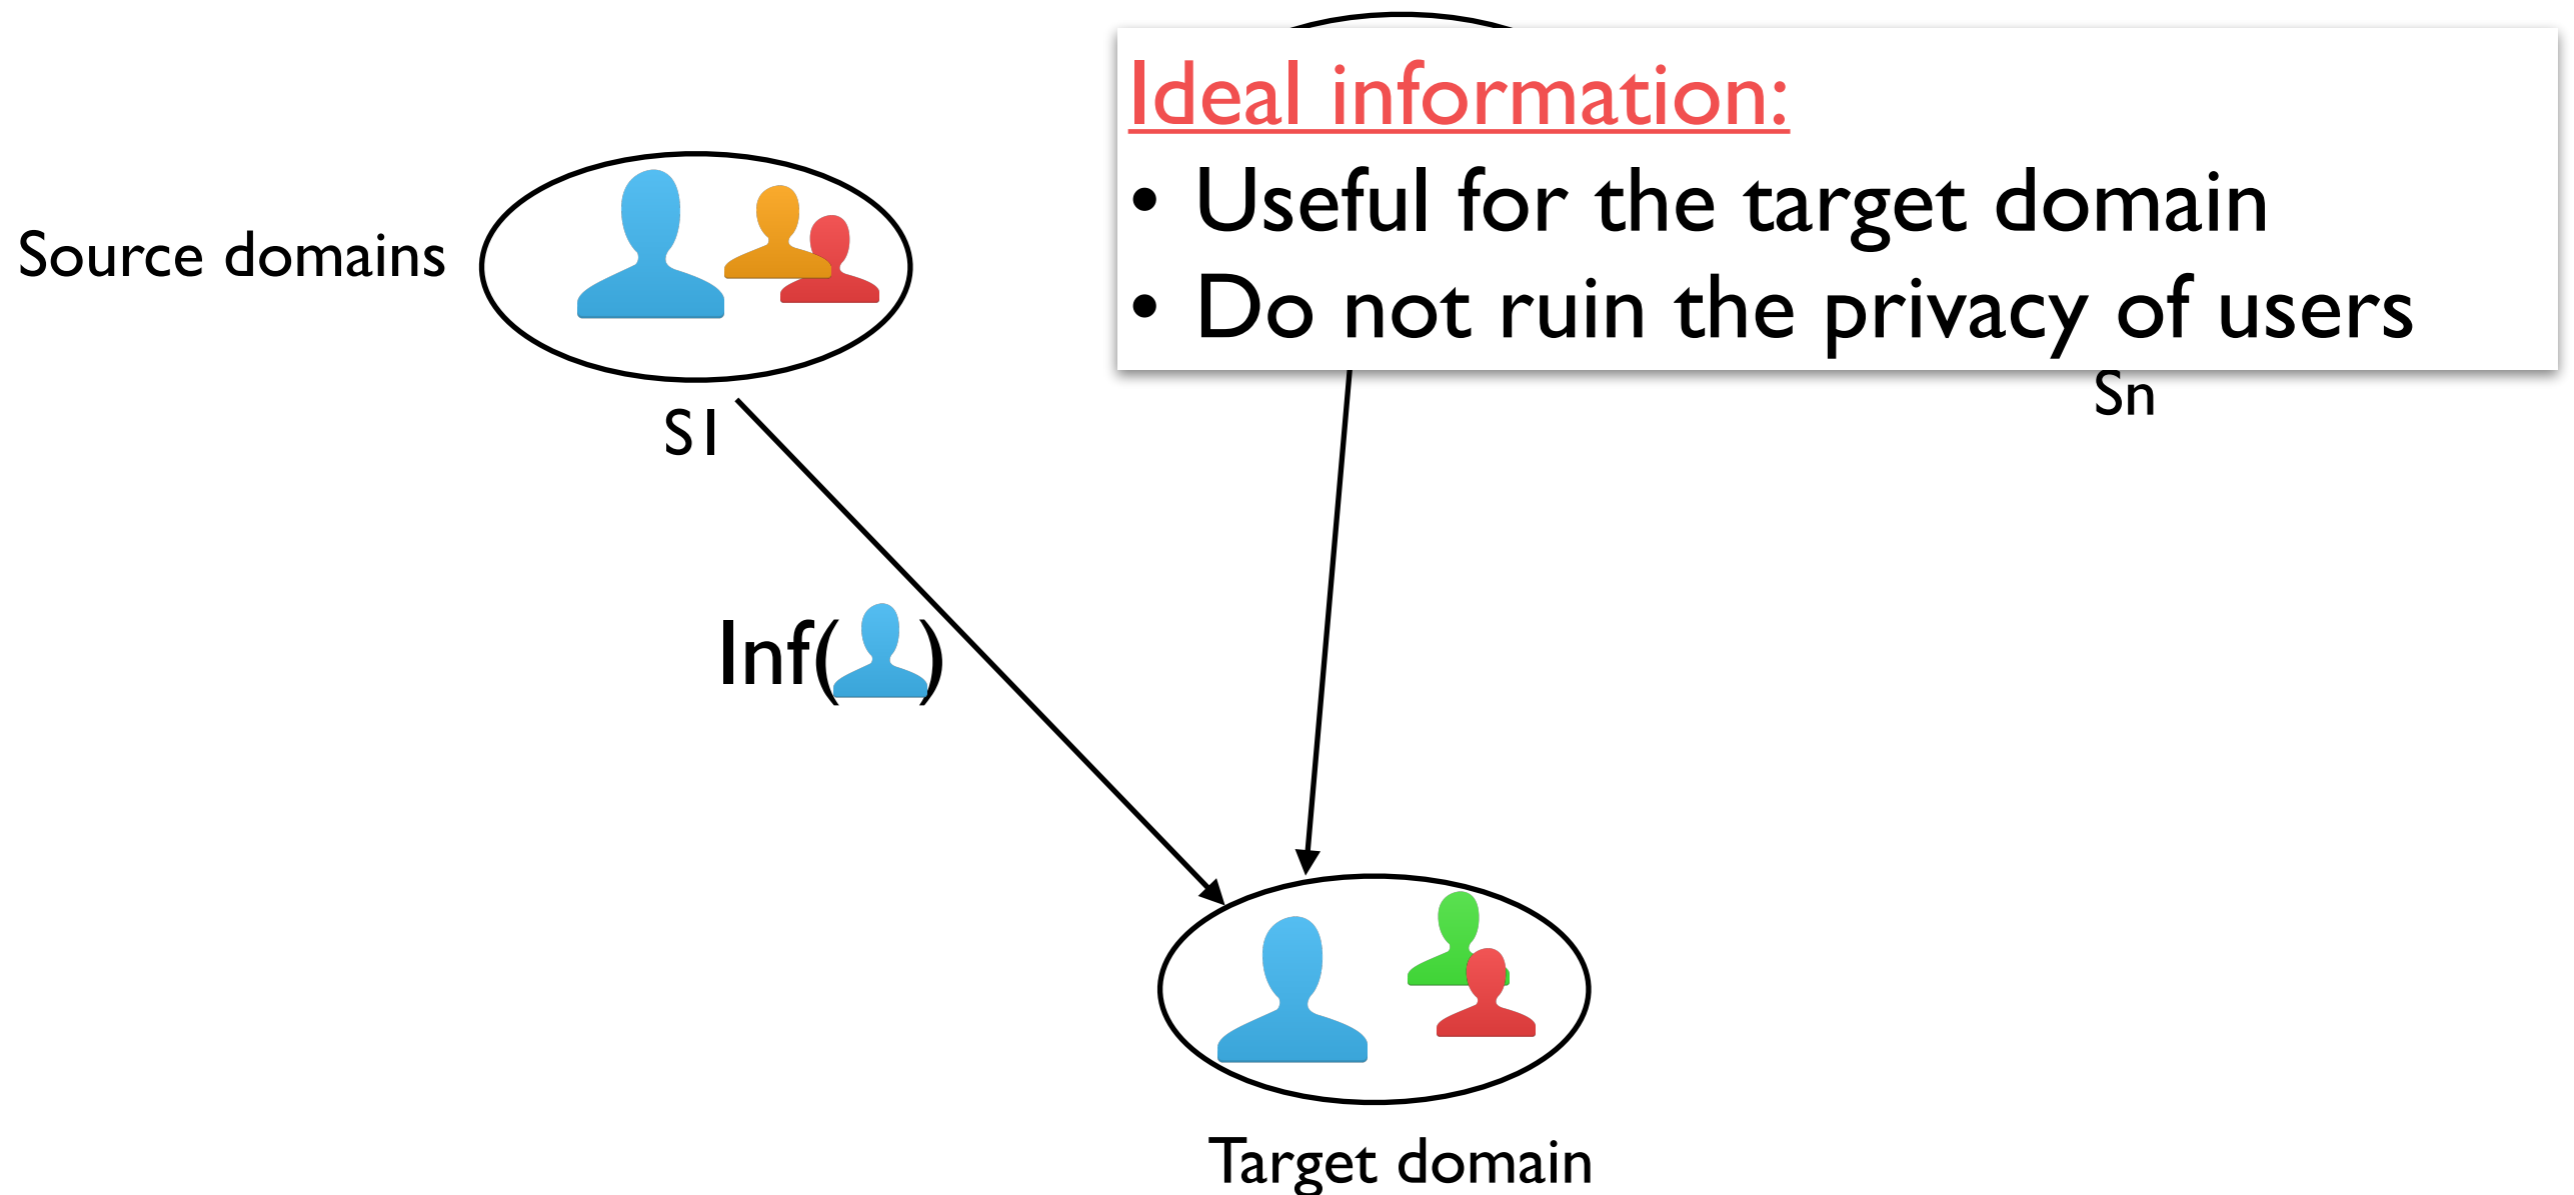
2. What information to export?



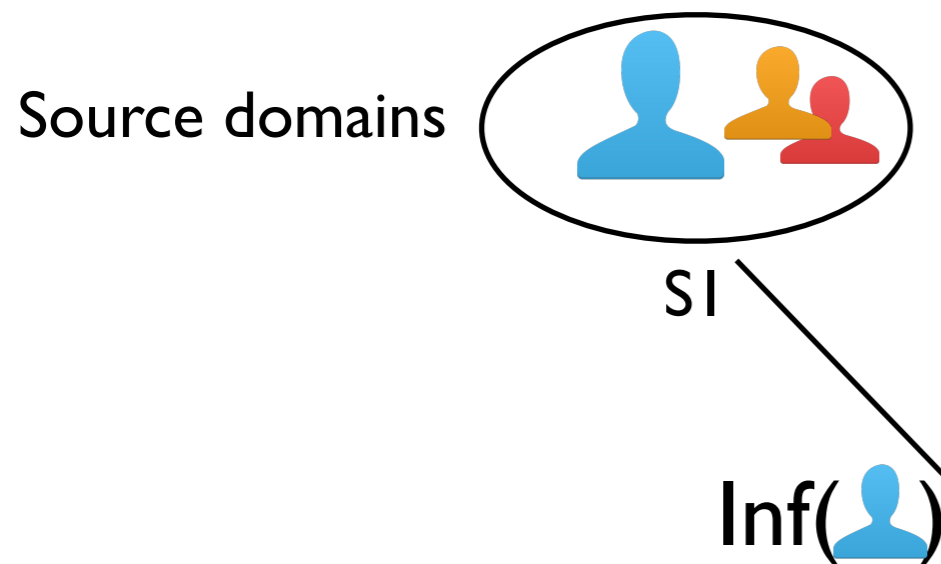
2. What information to export?



2. What information to export?



2. What information to export?



Ideal information:

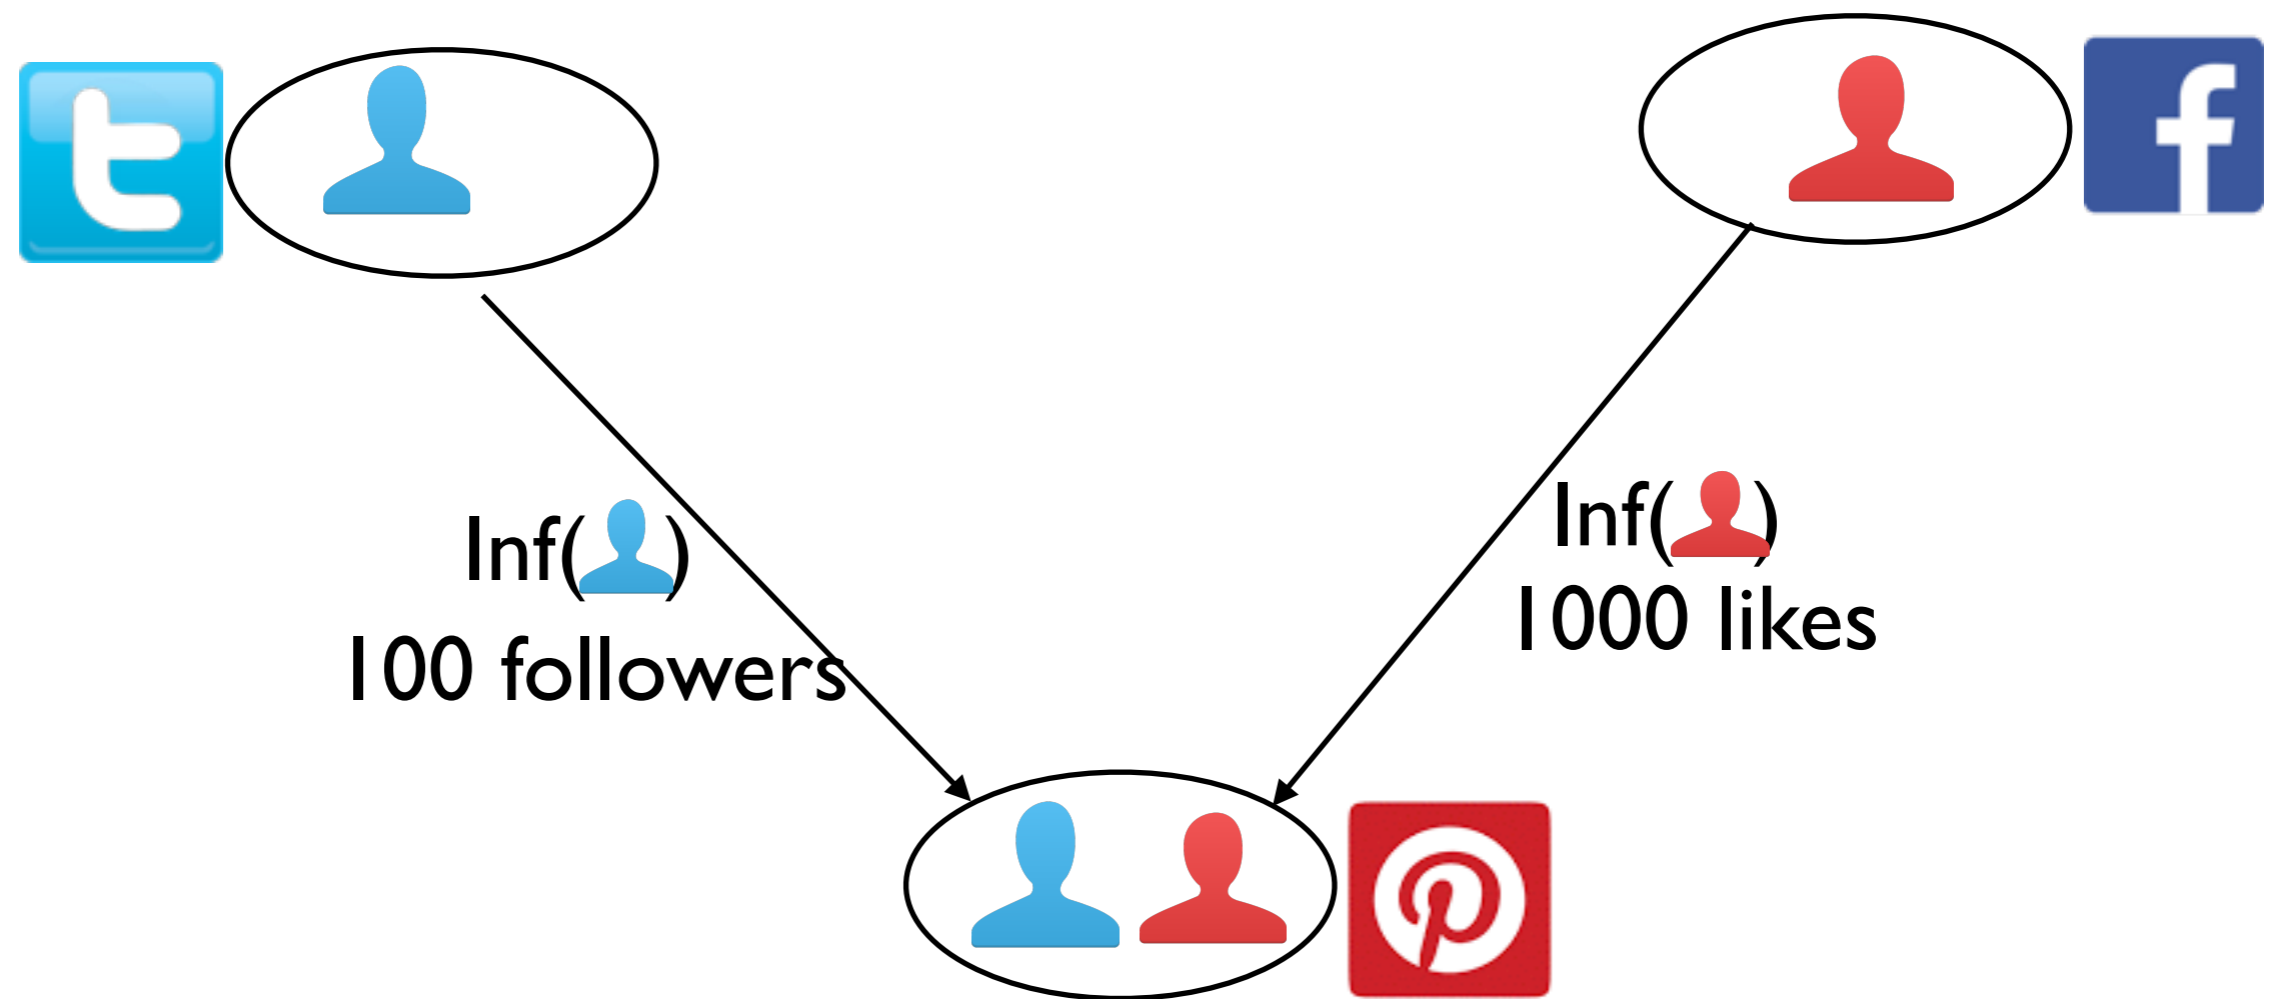
- Useful for the target domain
- Do not ruin the privacy of users

S_n

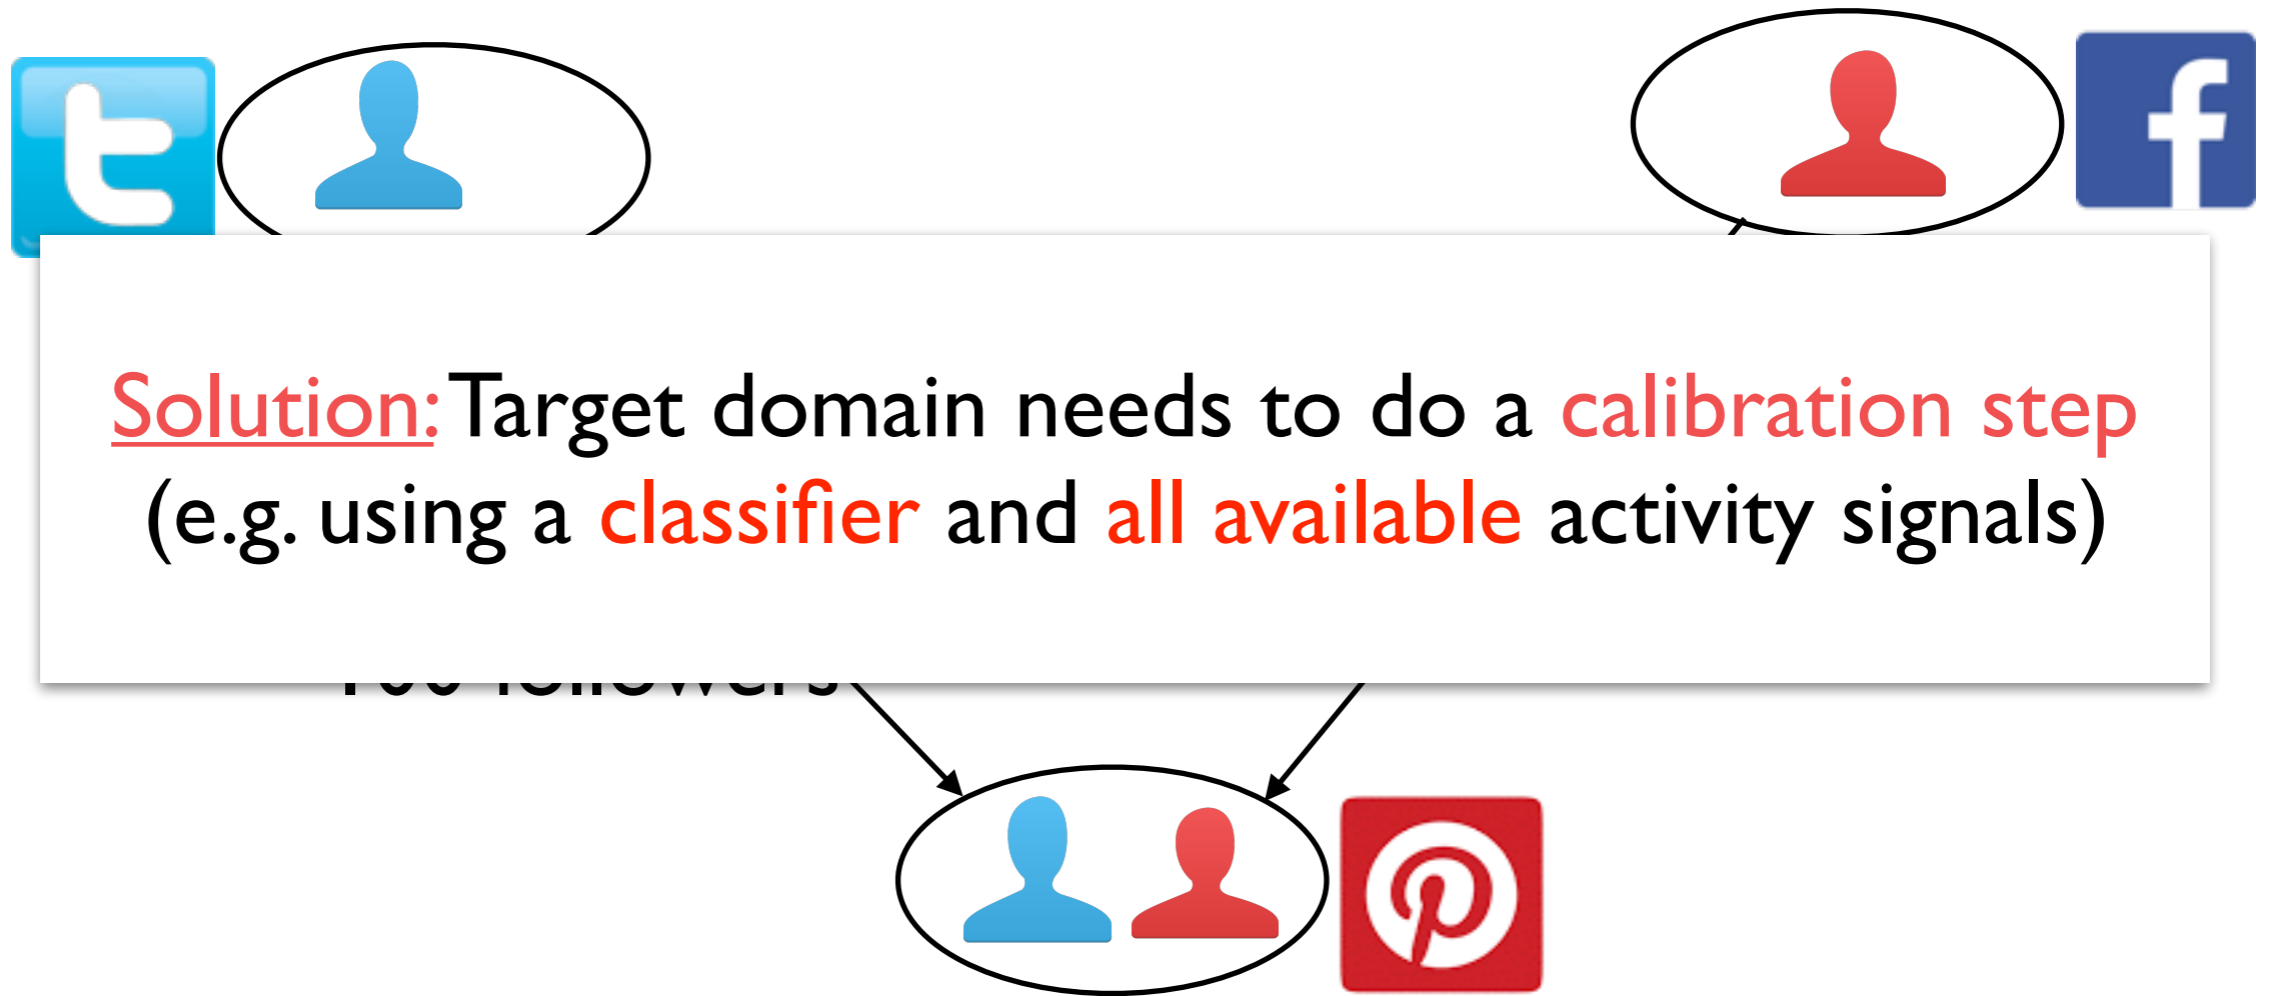
Solution:

- Each source domain can choose the information it wants to transfer
- The source domain can ask the permission of the user (e.g., OAuth)
- Break correlation with possible sensitive information (e.g., l-diversity, t-closeness)

3. How to interpret and combine the information?



3. How to interpret and combine the information?



This talk

1. Potential for inter-domain trust transfer
2. Inter-domain trust transfer framework
3. Leverage inter-domain trust transfer for identity curation

This talk

1. Potential for inter-domain trust transfer
2. Inter-domain trust transfer framework
3. Leverage inter-domain trust transfer for identity curation

Identity curation

Curated set of identities: a set of identities with high probability to be trustworthy

Why curate?

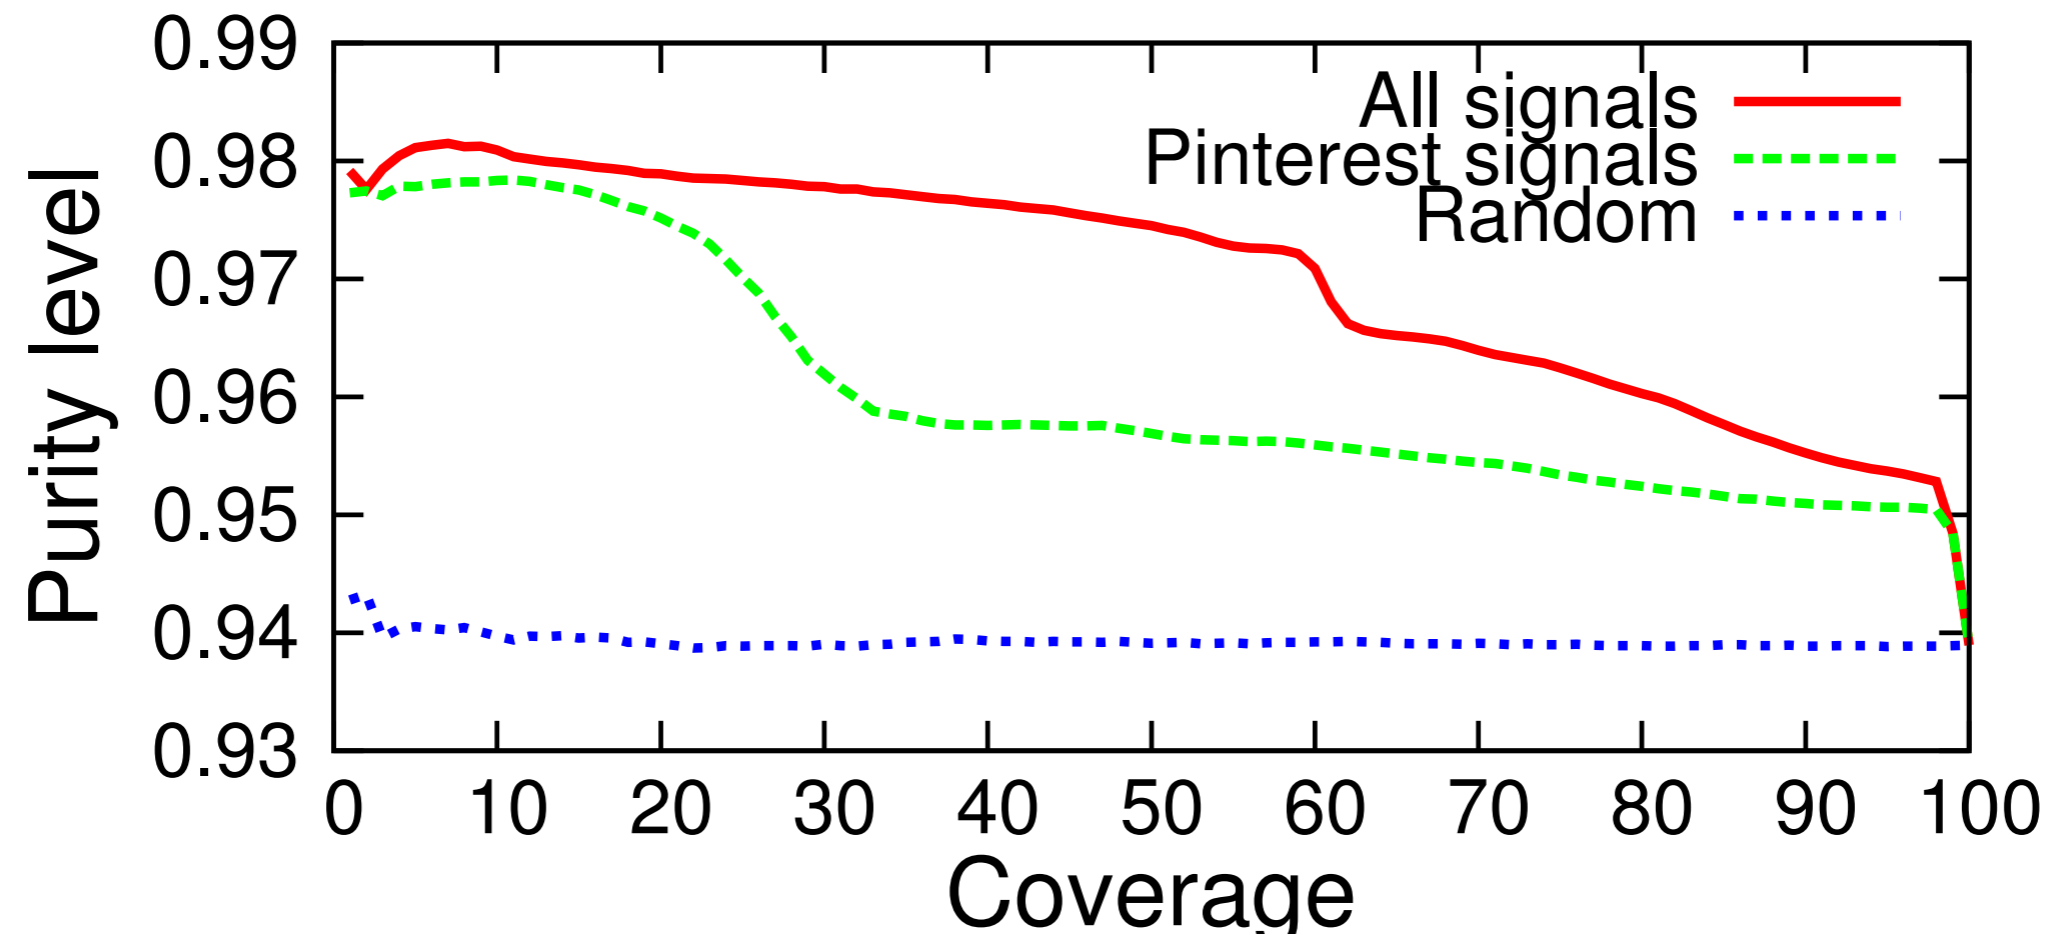
- I. Early access to elevated privileges
- II. Sybil-resilient content recommendation

Question:

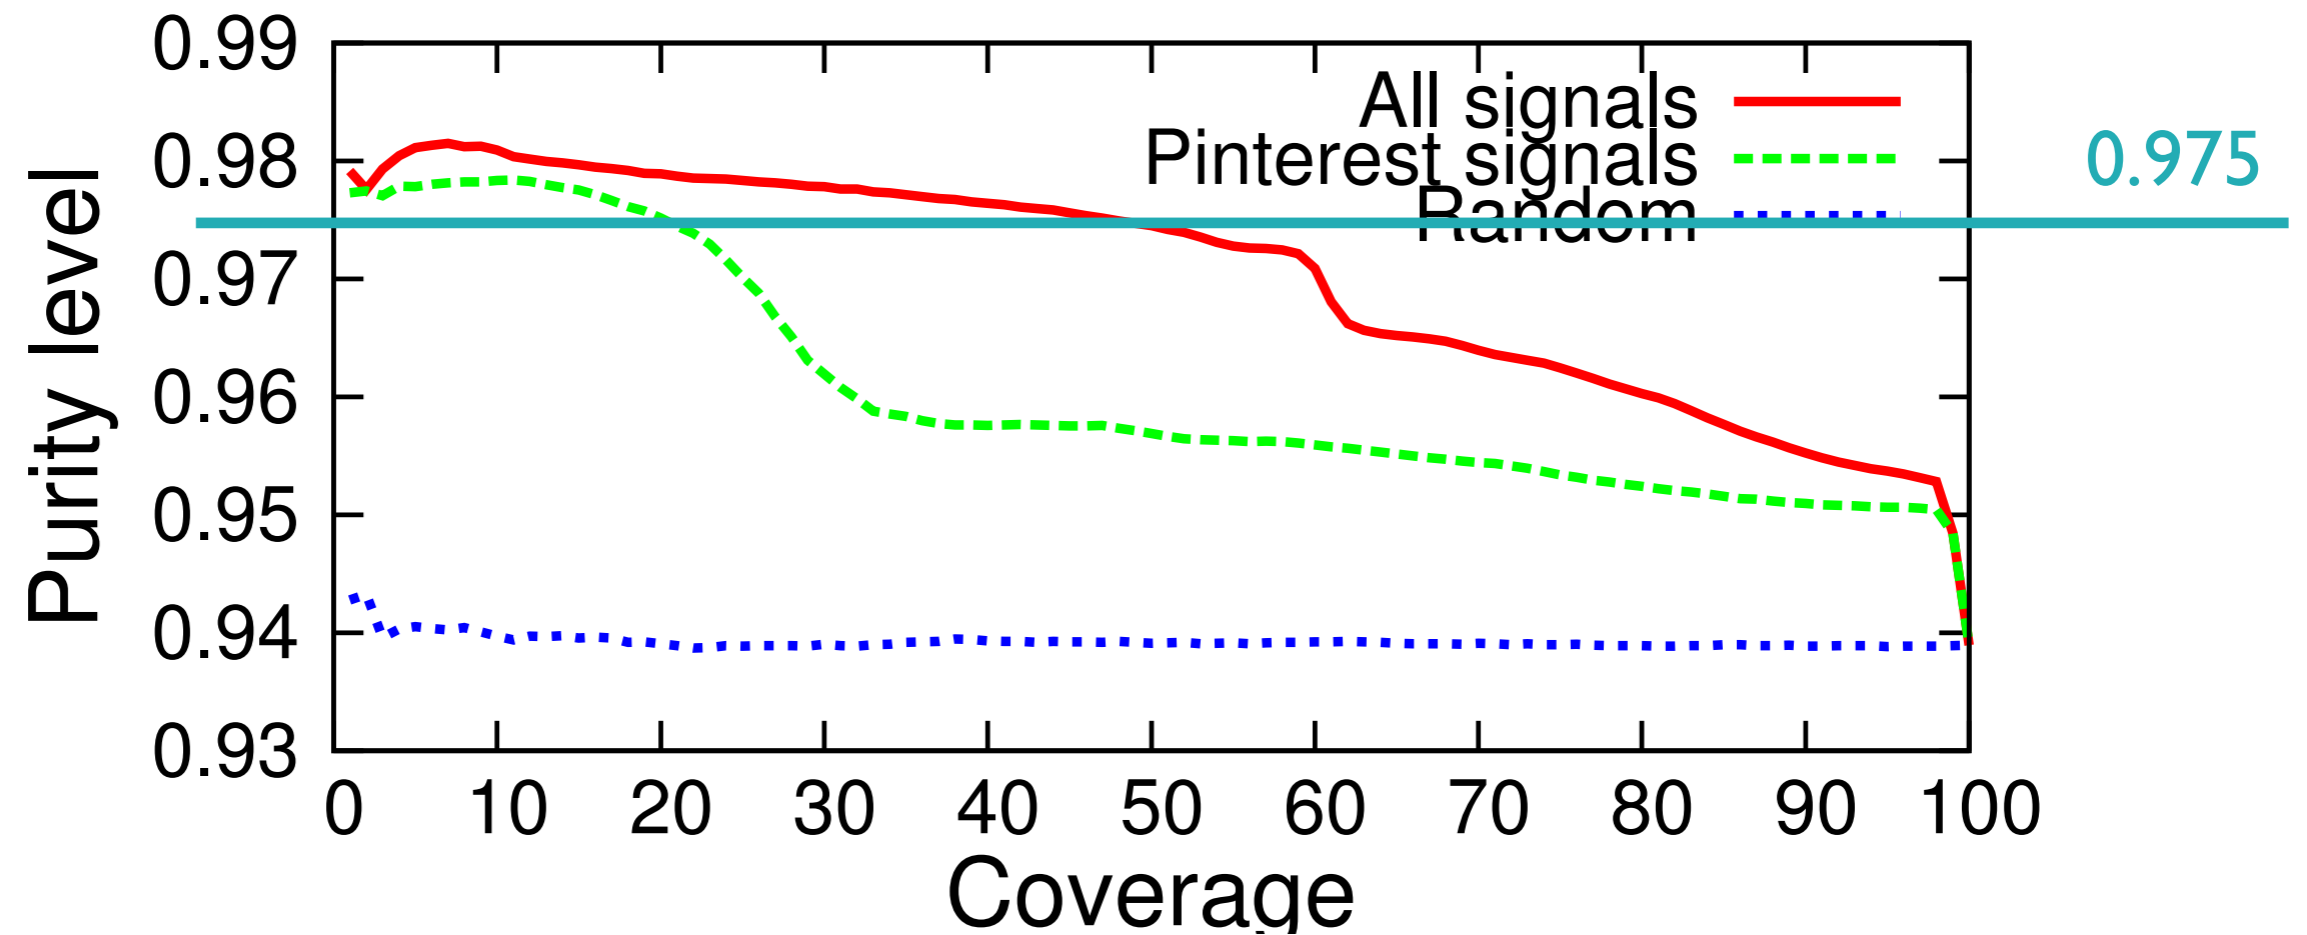
Can activity signals from Facebook/Twitter help Pinterest:

- Curate more identities?
- Curate identities early-on?

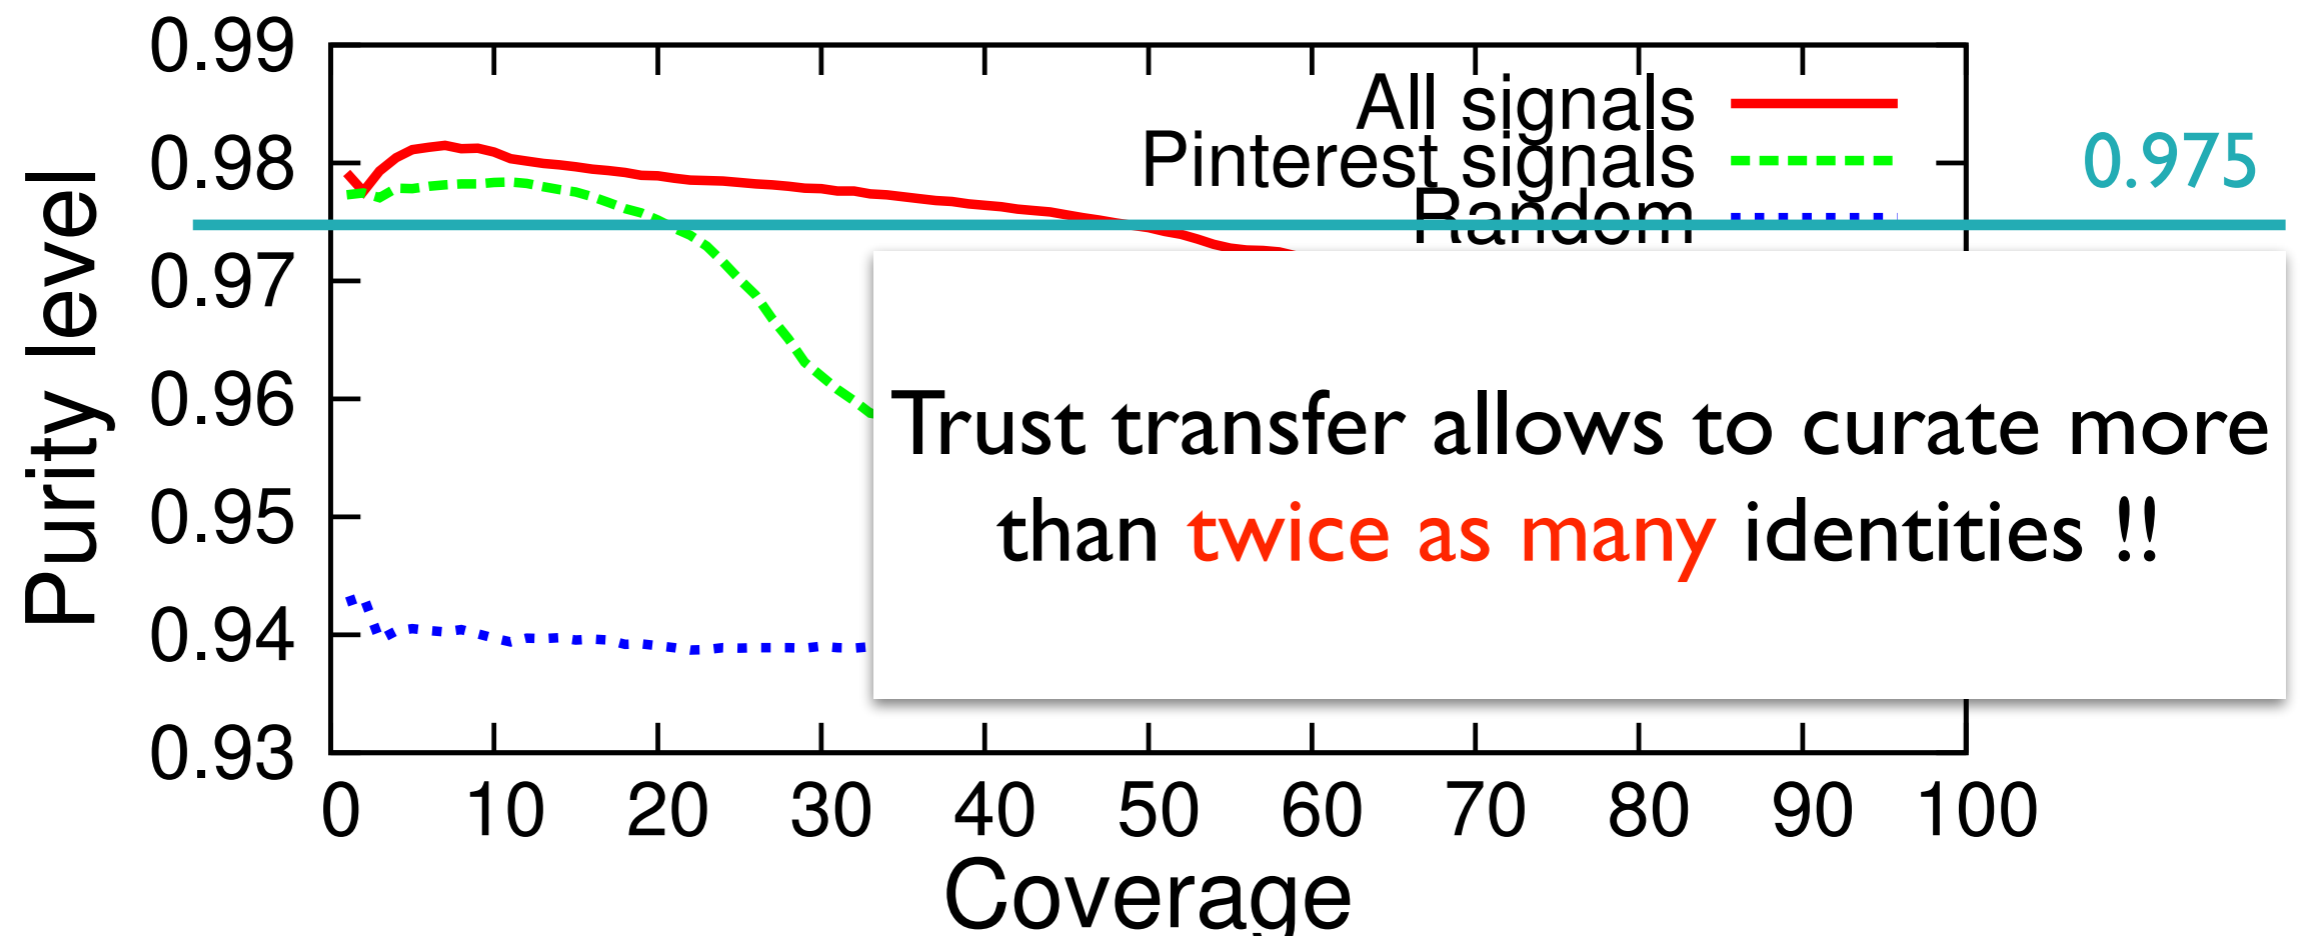
Pinterest can curate more identities




Pinterest can curate more identities



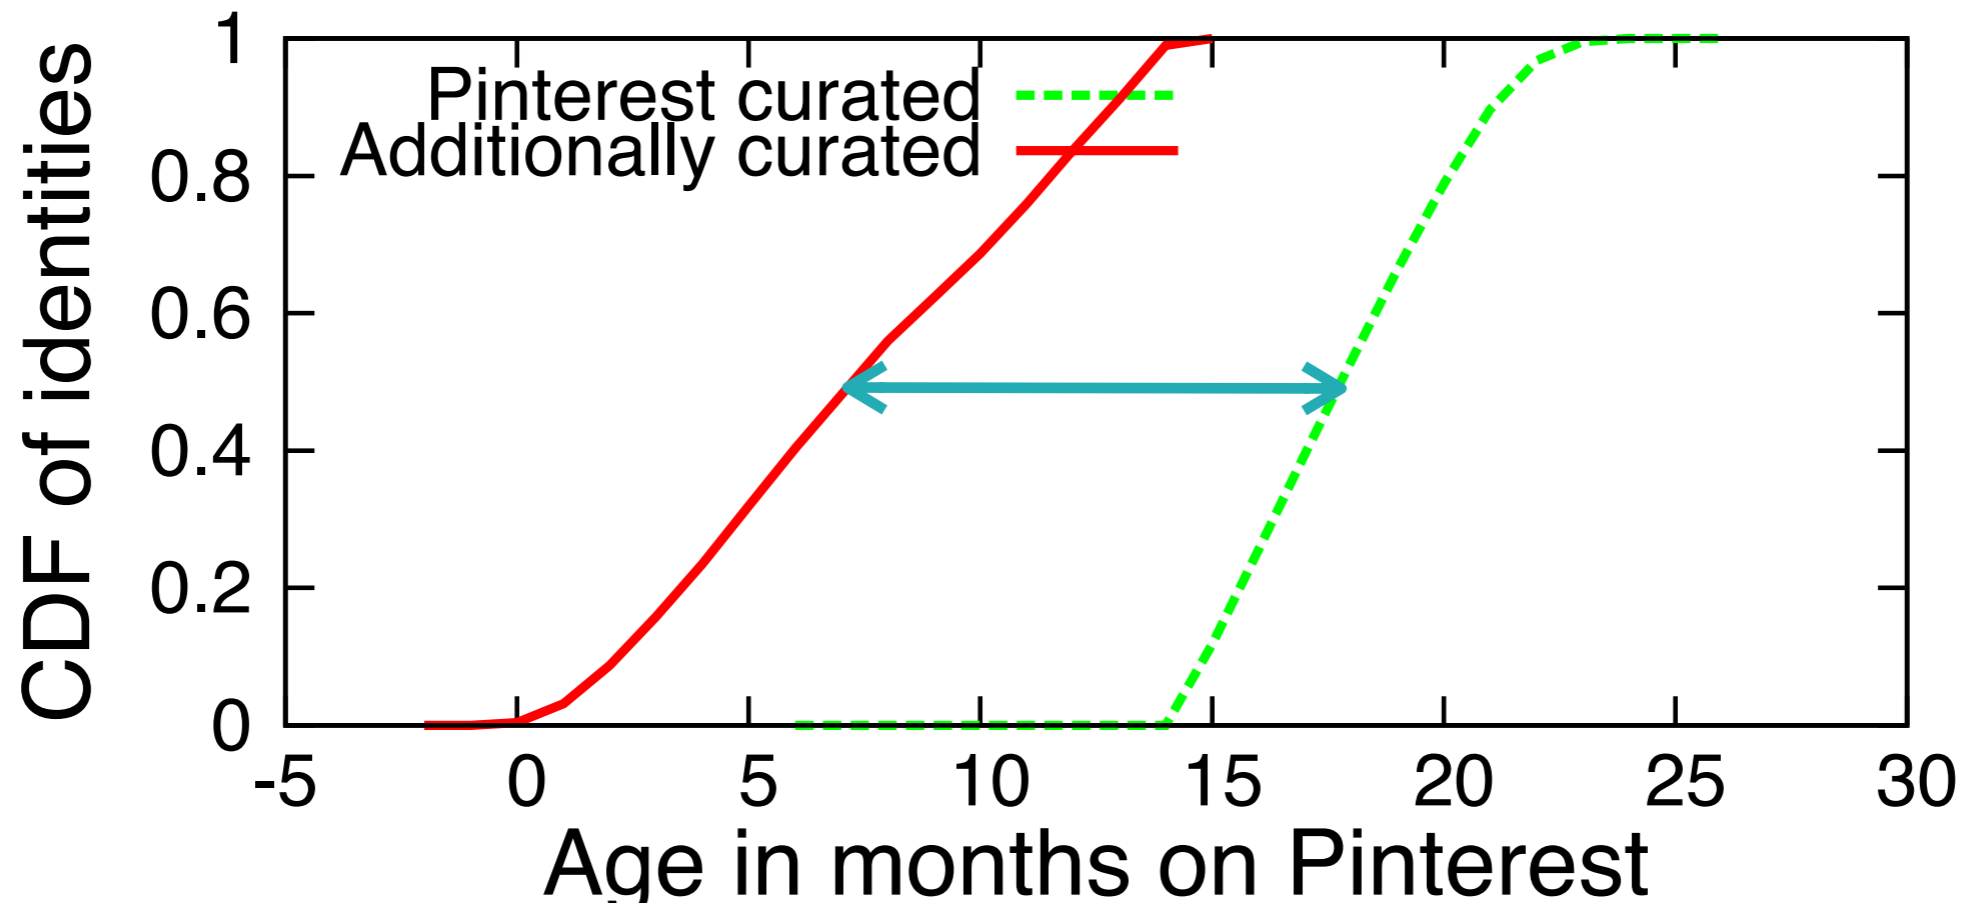
Pinterest can curate more identities



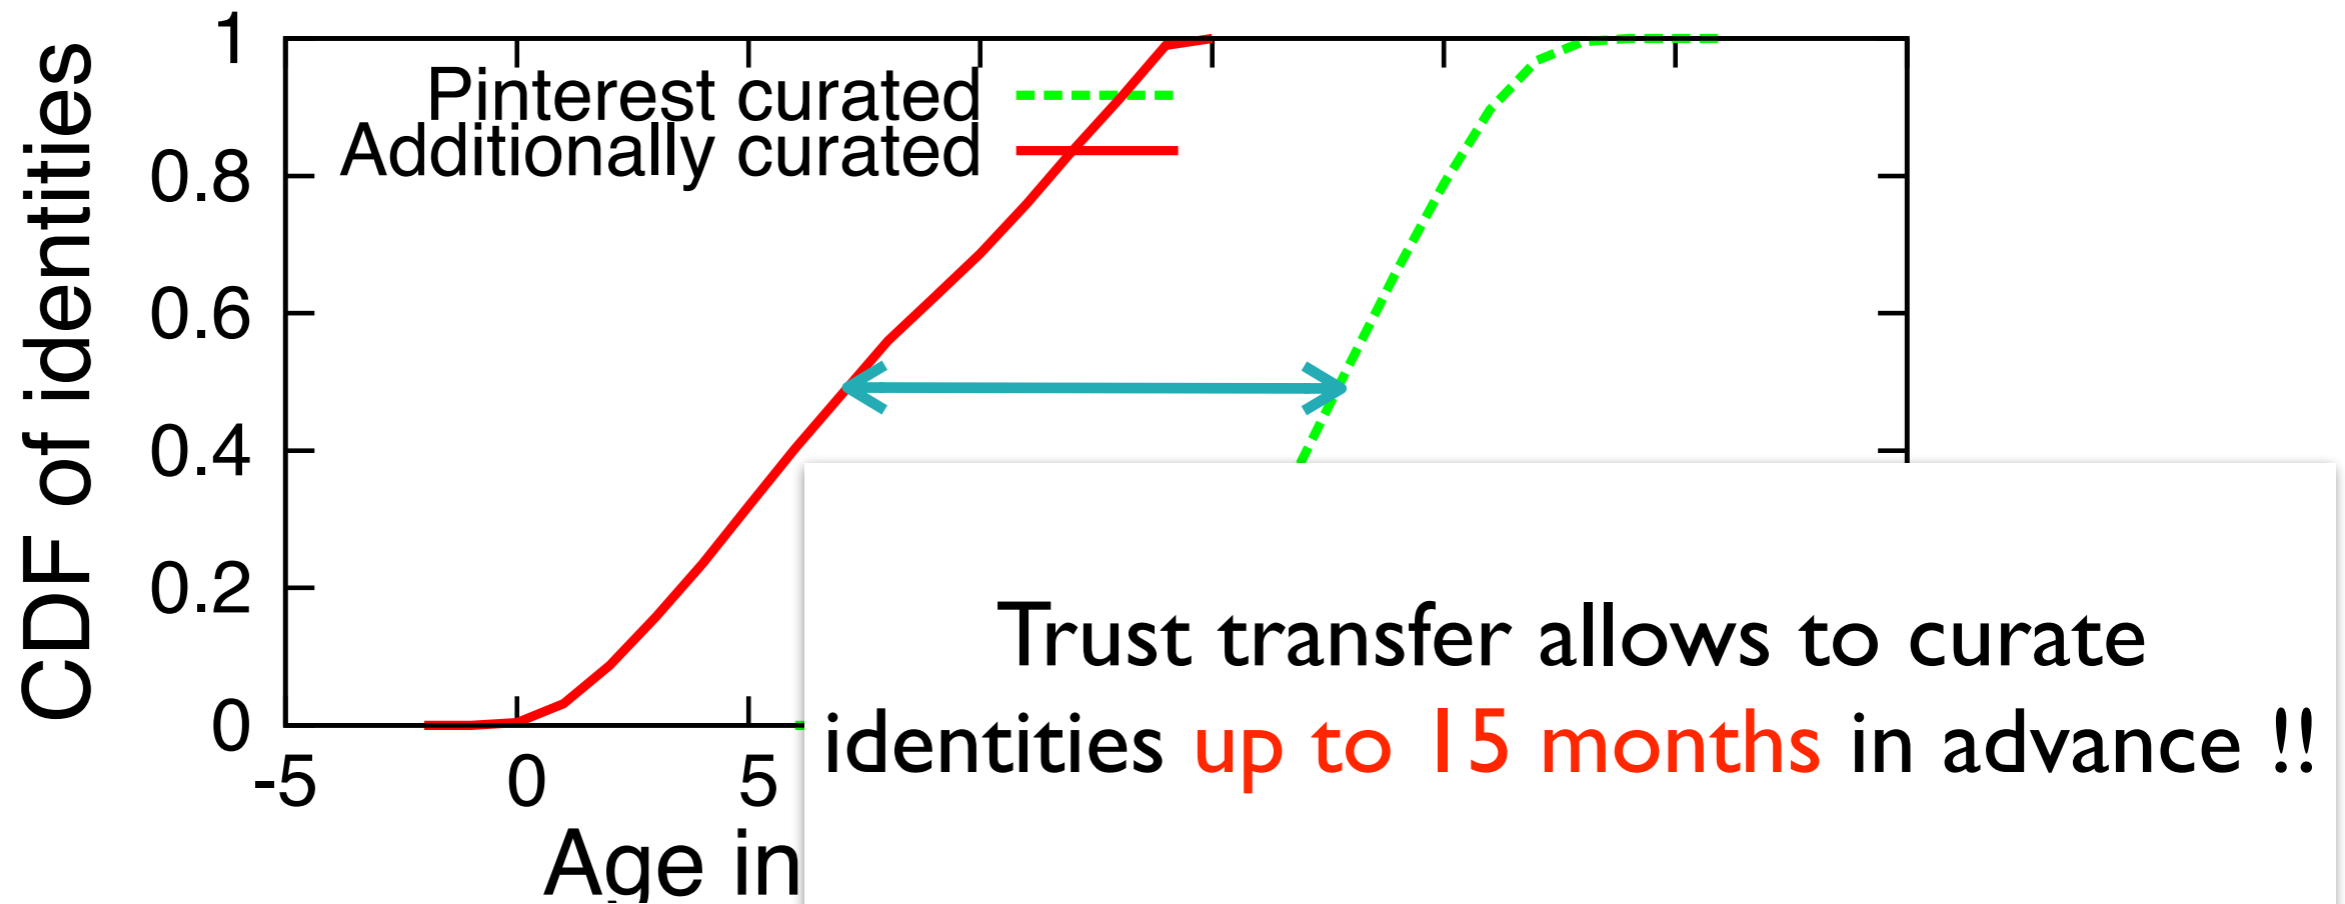
Pinterest can curate identities early-on



Pinterest can curate identities early-on



Pinterest can curate identities early-on



Conclusions

- Lack of external trust certificates for weak identities could be mitigated by user's weak identities on other domains
 - Users can **keep their identities anonymous!**
- Inter-domain trust transfer framework
 - Low deployment overheads
 - Without significant loss of privacy
- Evaluation on real-world domains is very promising!
 - Even simple activity signals go a long way