UNIVERSITY OF Southampton
School of Electronics and Computer Science

Data provenance in a distributed calculus

Vladimiro Sassone

work in progress with A. Francalanza, J. Rathke, and I. Souilah

Motivation

- (Meta)data is almost entirely neglected in the process calculi literature
- Track data provenance both for its important applications and as an challenging exercise in modelling (meta)data. We aim at simplicity:
  - data annotations representing provenance
  - structure, interpretation and management of provenance information
  - provenance tracking
- Provenance-based security (aspects: trust + data confidentiality and privacy)
  - Example: photography competition
- The overall ambition is to underpin practical development, like trust-policy languages and protocols, and provenance-middleware

Model features

- Two the central features of the basic model:
  1. values are annotated with their provenance
  2. provenance is kept up-to-date as computation proceeds
- Focus on one particular kind of provenance information:
  - The principals that influenced a value, and how they did it

Provenance model

Annotated data

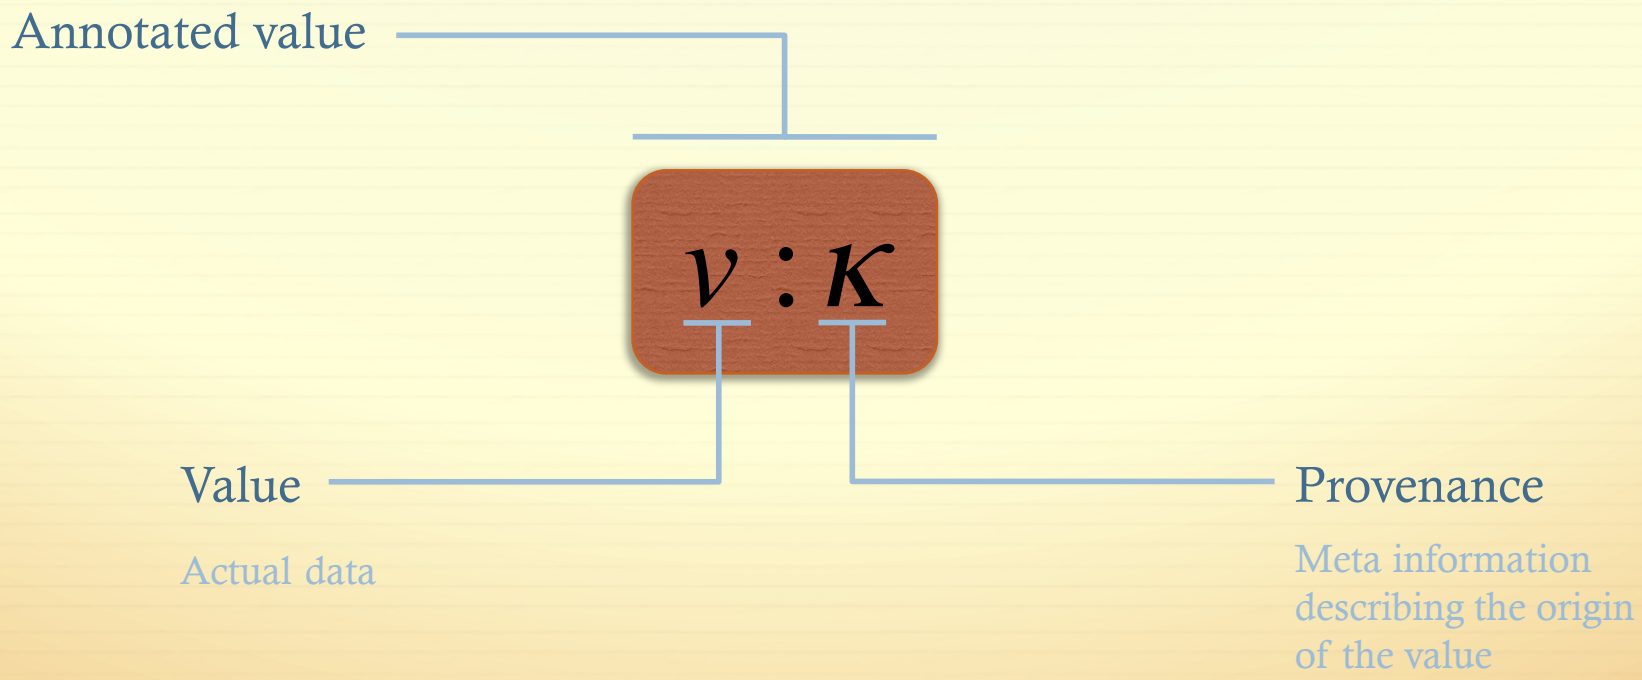

Provenance model

Structure and interpretation of provenance

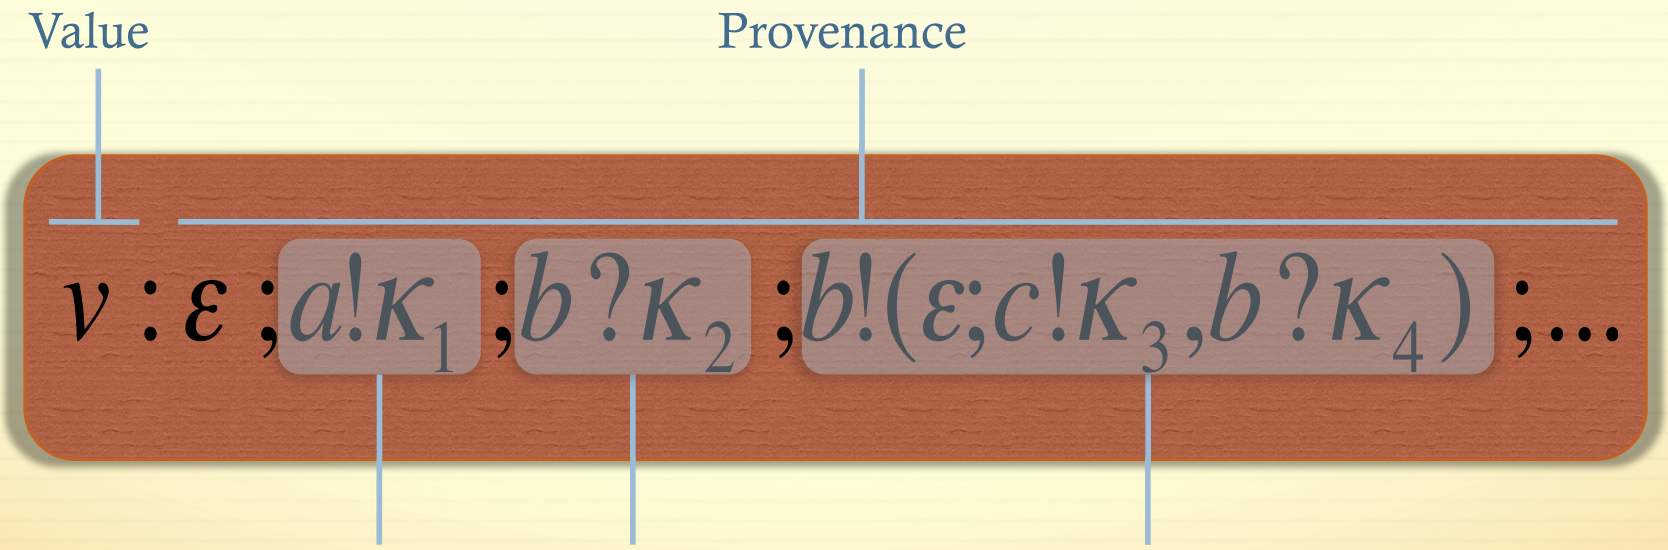

Provenance model

Structure and interpretation of provenance

ε (empty provenance) denotes value v originated here

Provenance model

Structure and interpretation of provenance

ε (empty provenance) denotes value v originated here

It was sent by a on a channel with provenance κ_1

Provenance model

Structure and interpretation of provenance

ε (empty provenance) denotes value v originated here

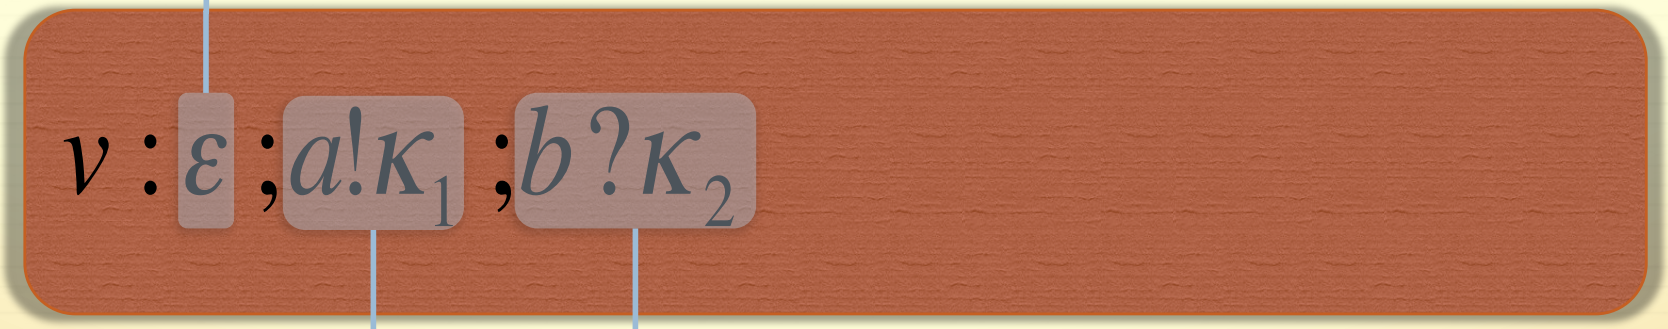

It was sent by a on a channel with provenance κ_1

Was then received by b on a channel with provenance κ_2

Provenance model

Structure and interpretation of provenance

ε (empty provenance) denotes value v originated here

And then sent by b on a channel that b received from c…

It was sent by a on a channel with provenance κ_1

Was then received by b on a channel with provenance κ_2

Provenance model

Provenance tracking

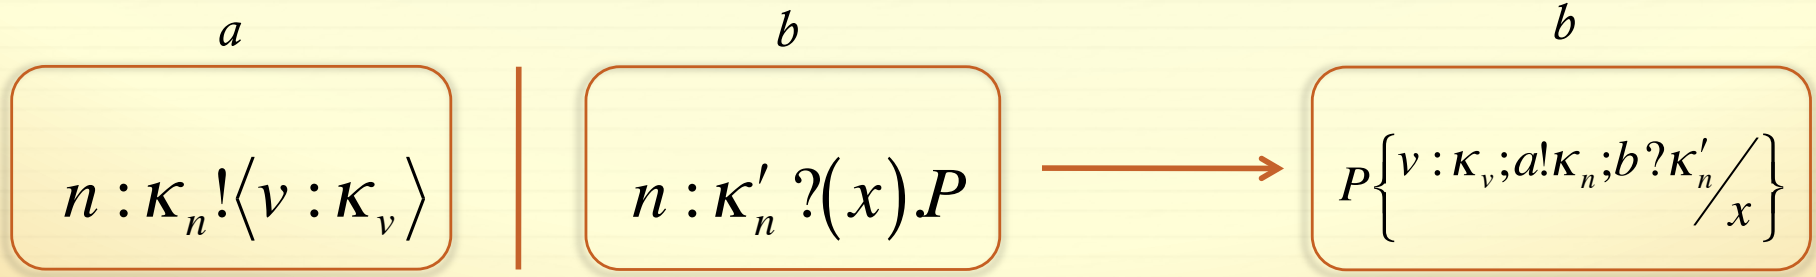

Provenance model

Provenance tracking

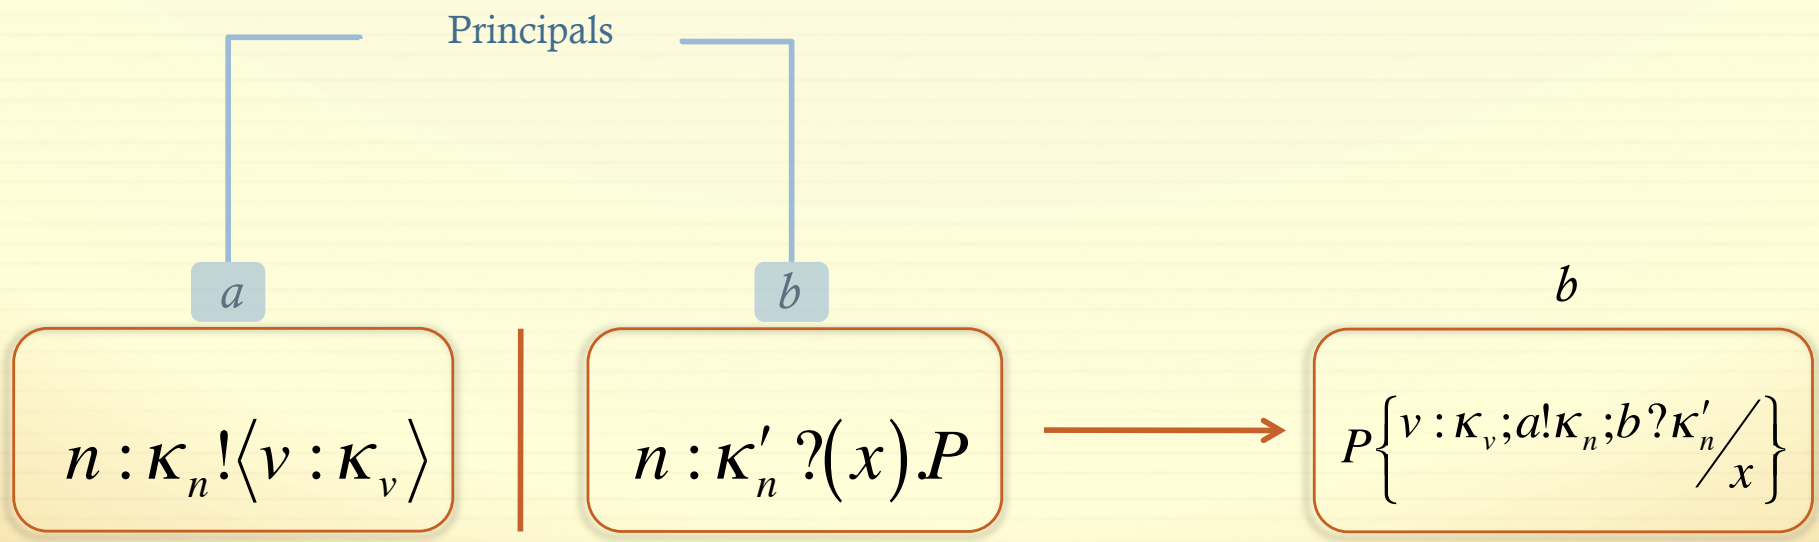

Provenance model

Provenance tracking

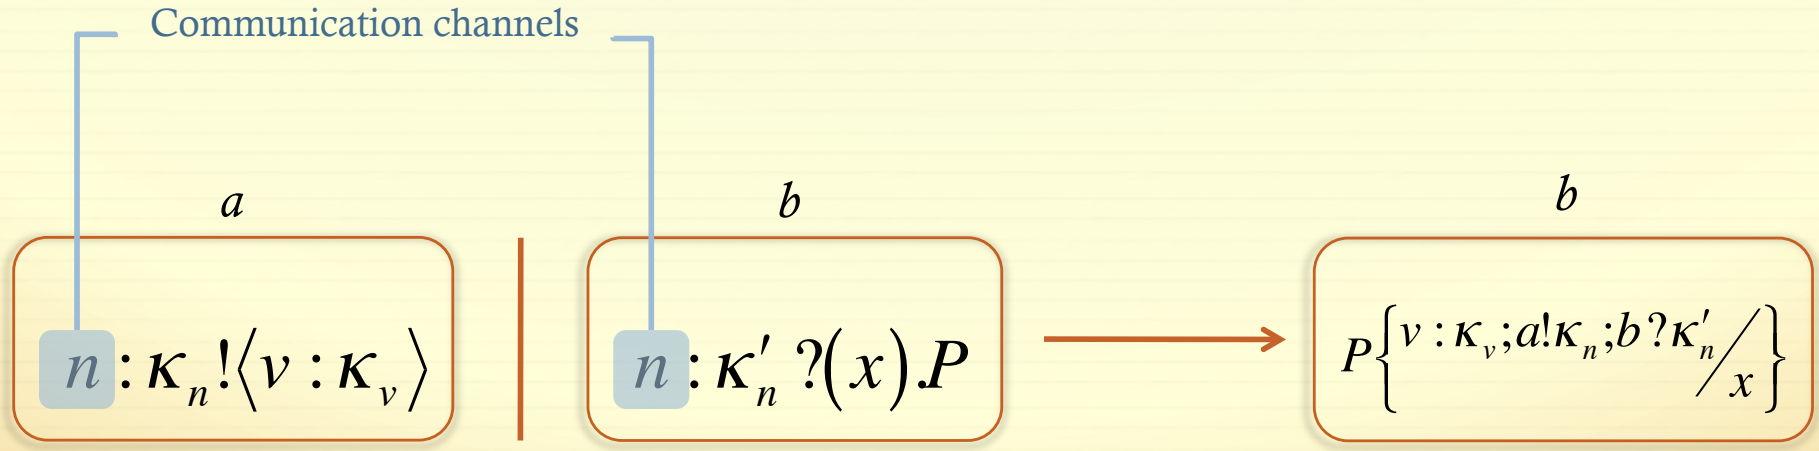

Provenance model

Provenance tracking

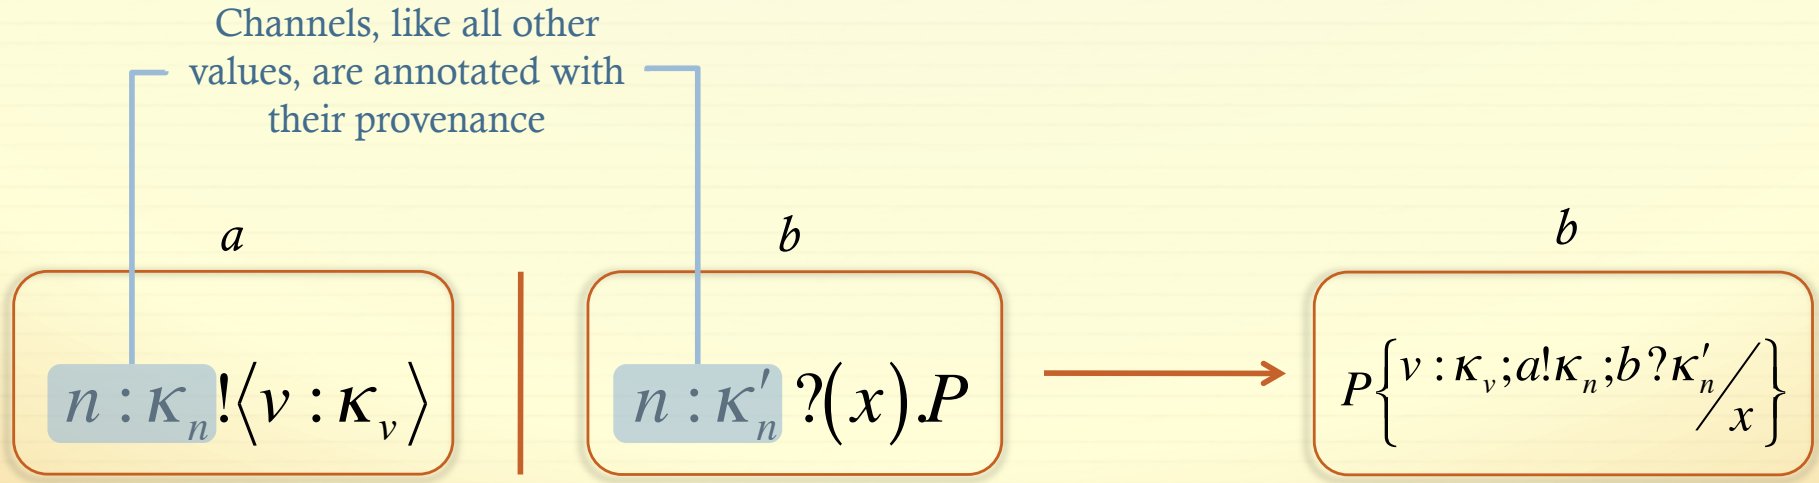

Provenance model

Provenance tracking

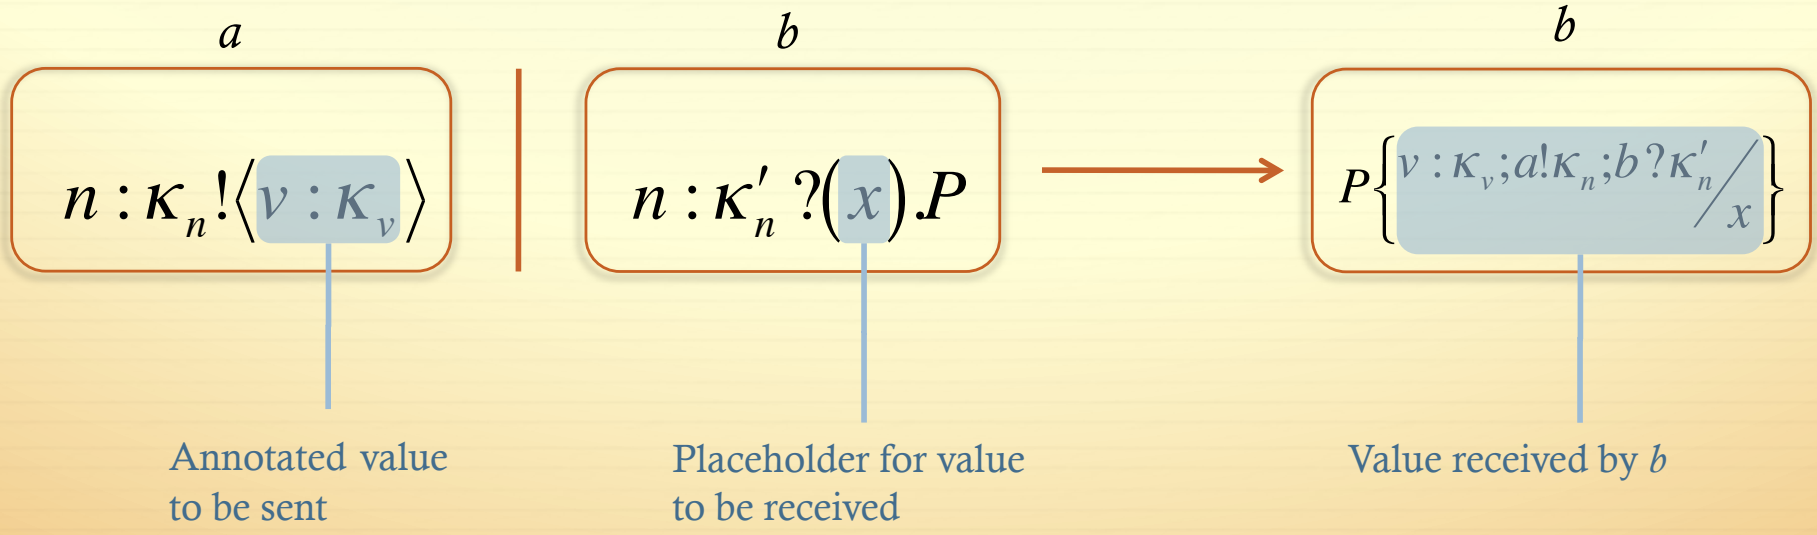

Provenance model

Provenance tracking

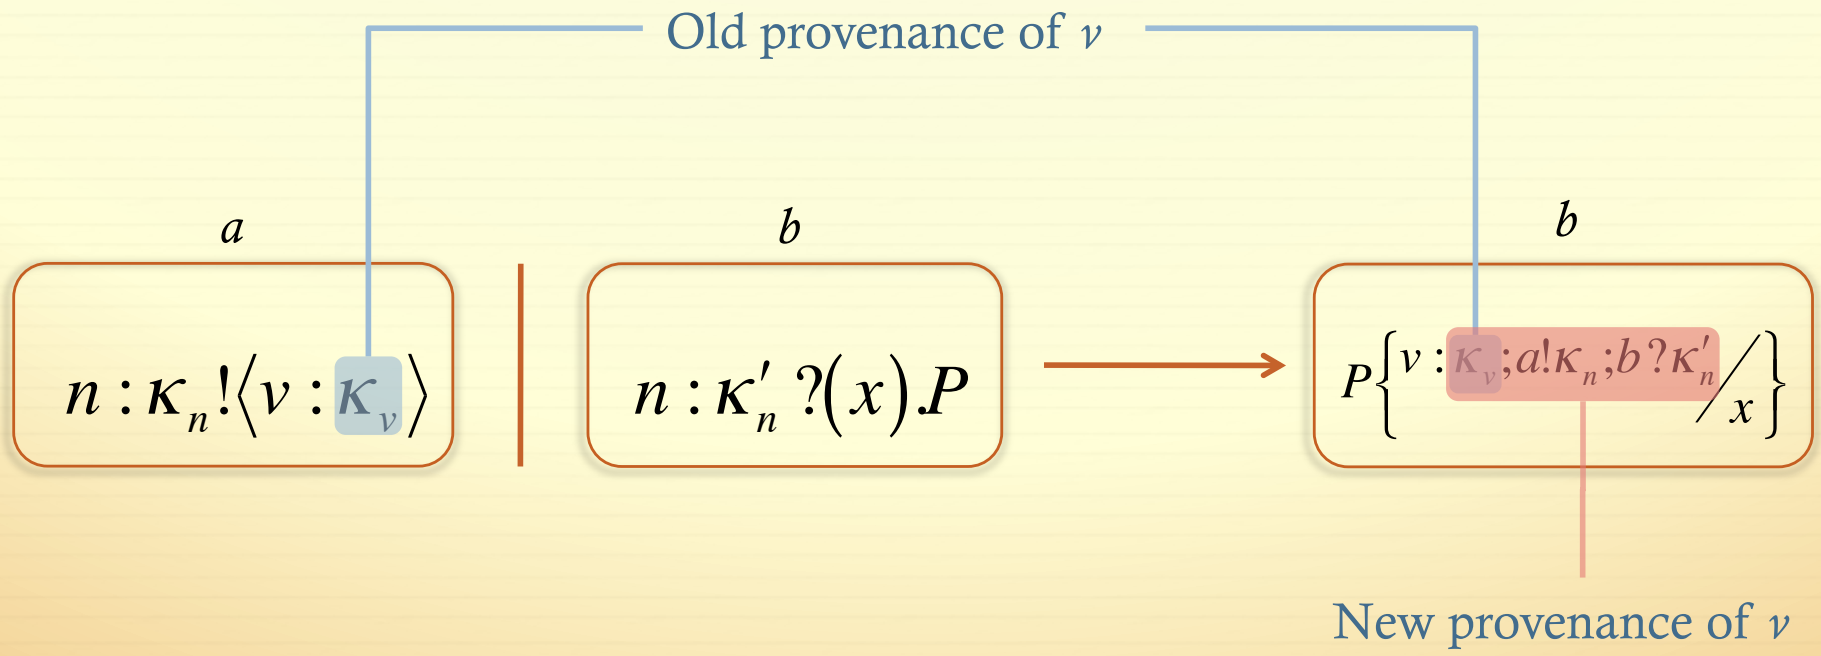

Provenance model

Provenance tracking

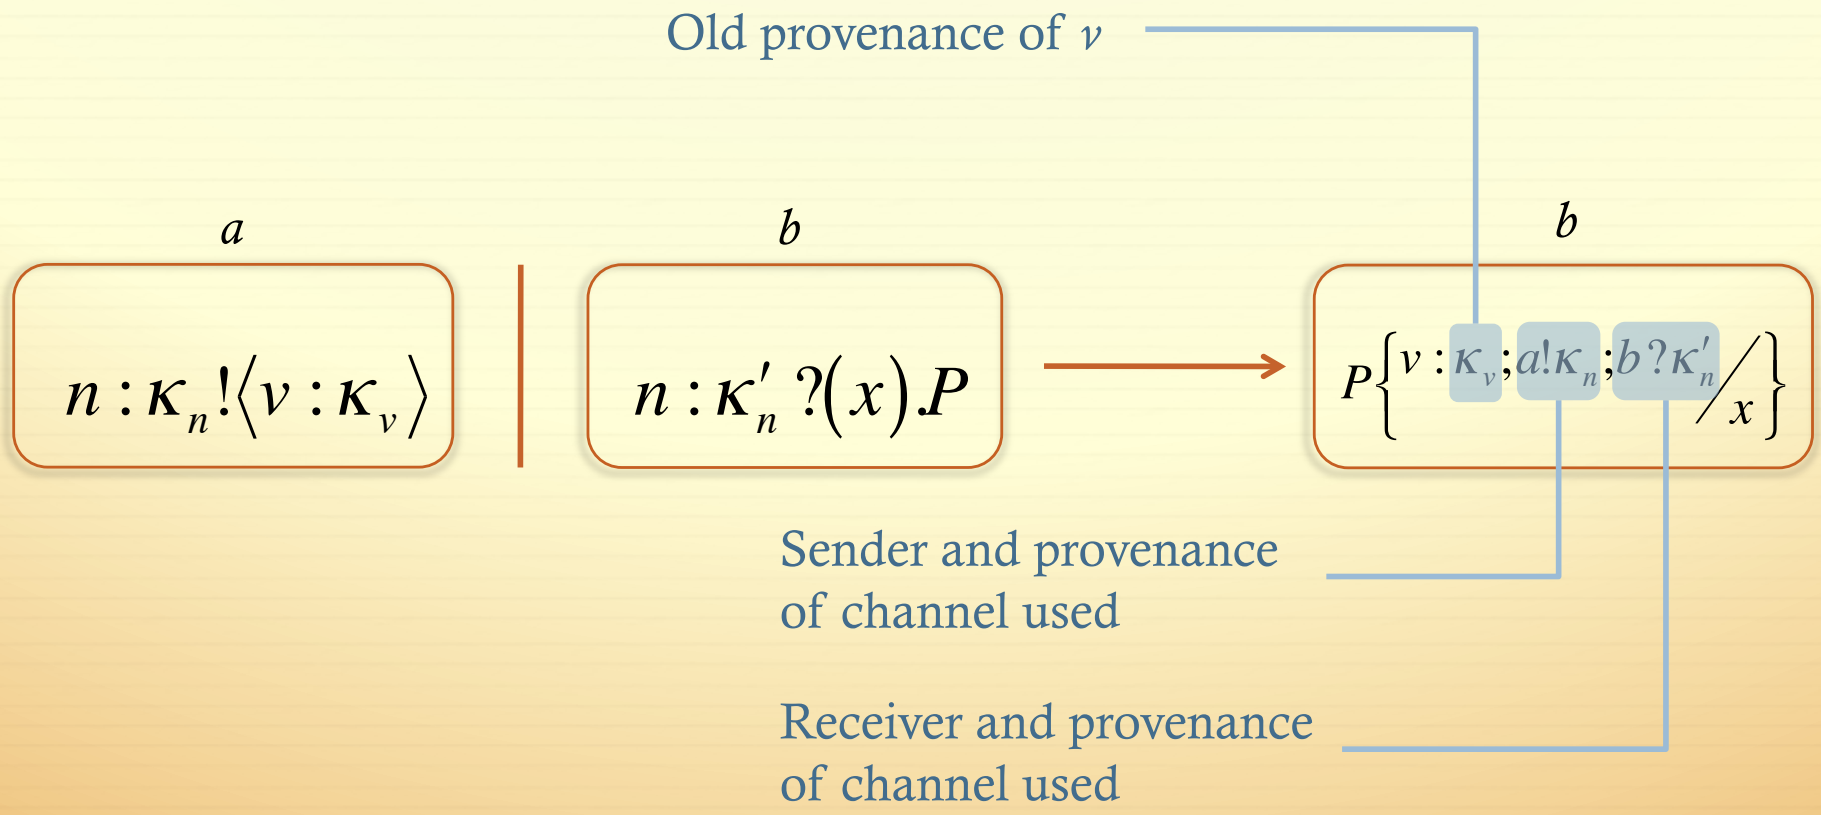

Confidentiality in provenance systems

- Data may be public, yet its provenance confidential, or vice versa
- Principals who may access data are not necessarily the same as those who may access its provenance
- In general, fine grained access control over provenance "histories" is needed as different parts of it have different sensitivity

Security requirements of data ≠ Security requirements of its provenance

Hiding provenance trees

Example: photography competition

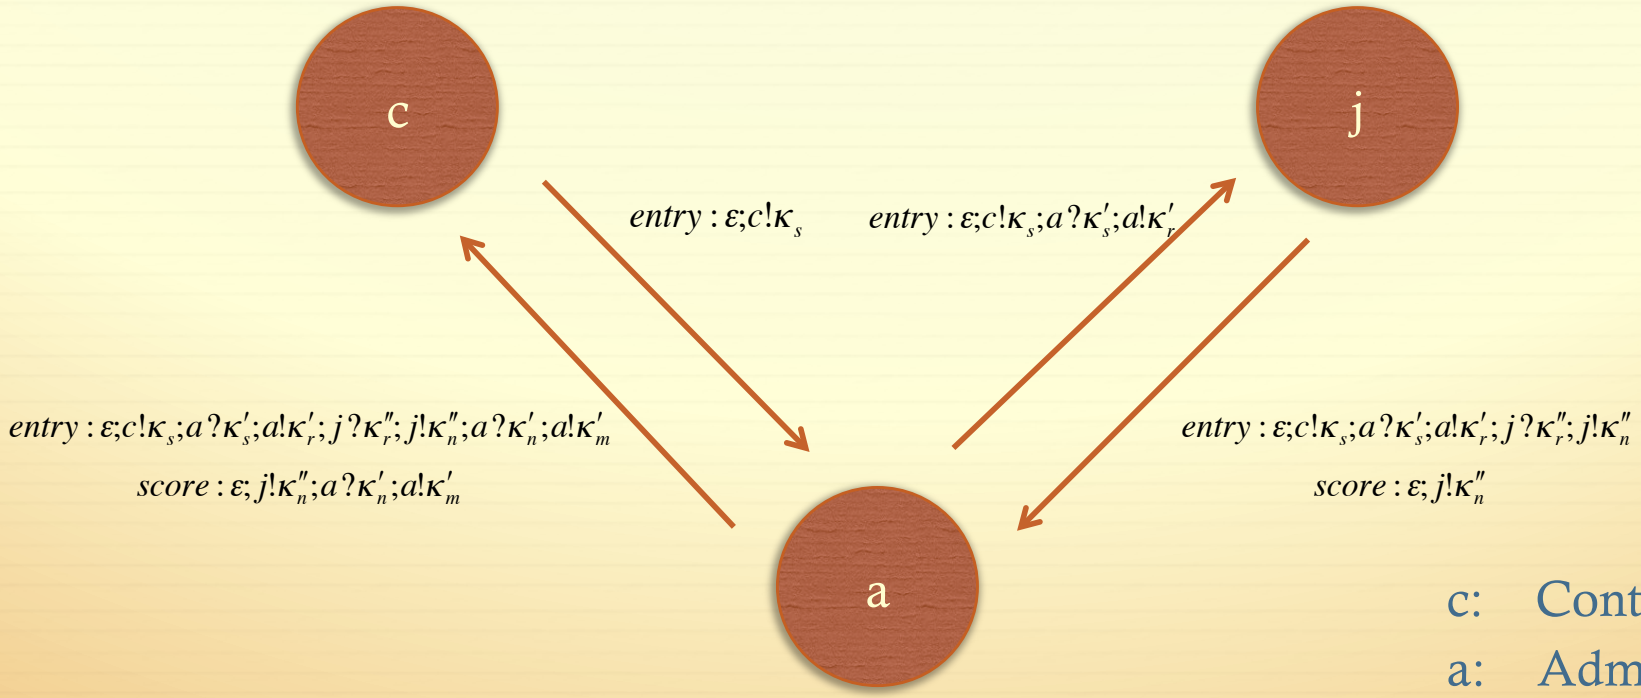

c: Contestant

a: Administrator

j: Judge

Hiding provenance trees

Example: photography competition

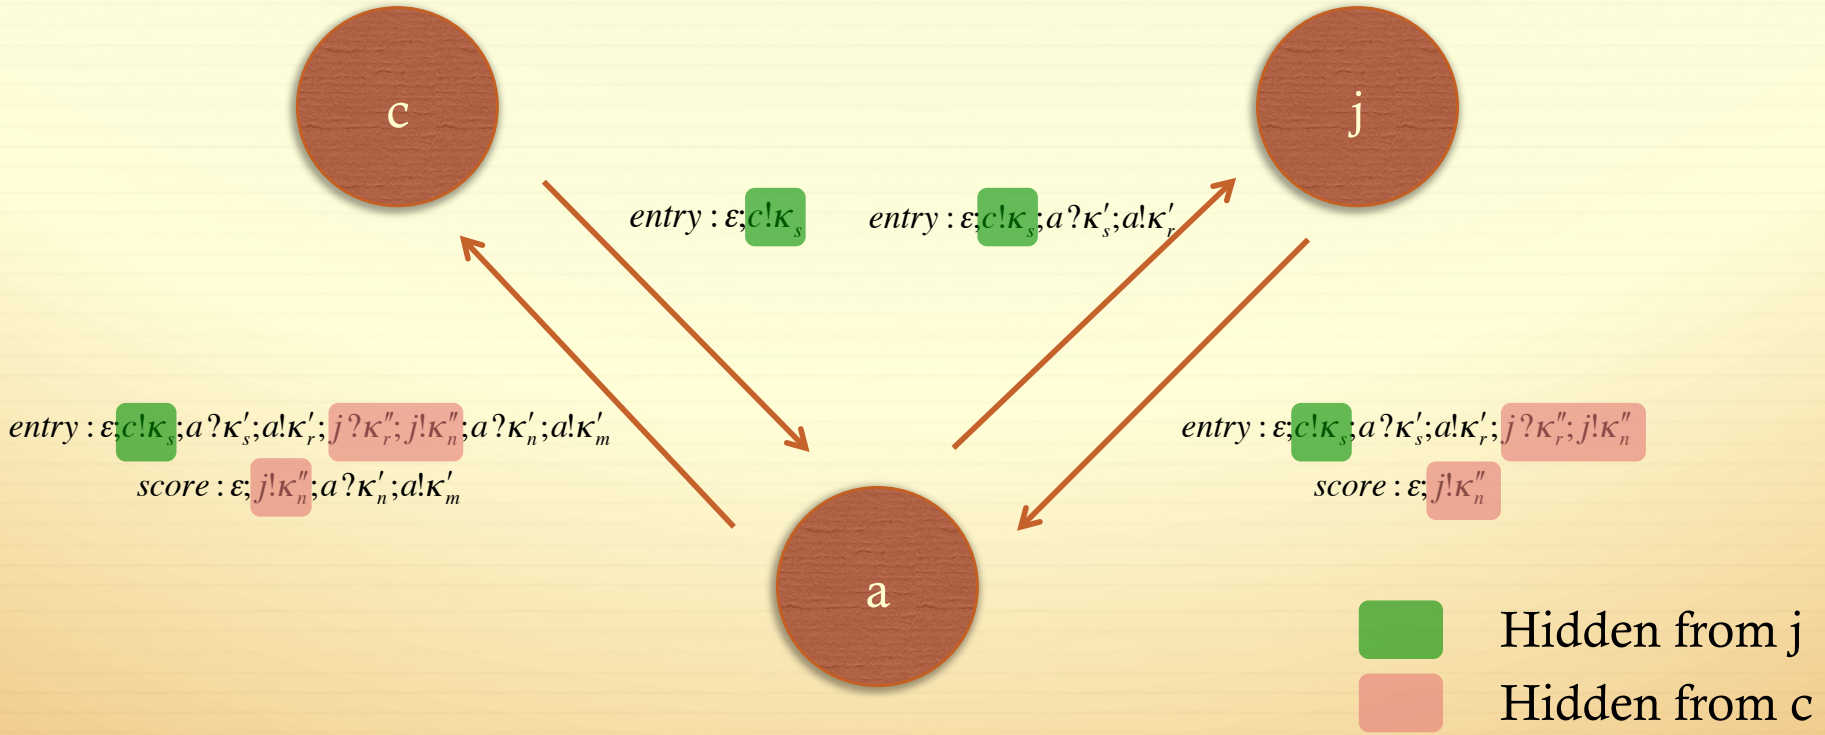

Confidentiality in provenance systems

a promising approach

- One value, multiple **views**
  - Different principals have different views of the same provenance list based on their privileges

$$entry : \varepsilon; c!\kappa_s; a?\kappa'_s; a!\kappa'_r; j?\kappa''_r; j!\kappa''_n; a?\kappa'_n; a!\kappa'_m$$

Confidentiality in provenance systems

a promising approach

- One value, multiple **views**
  - Different principals have different views of the same provenance list based on their privileges

$$entry : \varepsilon; c!\kappa_s; a?\kappa'_s; a!\kappa'_r; j?\kappa''_r; j!\kappa''_n; a?\kappa'_n; a!\kappa'_m$$

Confidentiality in provenance systems

a promising approach

- One value, multiple **views**
  - Different principals have different views of the same provenance list based on their privileges

$$entry : \varepsilon; c!\kappa_s; a?\kappa'_s; a!\kappa'_r; j?\kappa''_r; j!\kappa''_n; a?\kappa'_n; a!\kappa'_m$$

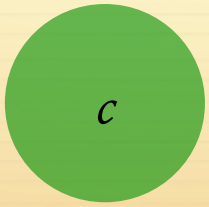

Confidentiality in provenance systems

a promising approach

- One value, multiple **views**
  - Different principals have different views of the same provenance list based on their privileges

$$entry : \varepsilon; c!\kappa_s; a?\kappa'_s; a!\kappa'_r; j?\kappa''_r; j!\kappa''_n; a?\kappa'_n; a!\kappa'_m$$

Confidentiality in provenance systems

a promising approach

- One value, multiple **views**
  - Different principals have different views of the same provenance list based on their privileges

$$entry : \varepsilon; c!\kappa_s; a?\kappa'_s; a!\kappa'_r; j?\kappa''_r; j!\kappa''_n; a?\kappa'_n; a!\kappa'_m$$

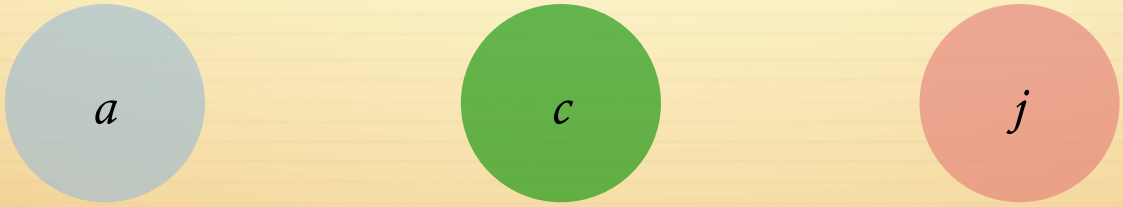

Confidentiality in provenance systems

a promising approach

- To achieve this, we use ***groups***
- Different principals belong to different groups
- Group membership determines what parts of a provenance list a principal has access to
- Principals
  - Can create new groups: $\text{new } G$
  - Can add other principals to their groups: $\text{add}(a,G)$
  - Can **restrict** access to particular parts of a provenance tree to a particular group: $\text{hide}(v : \kappa, G)$

Hiding provenance trees

photography competition

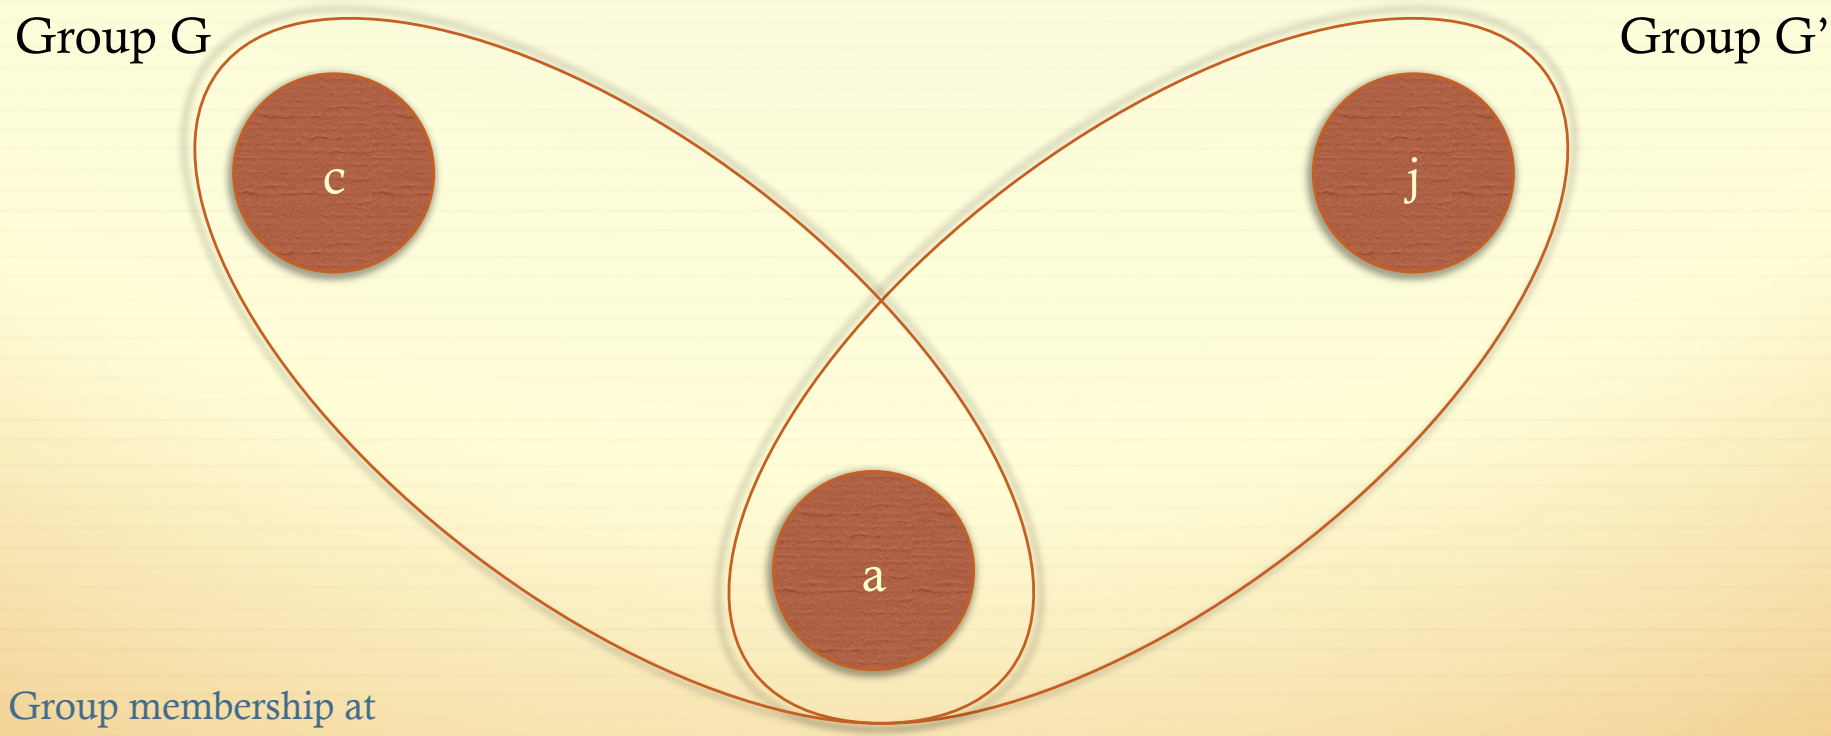

Group membership at start of competition

Hiding provenance trees

photography competition

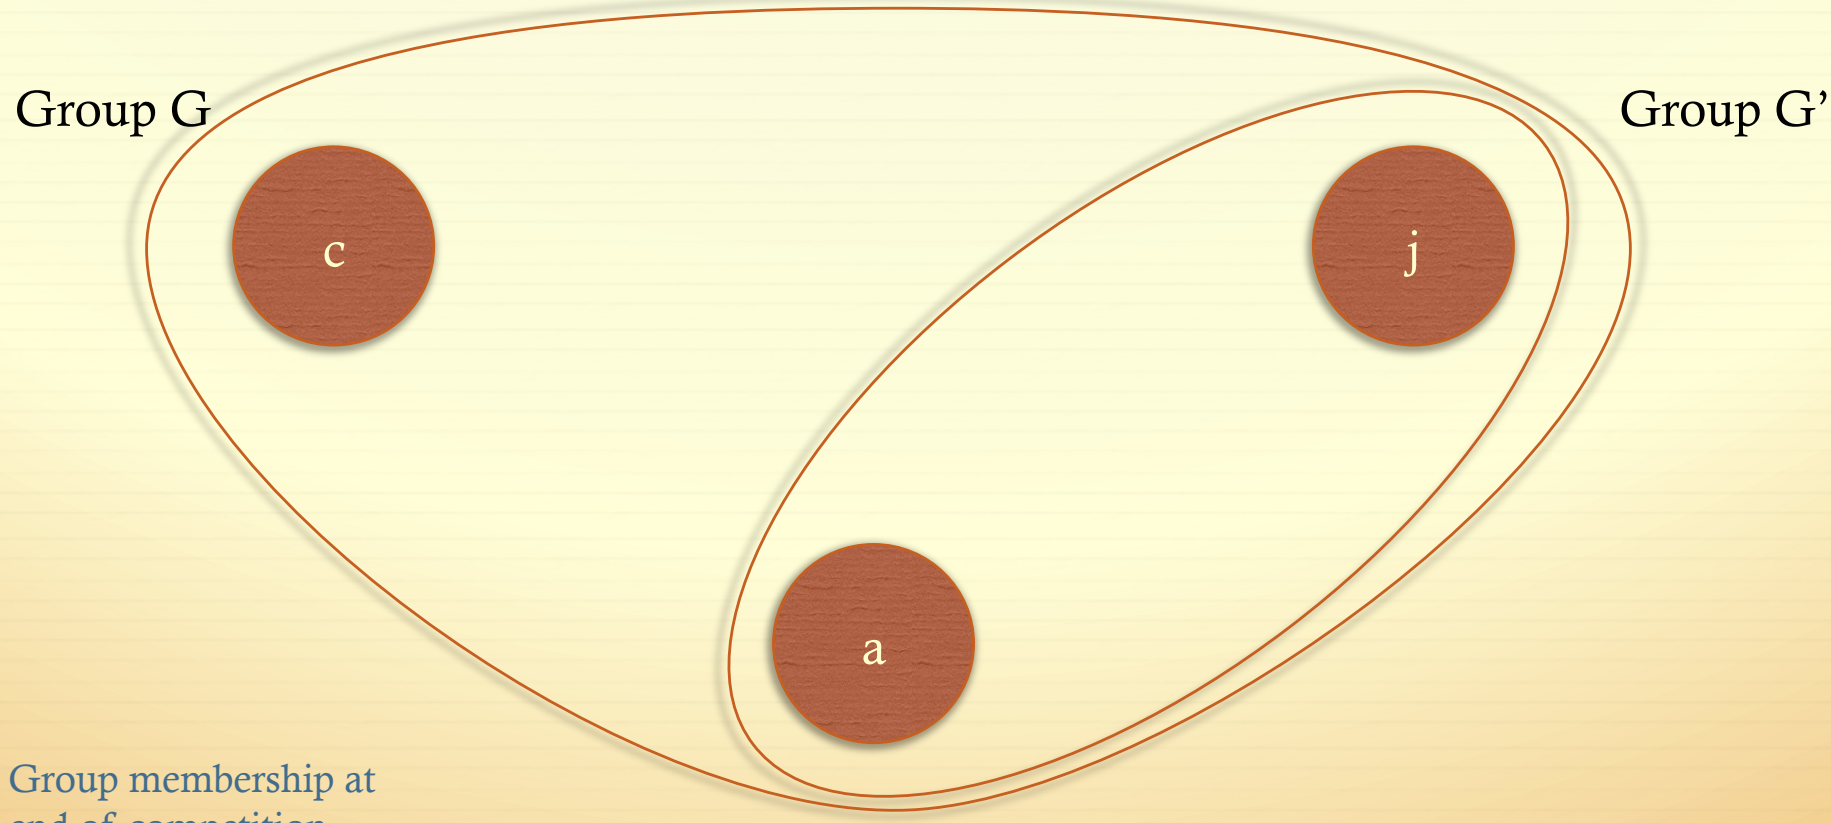

Group membership at
end of competition

Current work

- Correctness of provenance tracking: the provenance information determines the history of each piece of data accurately "enough"
  - Express this as a form of testing (on traces):

$$\forall S(\forall t \in [\![S]\!](t \rightarrow^* v : \kappa \implies \forall s \in [\![v : \kappa]\!](s|t \rightarrow^* \checkmark)))$$

- Using provenance:
  - Provenance queries vs pattern restricted input
  - Trust in quality of data based on trust in principals and provenance of data
- Policies and types