Trust and Reputation Systems Part 1

Audun Josang
QUT, Australia
a.josang@qut.edu.au
http://www.fit.qut.edu.au/~josang/

Tutorial overview

- Wednesday 13 September
  - Basic trust concepts
  - Trust classes and trust semantics
  - Principles for building trust and reputation systems
    - Network architectures
    - Reputation engines (binomial)
- Thursday 14 September
    - Reputation engines (multinomial)
    - Trust computation engines
  - Commercial and online systems
  - Problems and proposed solutions
  - Concluding remarks

Basic trust concepts

Complexity of trust

Trust is a complex concept with multiple meanings

Concept:	Counts:	Webster's (1981)	Oxford's (1989)
Cooperation	# defs. # lines	3 14	2 75
Trust	**# defs. # lines**	**9 112**	**18 633**
Love	# defs. # lines	17 82	28 1670

Source: McKnight & Chervany 1996

Manifestations of trust

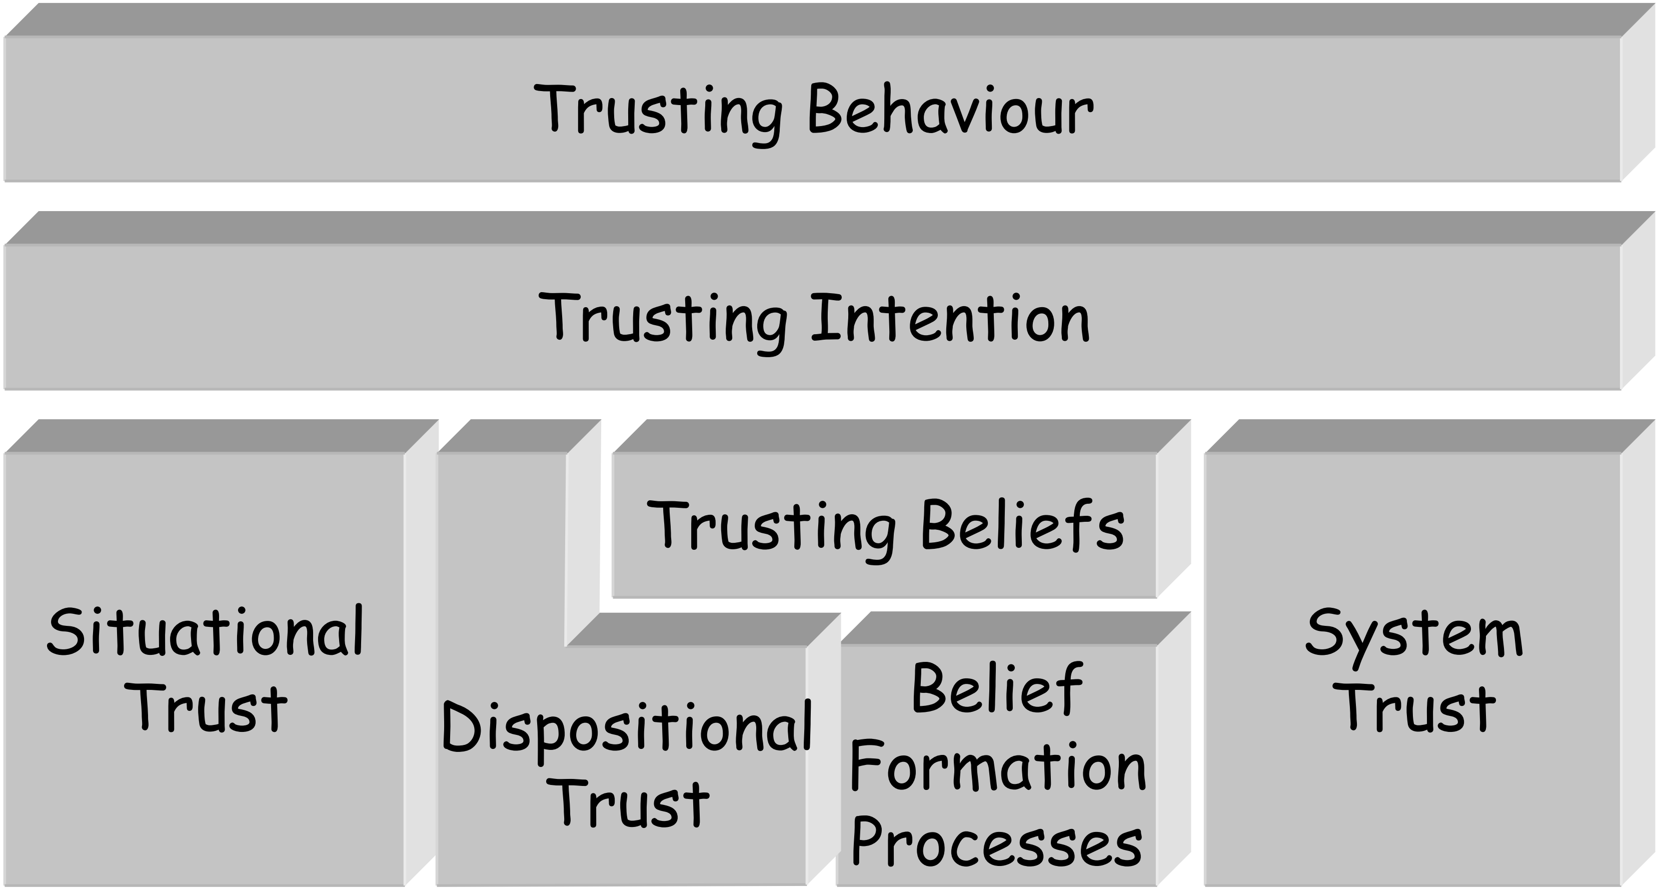

Source: McKnight & Chervany 1996

Two definitions of trust

- Reliability trust
  - The **subjective probability** by which an individual, *A*, expects that another individual, *B*, performs a given action on which its welfare depends. (Gambetta 1988)

- Decision trust
  - The **willingness to depend** on something or somebody in a given situation with a feeling of relative security, even though negative consequences are possible. (McKnight & Chervany 1996)

Would you trust this rope?

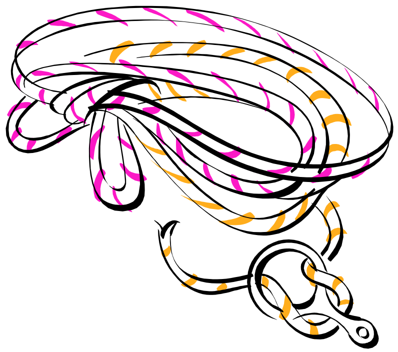

For what?

To climb down from the 3rd floor window of a house

The rope looks a bit old

Fire drill:	**No!**

Real fire:	**Yes!**

Computational trust

- Most computational models assume reliability trust.
- Decision trust not often modelled
- Decision trust can be complex, and needs to take many additional factors explicitly into account, e.g. utility, risk, risk attitude, reliability.
- Examples of decision trust models:
  - Manchala (1998)
  - Josang & Lo Presti (2004)

Manchala's Risk-Trust Matrix

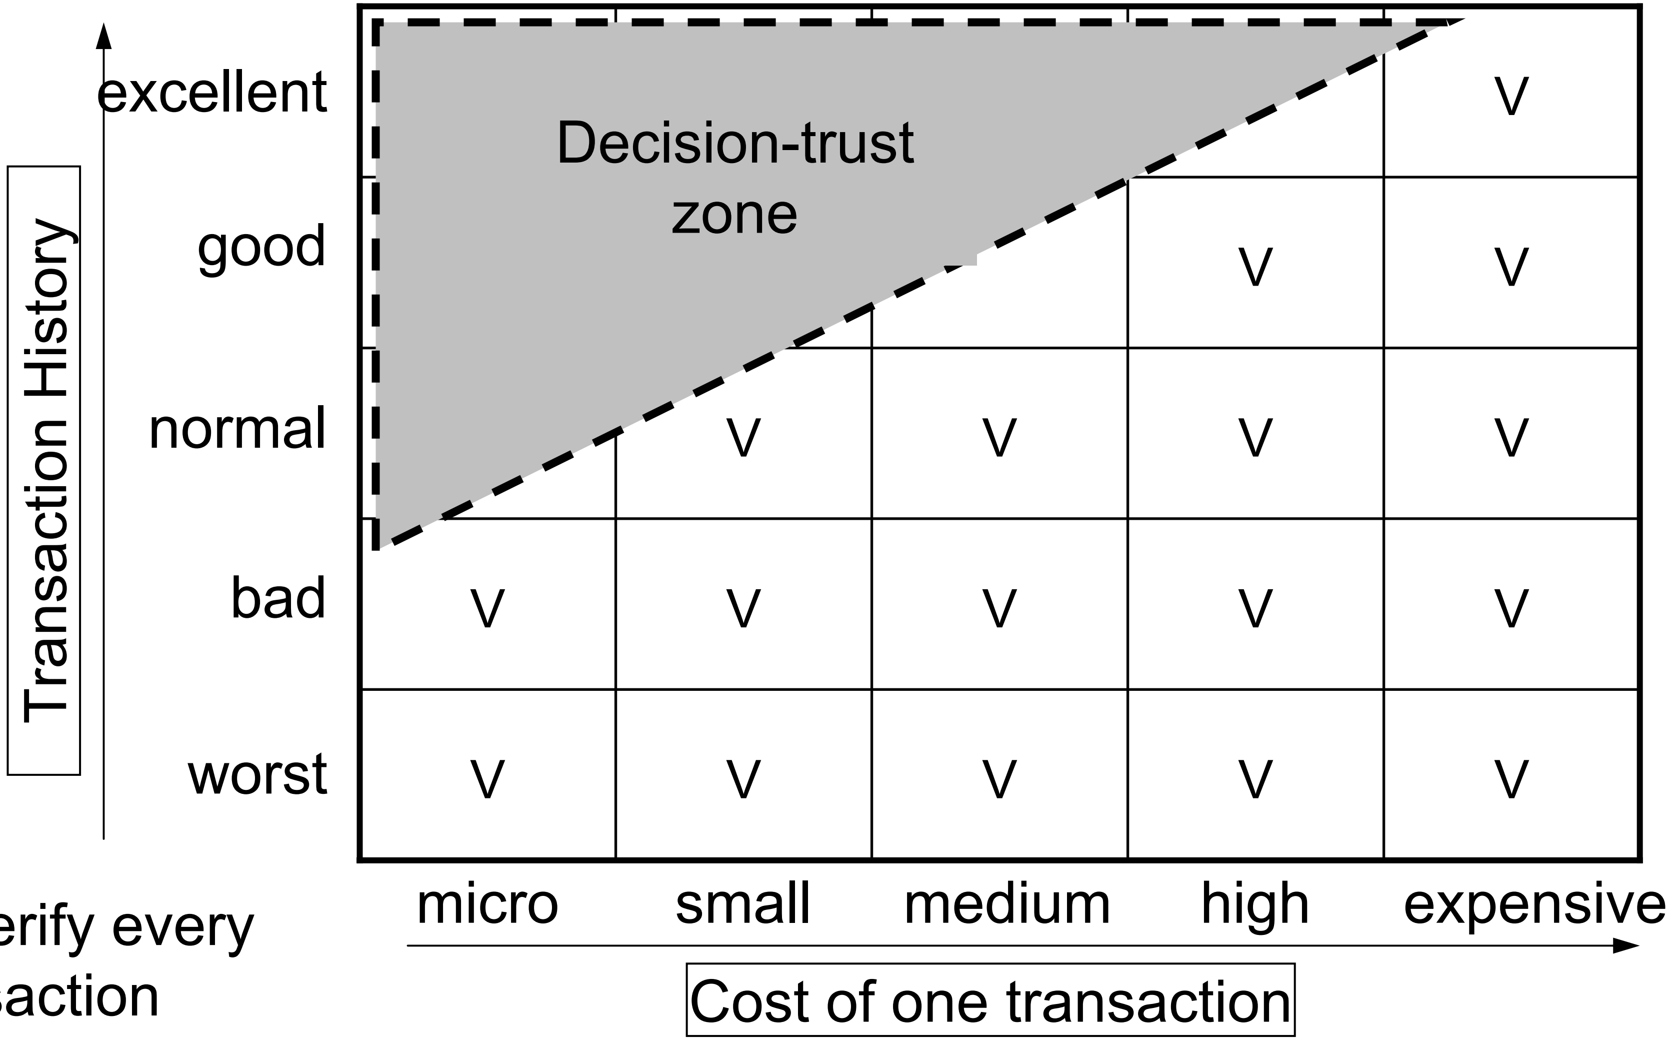

V: Verify every transaction

Trust and economic modelling

- Trust adds nothing new
  - (Williamson: Calculativeness, Trust and Economic Organisation, 1993)
- Many advanced economic models for decision making, based on
  - Reliability
  - Utility (subjective and objective)
  - Risk and risk attitude
  - etc.
- The original elements of computational trust modelling comes from the **architectures** for communicating and processing information relating to trust and decision making

When is trust relevant?

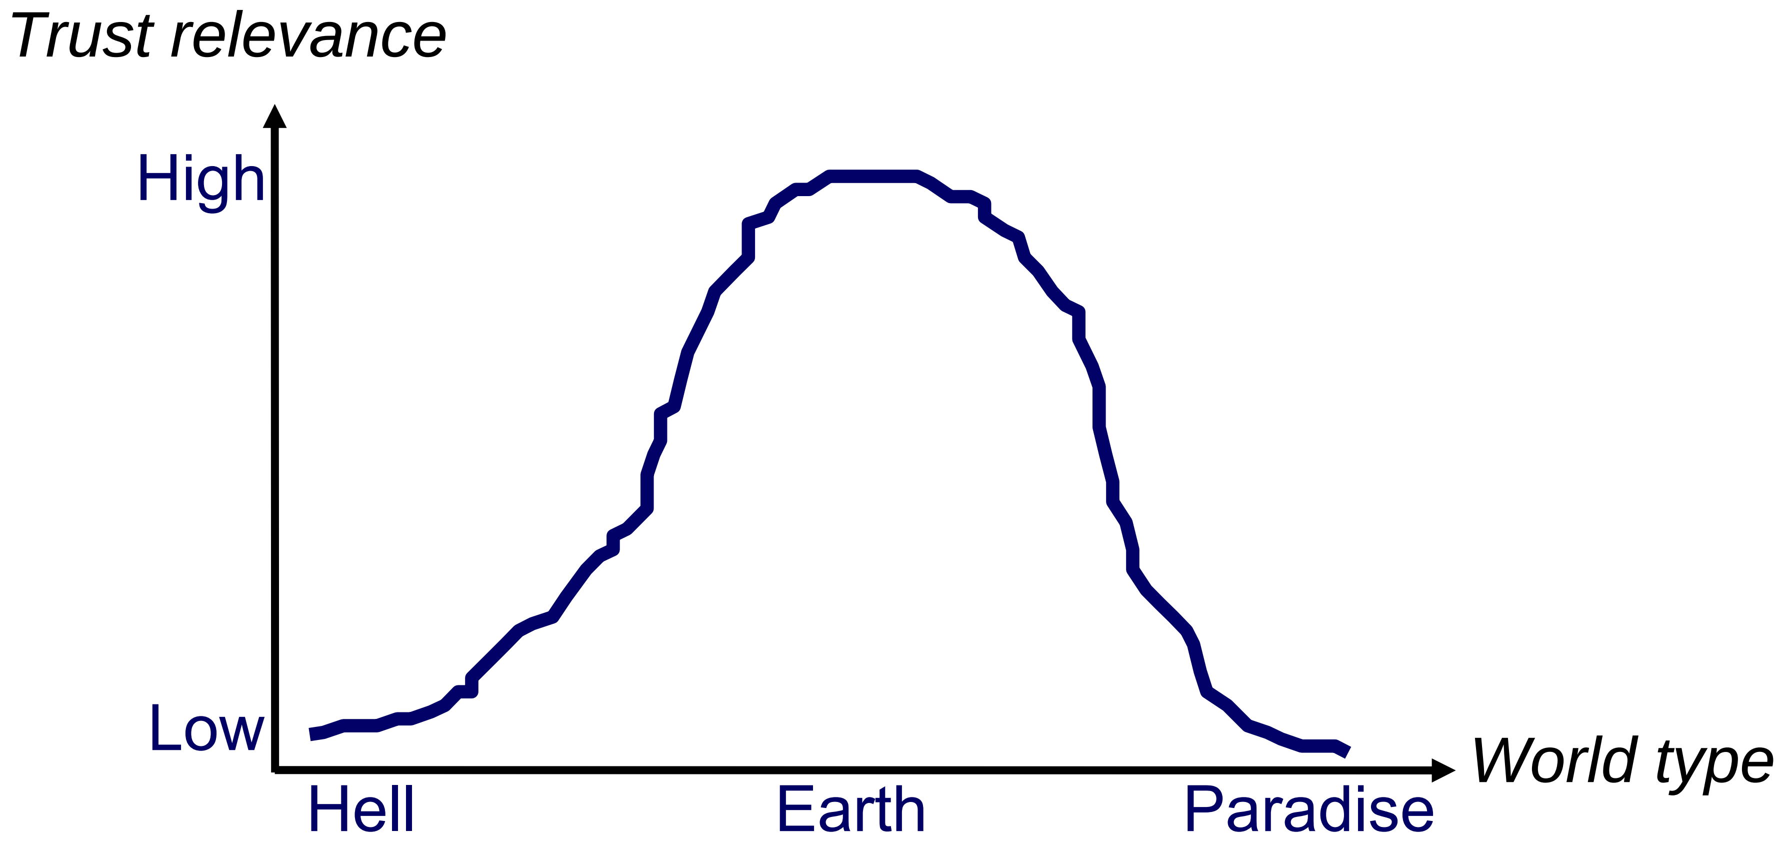

Trust related phenomena

- Dependence
- Belief
- Uncertainty
- Risk
- Risk attitude
- Decision
- Dynamics
- Subjectivity

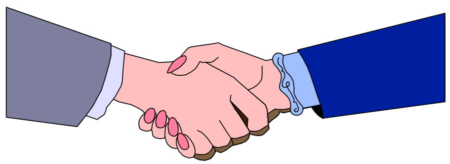

Hormones and trust

The hormone *oxytocin*

- is released after trusting behaviour, and
- stimulates trusting behaviour

Figure 1: OT and total amount sent by DM2

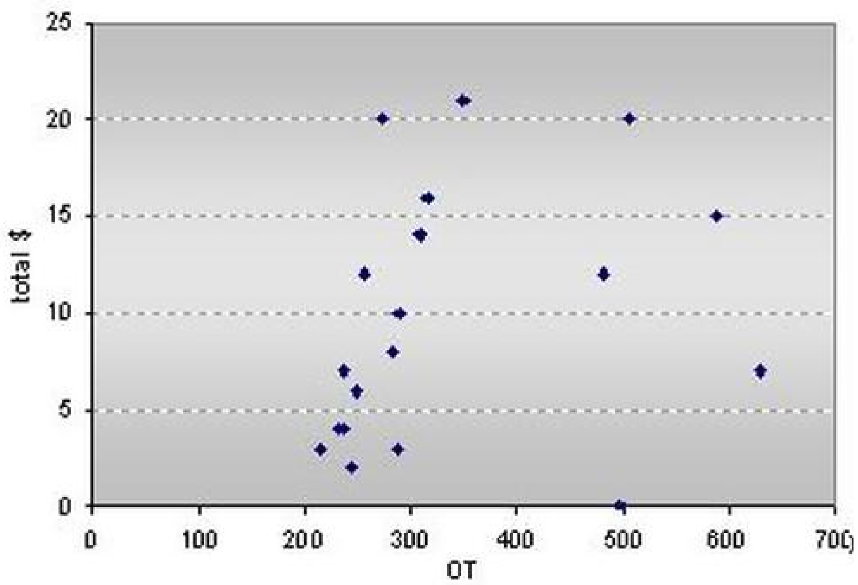

from Zak *et al.*, 2003

Importance of trust

- Progress requires collaboration
- Potential collaboration partners must make decisions involving risk and uncertainty
- Fear of negative consequences is an obstacle for collaboration
- Trust
  - is the perception that the risk is acceptable
  - Is a catalyst for human cooperation
  - influences type and size of organizations
  - represents social capital in a community

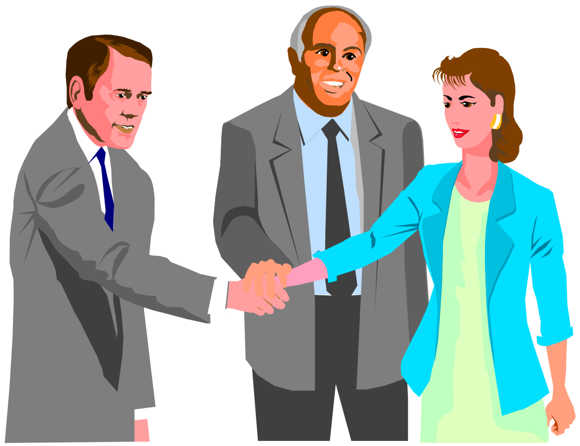

Trust is a relationship

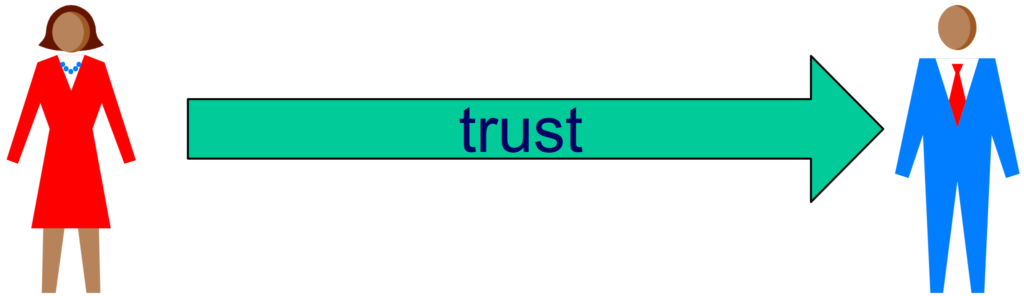

- Trusting party
  - Also called
    - “relying party”
    - “trustor”
  - Is in a situation of
    - Dependence
- Trusted party
  - Also called
    - “trustee”
  - Is in a situation of
    - Power
    - Expectation to deliver

Two sides of trust management

Trusting party

Wants to **assess** and make **decisions** w.r.t. the dependability of the trusted party for a given transaction and context

Trusted party

Wants to **represent** and put in a **positive light** own competence, honesty, reliability and quality of service.

A definition of reputation

- Reputation is what is generally said or believed about a person's or thing's character or standing. (Concise Oxford Dictionary)
  - (Reputation of *B*)= Average[Reliability Trust in *B*]

Reputation and trust

<u>REPUTATION</u>

- Public info
- Common opinion
- Not necessarily objective

<u>TRUST</u>

- Both private and public info
- Private info carries more weight
- Subjective

- *"I trust you because of your good reputation"*
- *"I trust you despite your bad reputation"*

Reputation aspects

- Default / base rate reputation of a group’s member
  = the group’s reputation
- A group’s reputation
  = average reputation of its members
  (not always true, one bad example can destroy …)
- Reputation of well-known companies transfer from the real to the online world.
- Reputation of lesser known companies is built on what others say about them online
- Reputations are not a function of right or wrong, but of perception, whether correct or not.

We trust what we depend on

Trust in people & organisations

Trust in legal, social and market institutions

Trust in ICT

Why is trust so popular?

- Metaphorical trust expressions
  - IT security people like metaphors:
    - E.g. firewall, honeypot, virus, Trojan horse, digital signature
  - Trust expressions serve as simple metaphors for complex security concepts, e,g. , …, ***trusted code, circle of trust***, …
- Trust has very positive connotations
  - Trust expressions are ideal as marketing slogans

Trust expressions can be difficult to intuitively understand

Trust as an abstract security layer

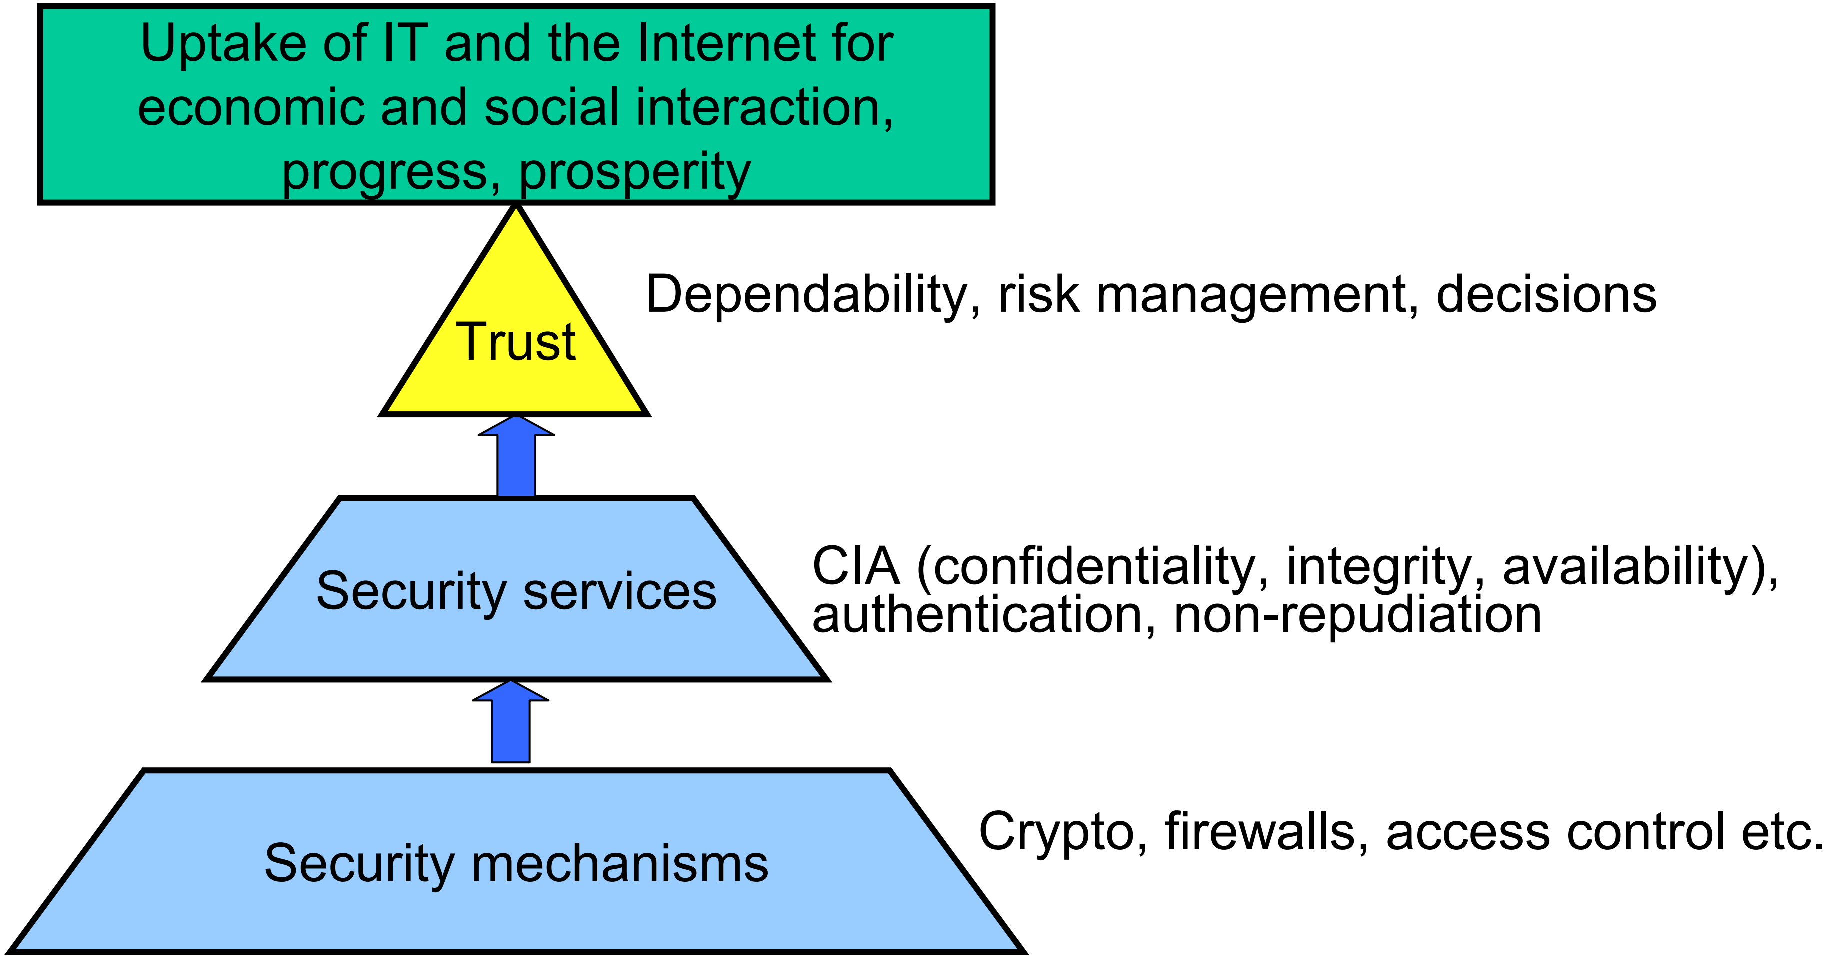

Trust as assumptions and primitives

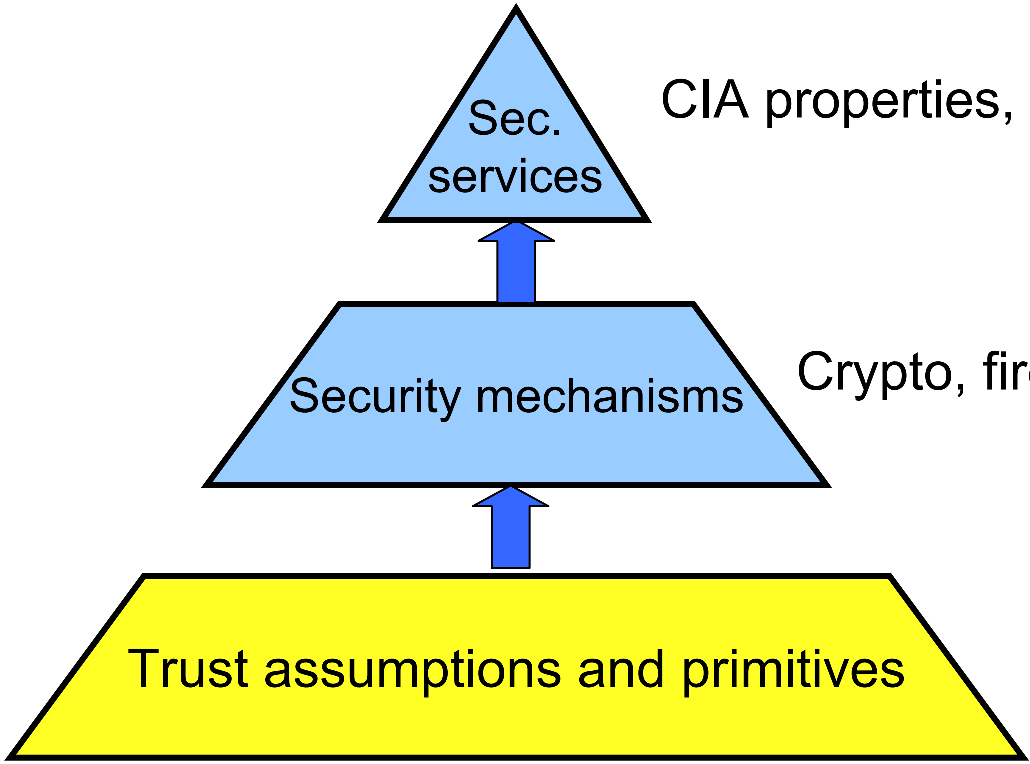

CIA properties, authentication, non-repudiation

Crypto, firewalls, access control etc.

- Algorithm strength
- Password confidentiality
- Platform security
- Certificates and credentials
- Tokens and attributes

Trust and IT security

Trusted code

Trustworthy computing

Trust management

WS Trust

Trusted Computing Base

Trusted system

Trusted computing

Trust bar

Trust eco-system

Trust negotiation

Trust provider

Circle of trust

Trusted Third Party

Hard v. soft security

- Security is the protection from harm
- Traditional **information security:**
  - Confidentiality, integrity & availability of info assets
  - Hard security
- What about deceit and poor quality services?
  - Problem is inversed, information assets can harm
  - Traditional security provides no protection
  - Trust and reputation systems provides protection
- **Trust and reputation systems:**
  - Soft security

Hard v. soft security

Hard Security

- Focuses on the assets and the methods to protect those assets from attackers
- Goal: to preserve the CIA properties of assets.
- Attacker agnostic

Soft Security

- Focuses on the attackers, and collaborative methods to identify and sanction them
- Goal: To stimulate quality assets and service providers
- User agnostic

Trust and access control

- Access control paradigm:
  - The resource owner grants access authorisation
  - The system verifies authorisation before access
- **Trusted user** = authorised user
- **Trusted code** = code running as system
- **Untrusted code** = code running in a sandbox
- **Semitrusted code** = some more access rights
- Access credentials can be exchanged and evaluated mechanically $\Rightarrow$ **trust negotiation**
- Access authorisation can be delegated in a transitive fashion $\Rightarrow$ **transitive trust**

AC Conceptual diagram

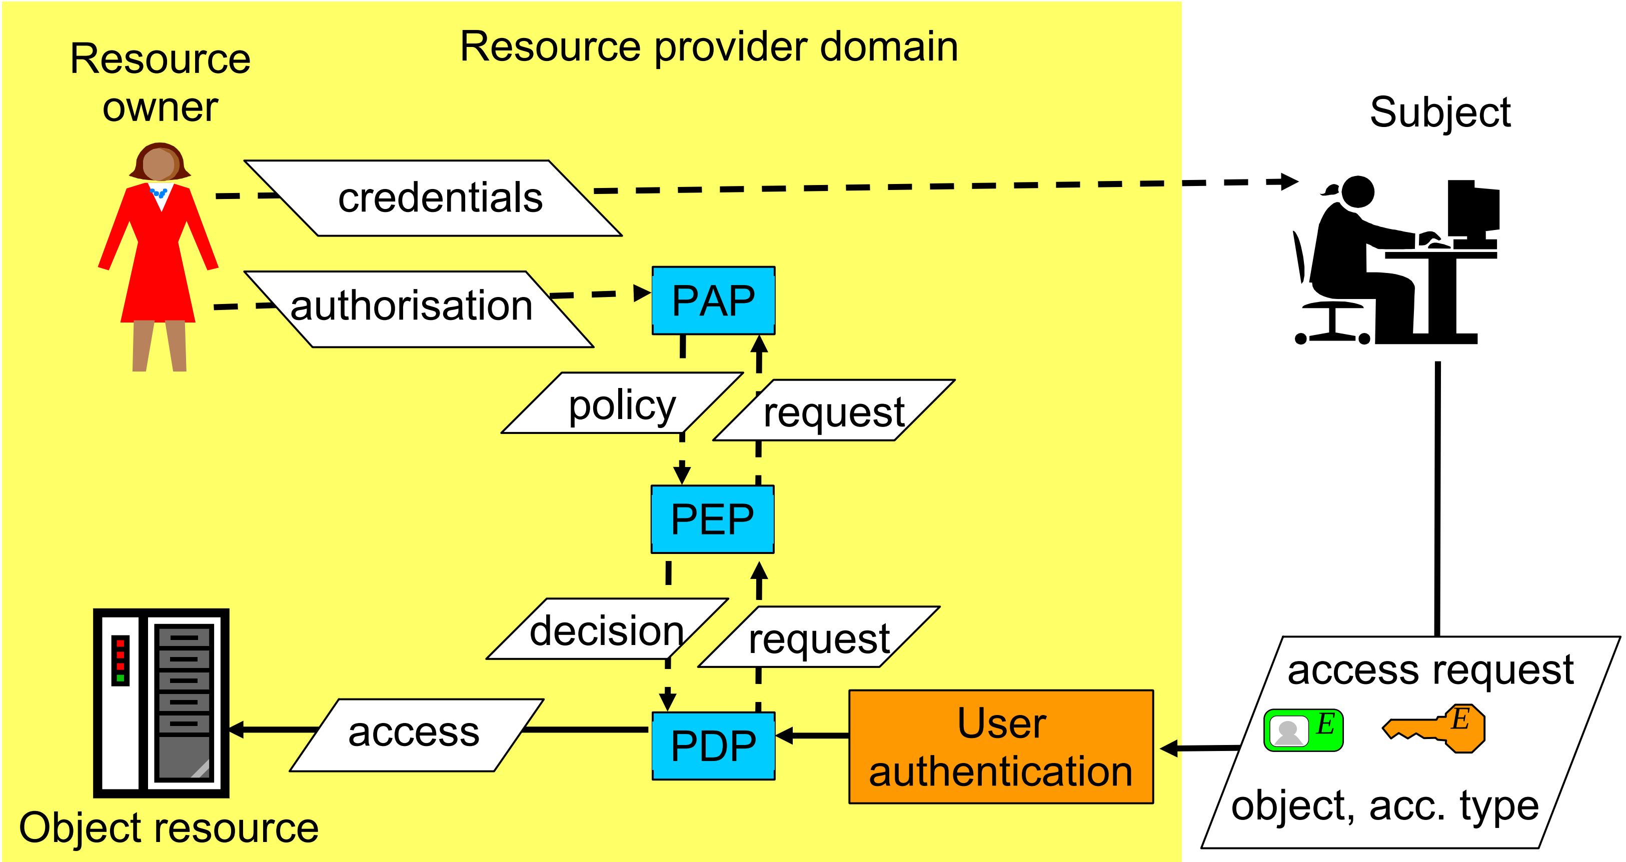

Legend PAP: Policy Administration Point (WS-Security terminology and architecture)

PEP: Policy Enforcement Point http://www.oasis-open.org/specs/index.php

PDP: Policy Decision Point

Distributed access control

originally called trust management (1996)

- Idea: “Who can I trust to access my resources?”
- Access authorisation can be delegated in a distributive fashion

> Trust management is supposed to be an incredibly vague and provocative term invented by Matt Blaze. I don’t know whether he intended it that way, but it comes natural to him

Joan Feigenbaum, AT&T Labs

Trust expressions

- **Trusted computing** = Computing platform with additional security hardware
- **Trustworthy computing** = Microsoft marketing slogan
- **Trust eco-system** = Microsoft marketing slogan
- **WS Trust** = WS Security standard specifying how to generate security tokens
- **Trust Bar** = Mozilla browser toolbar
- **Circle of trust** = Liberty Alliance term for group of organisations that enter into identity federation agreement
- **Trust provider** = Certificate Authority
- **Trusted Third Party** = Entity assumed to keep secrets

Perception and reality; The subjective perspective

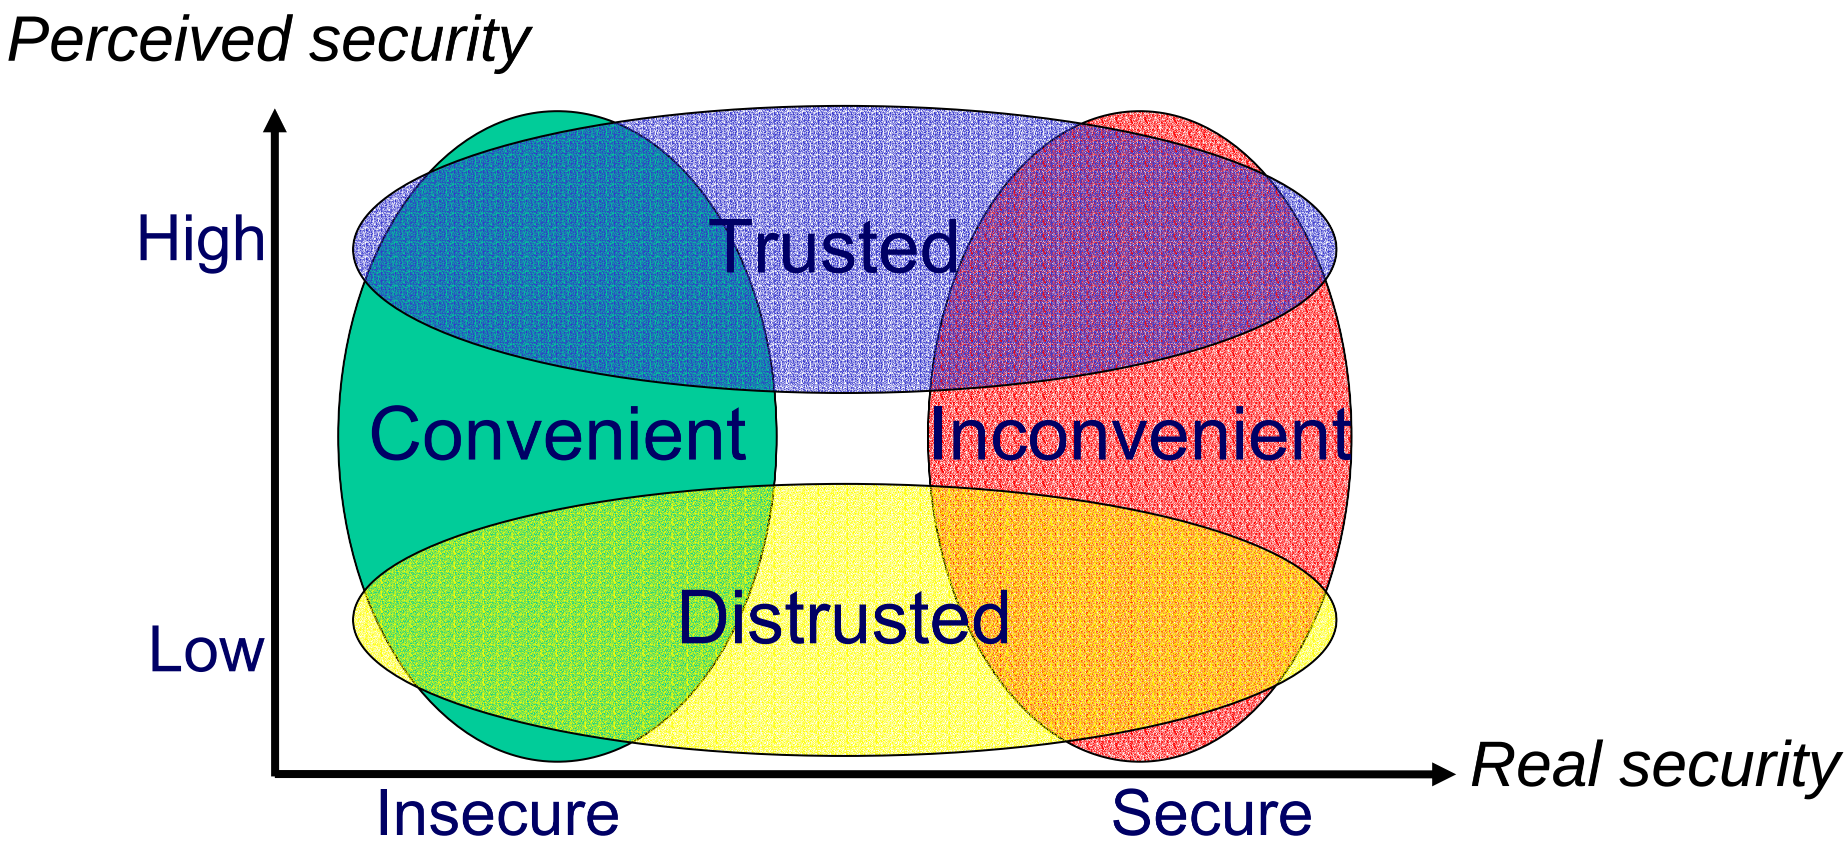

Real security is bad for e-business

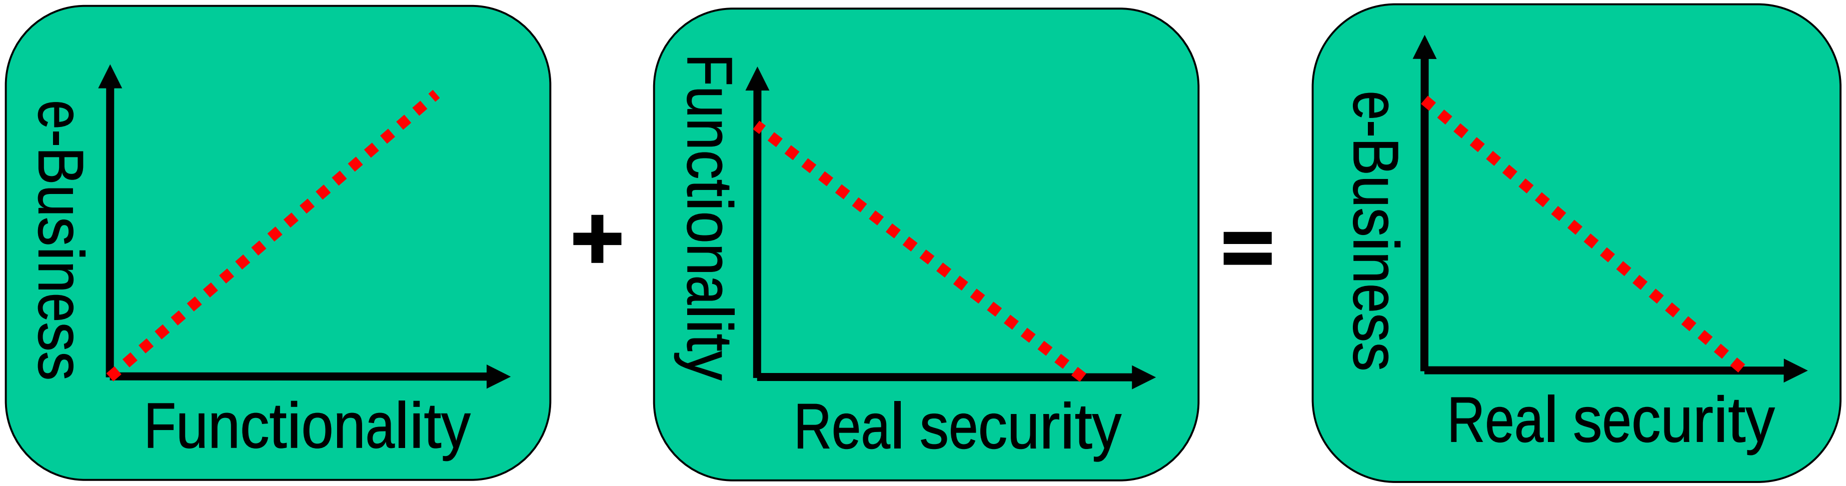

- e-business revolution not possible with real security
- Thank God the Internet isn't secure

Perceived security is good for e-business

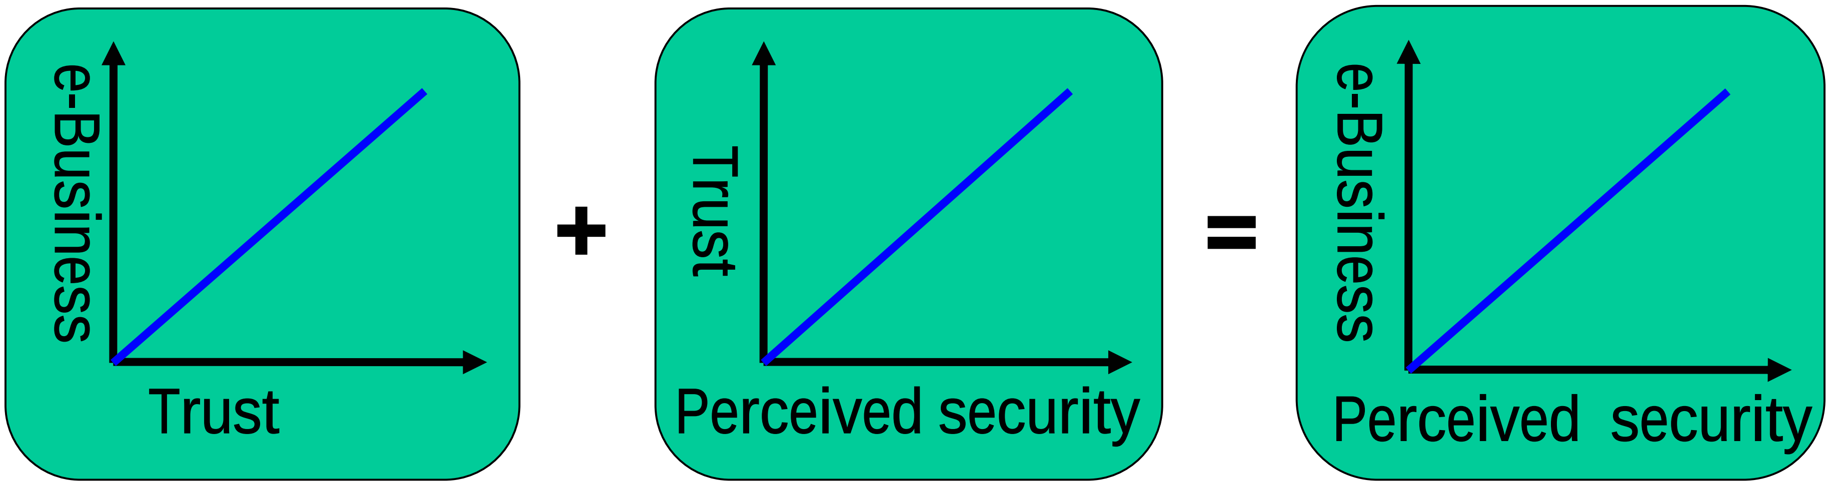

- e-business growth needs perceived security

The security dilemma

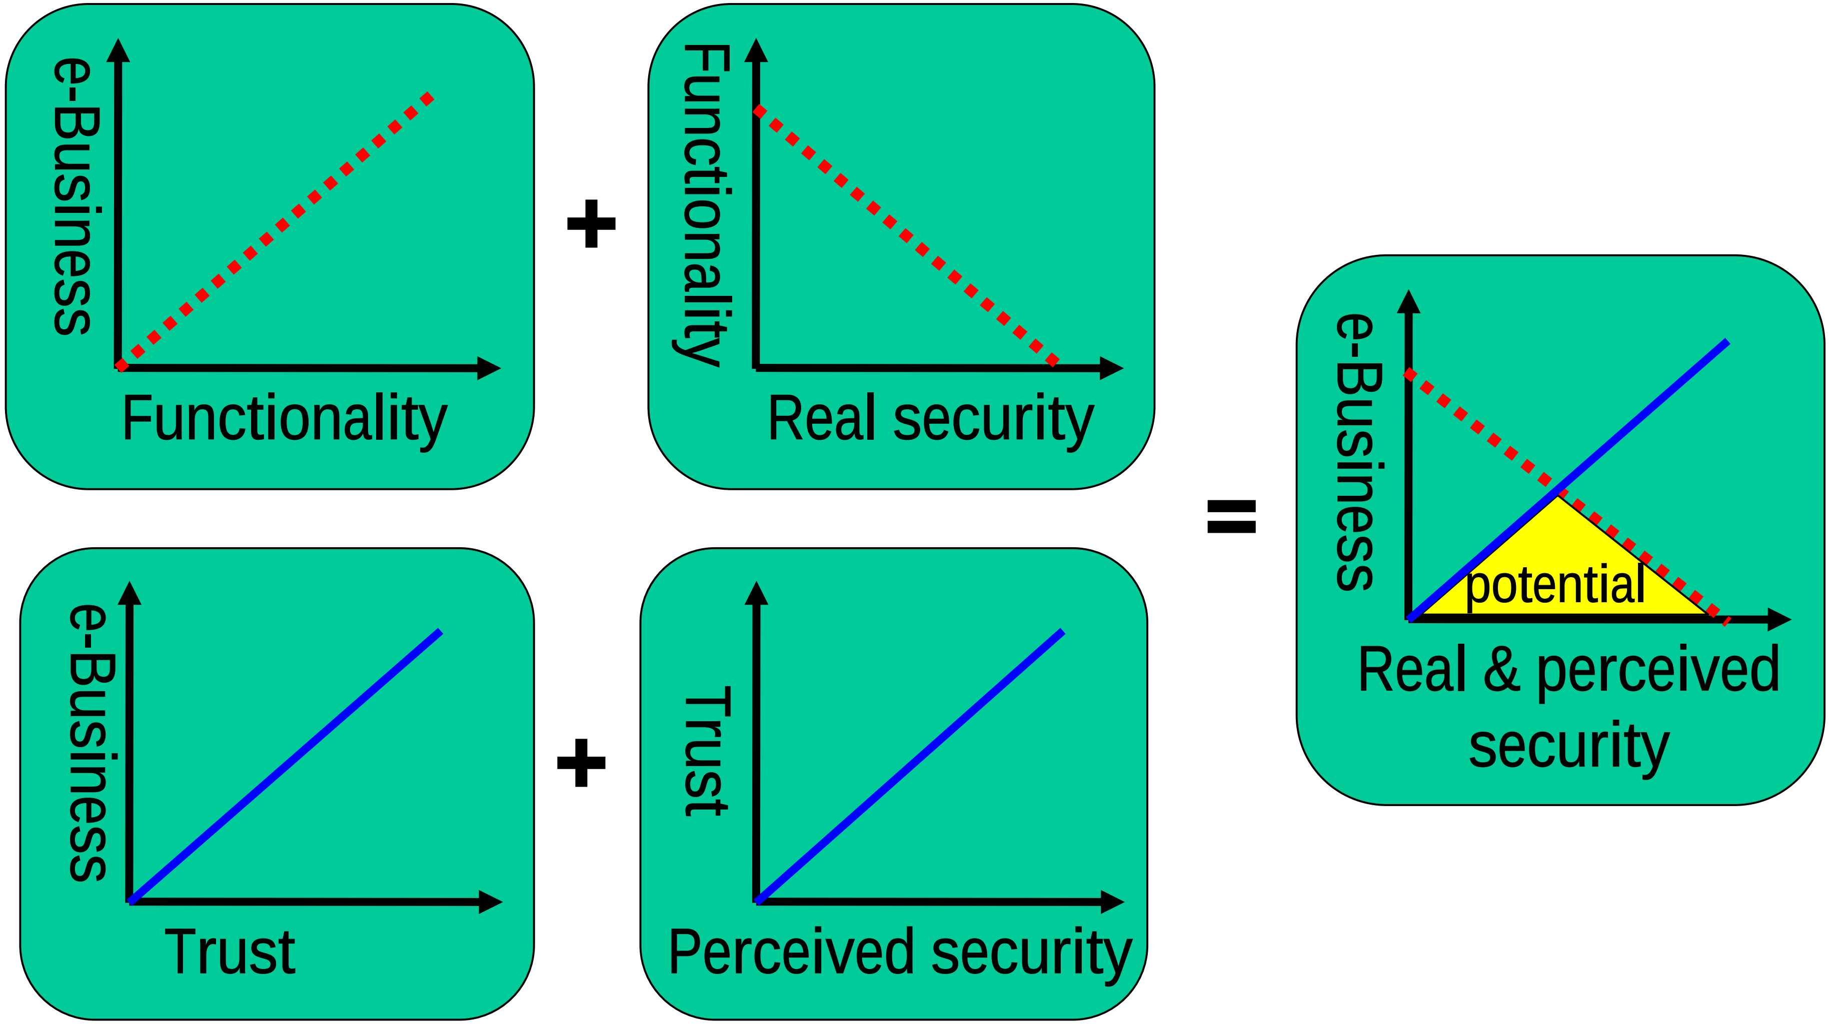

Jøsang's law of security and e-business

The potential of e-business is bounded by:

- The lack of functionality caused by real security
- The lack of trust caused by perceived insecurity

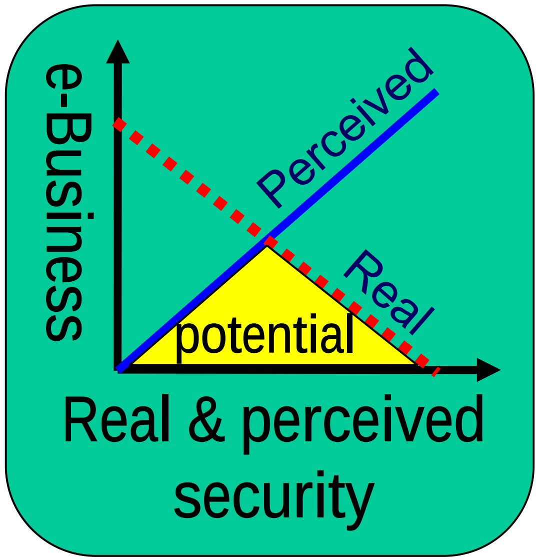

Trust classes and semantics

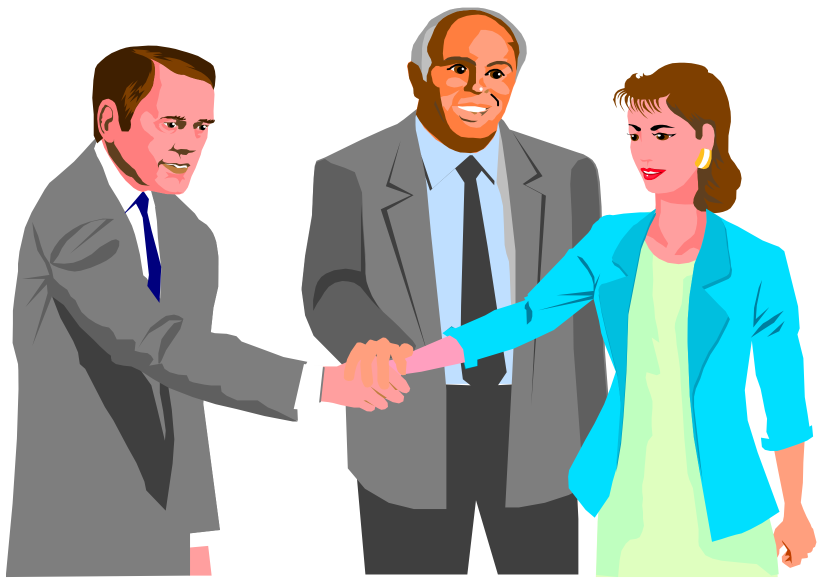

The Trust Scope

- For what something is trusted
- A particular trust scope can for example be
  - *“to be a good car mechanic”*
- Trust scopes can be specific or general
- Trust scopes can be related
  - i.e. if an entity is trusted for a specific scope, it can be assumed trustworthy for other scopes as well
- Hard to determine dependence between trust scopes
- Other terms used with the meaning of trust scope:
  - Trust context, Trust purpose, Subject matter

Classification of trust scopes

- Provision trust
  - Relying party’s trust in a service, or a service provider.
- Access trust
  - Service provider’s trust in entities requesting access to resources and services.
- Identity trust
  - Belief that an entity’s identity is as claimed
- Delegation trust
  - Trust in a agent to make trust decisions on behalf of the relying party
- Context trust
  - Belief that the necessary systems and institutions are in place in order to support a transaction that involves risk

(Source: Grandison & Sloman)

Combination of scope and assessment

- Scope dimension: *“Specificity – generality”*
- Assessment dimension: *“Subjectivity – objectivity”*

Scope: **Assessment:**	**Specific, vector based**	**General, synthesized**
Subjective	Survey questionnaires	eBay, elections
Objective	Product tests	Synthesized general score from product tests, D&B rating

trust and reputation measures, with examples

Trust mechanisms and processes

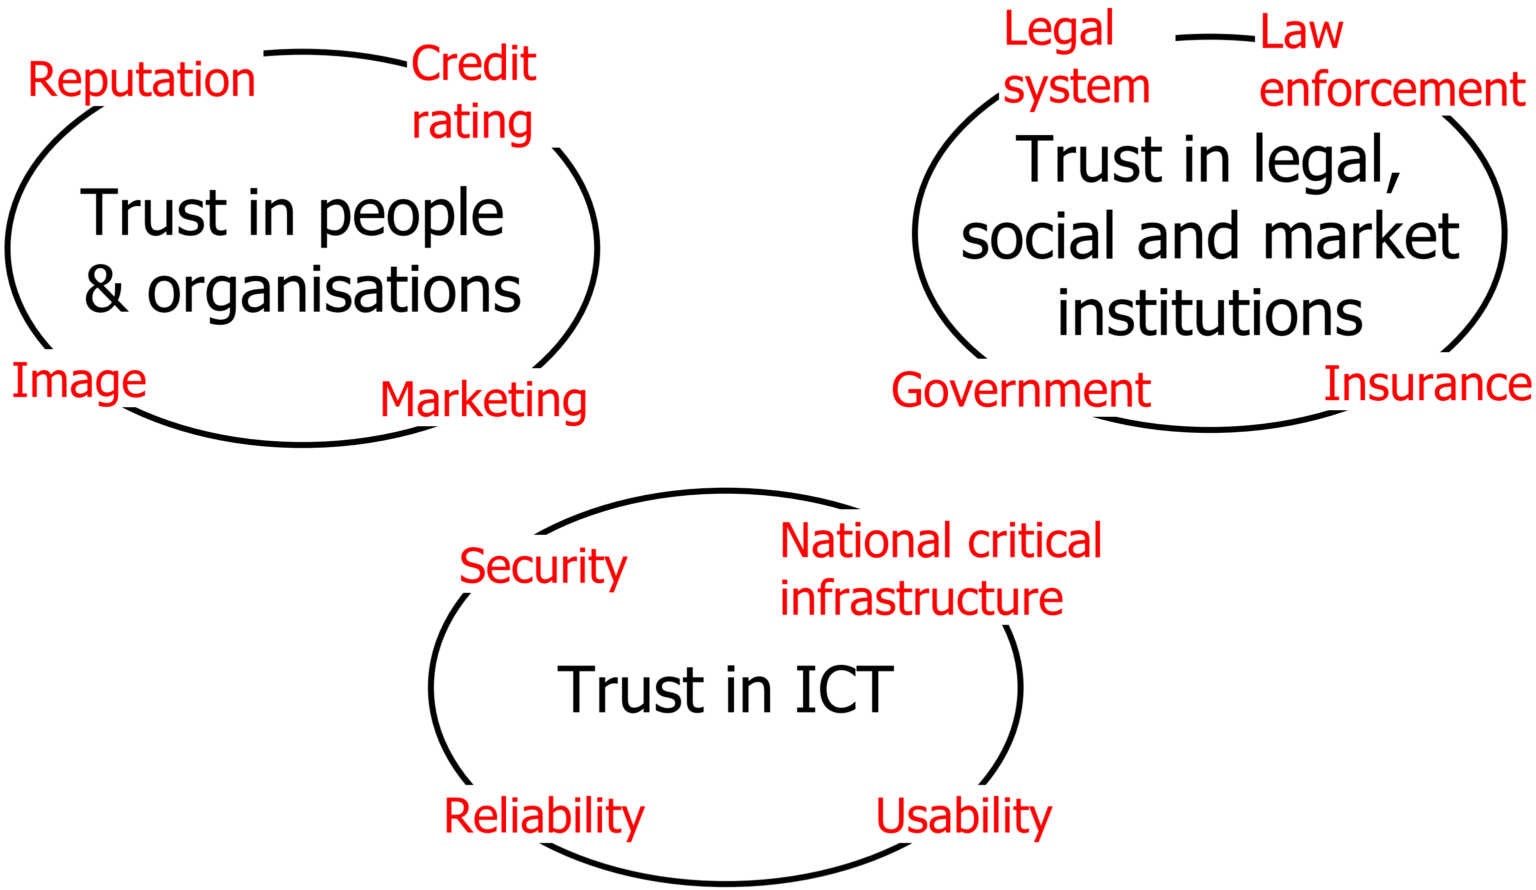

Trust and decision making

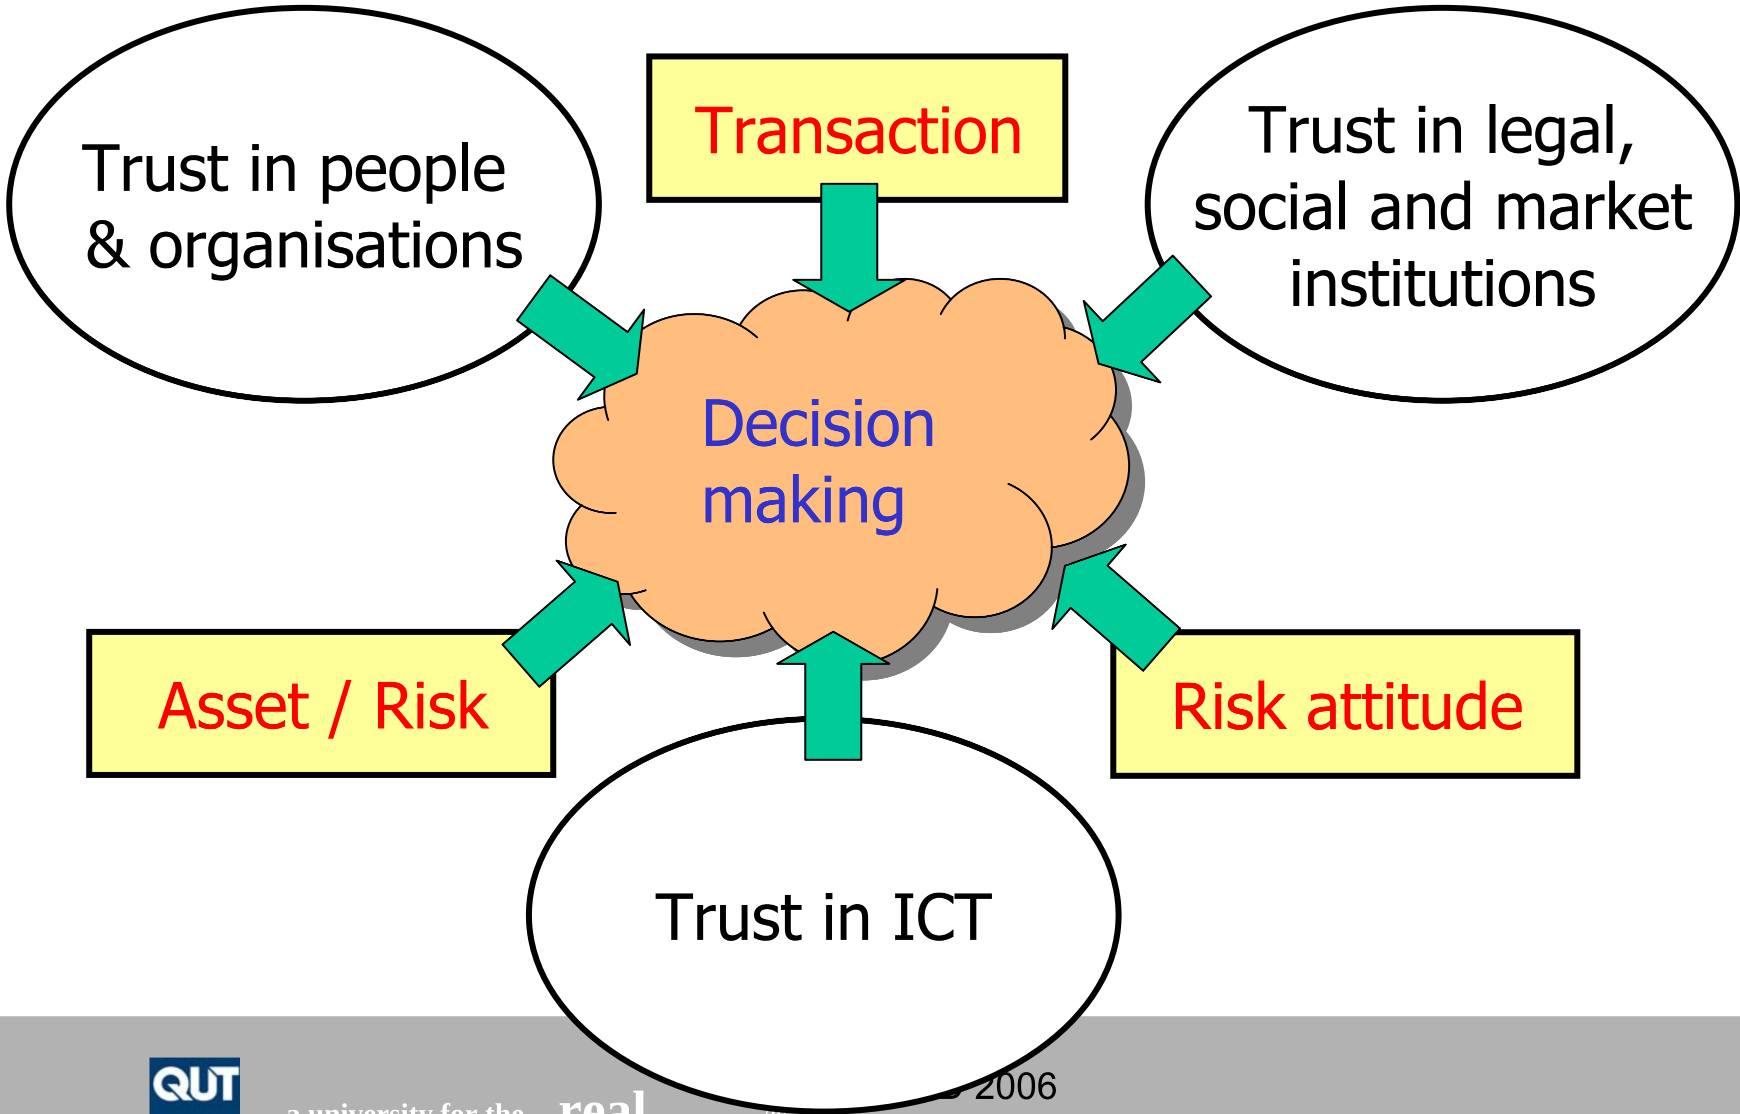

Extrinsic and intrinsic trust

Extrinsic Factors

- Cognitive
- Observed
- Recommendation
- Reputation
- External evidence
- Easy to manufacture

Intrinsic Factors

- Affective
- Experienced
- Intimate relationship
- Internalised pattern
- Take time to build
- Override extrinsic

A model for e-commerce trust

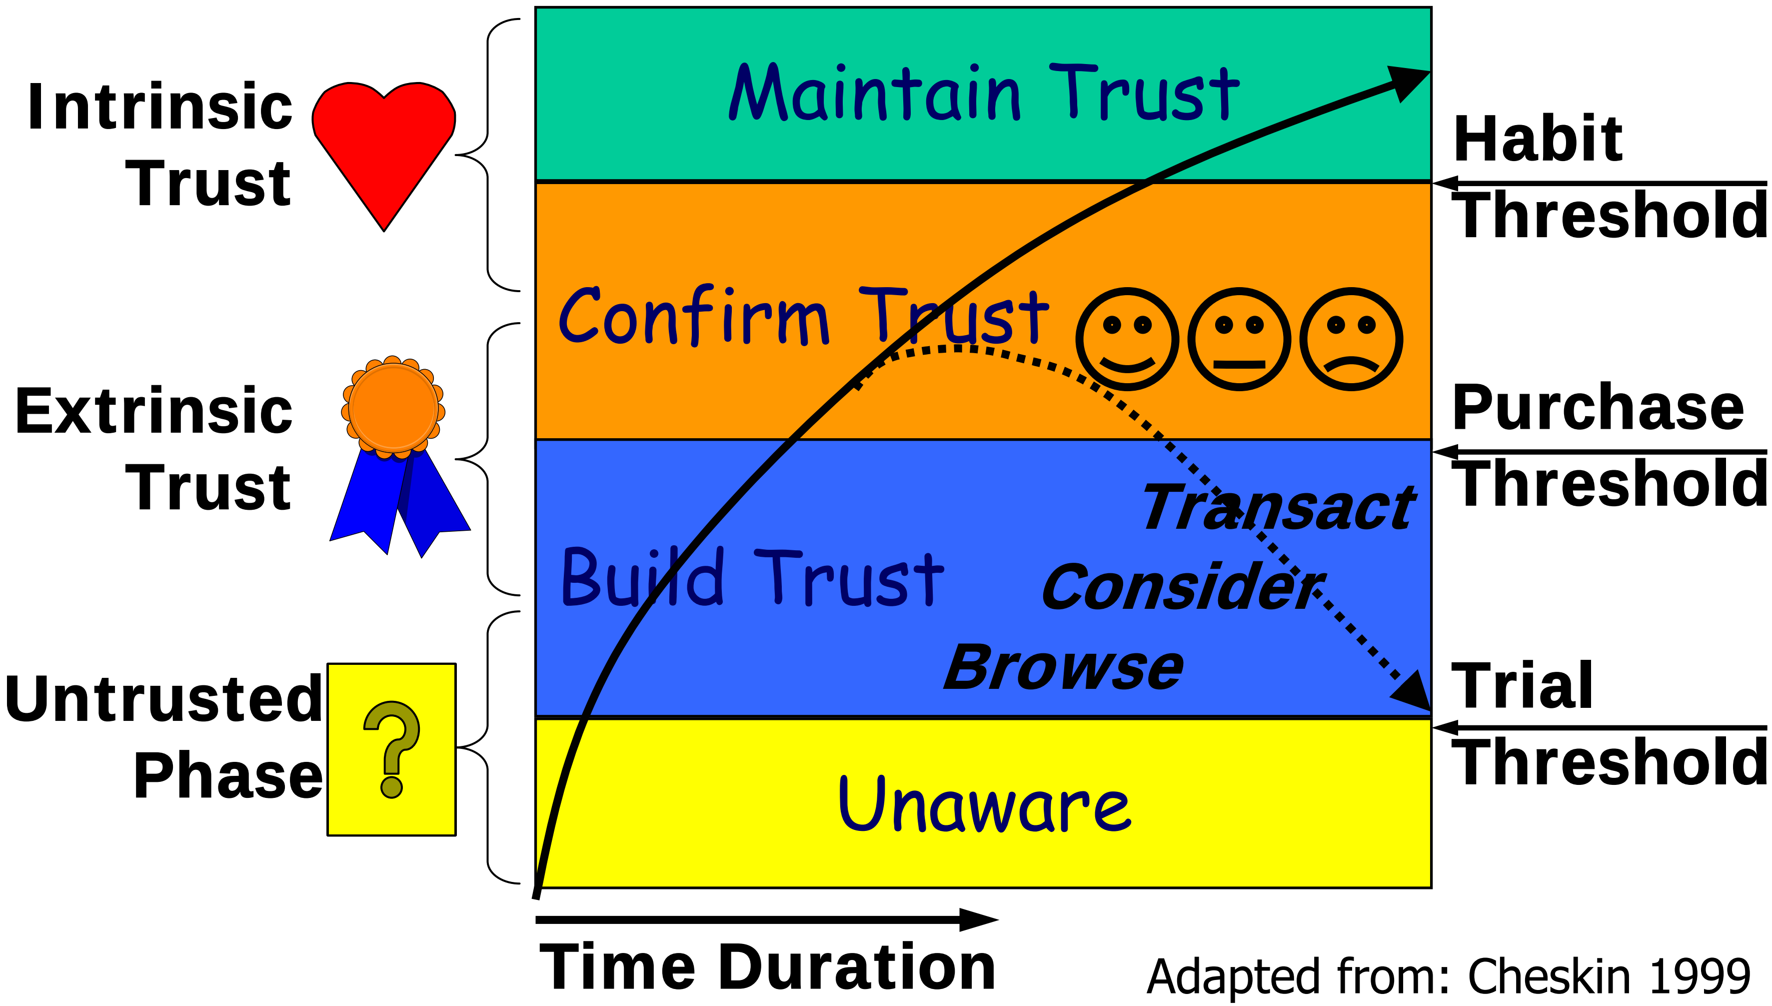

Adapted from: Cheskin 1999

Trust transitivity

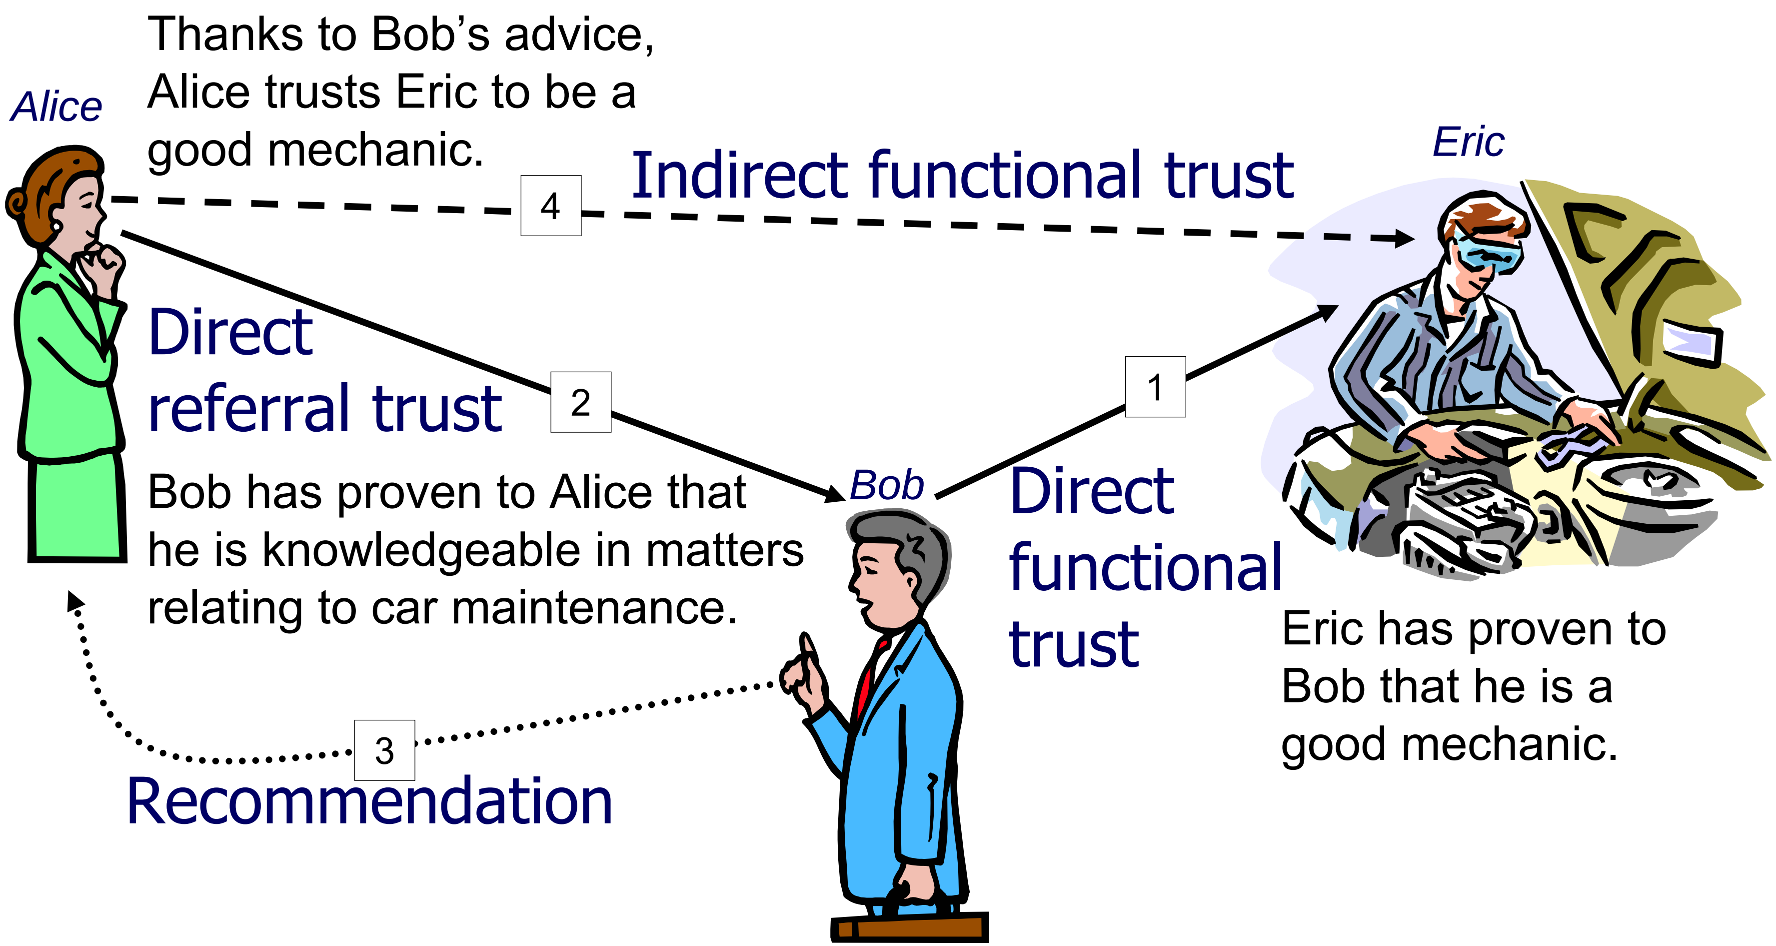

Variants of the same trust scope

1. Functional trust
   – Belief in an entity's ability (and willingness) to carry out or support a specific function (the scope) on which the relying party depends
2. Referral trust
   – Belief in an entity's ability and willingness to recommend another entity with respect to 1).

Trust types of the same scope

- Direct trust
  - Trust resulting from direct experience with the trusted party
- Indirect trust
  - Trust resulting from recommendation from other third parties

Basic trust diversity dimensions

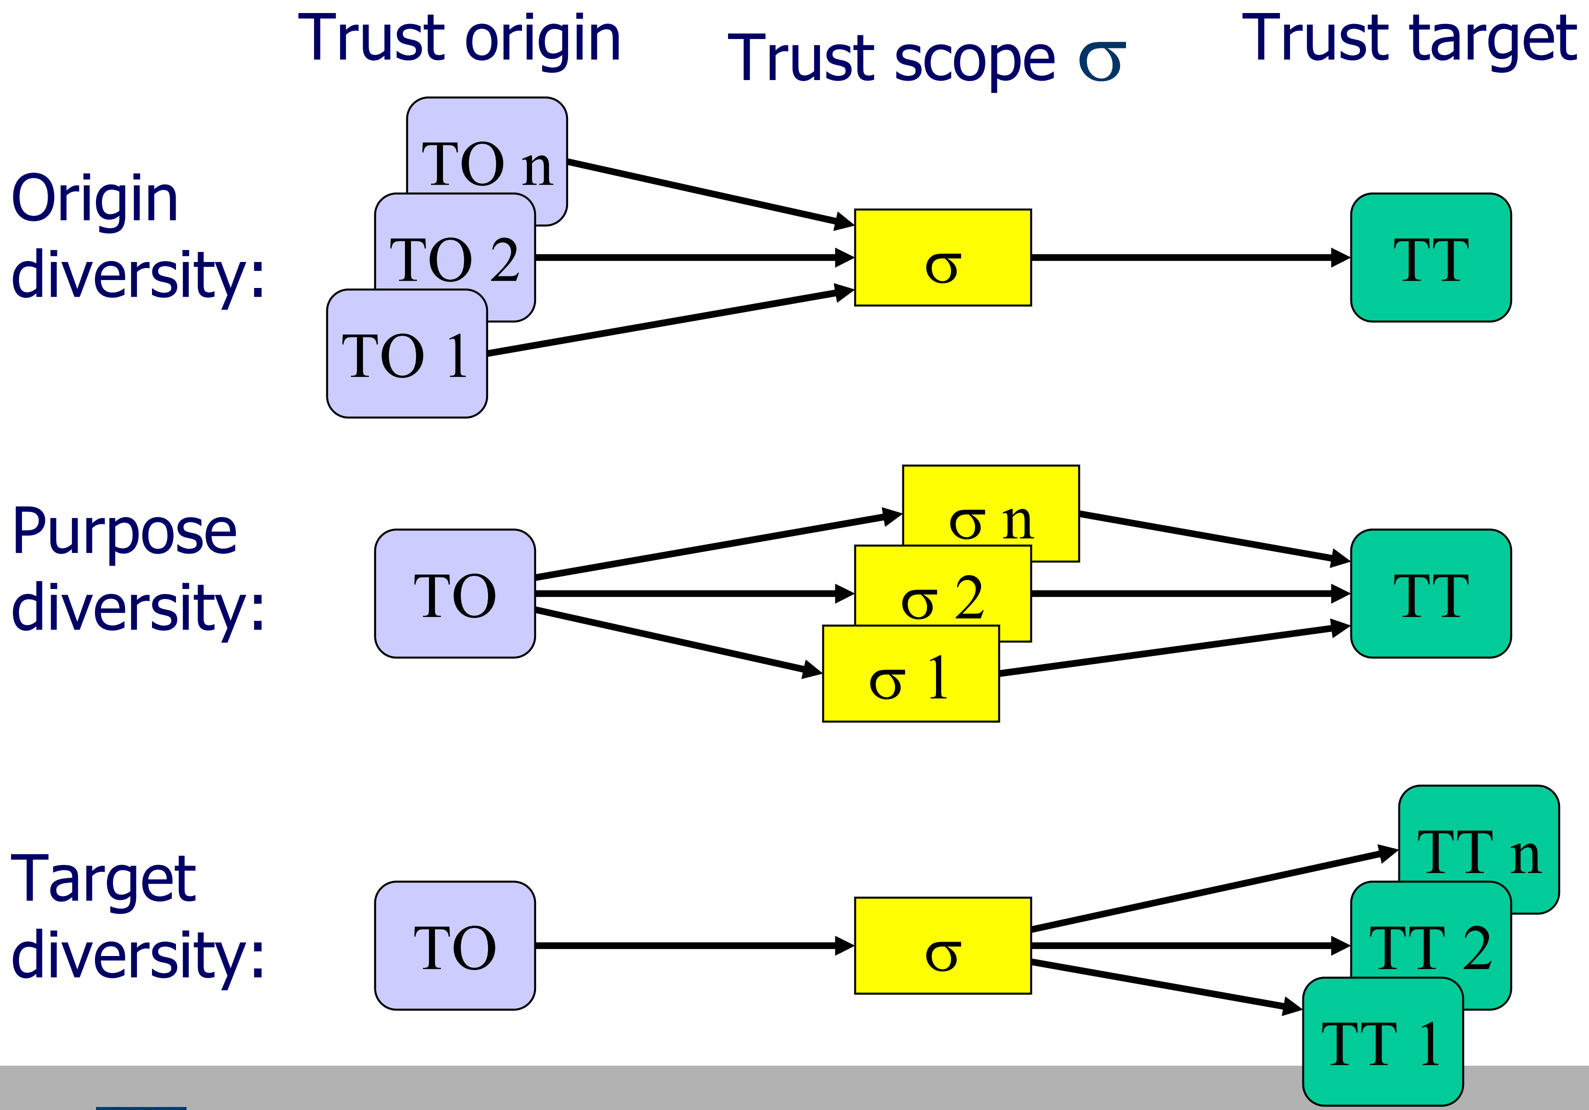

Additional trust dimensions

- Trust measure: μ
  - Binary (e.g. “Trusted”, “Not trusted”)
  - Discrete (strong-, weak-, trust or distrust)
  - Continuous (percentage, probability, belief)
- Time: τ
  - Time stamp when trust was assessed and expressed. Very important as trust generally weakens with temporal distance.

Valid transitive chains

- Every leg in the chain contains the same trust scope [σ]. (It doesn't make any sense otherwise!)
- The last trust link is **direct functional** trust [dfσ].
- All other trust links are **direct referral** trust [drσ].

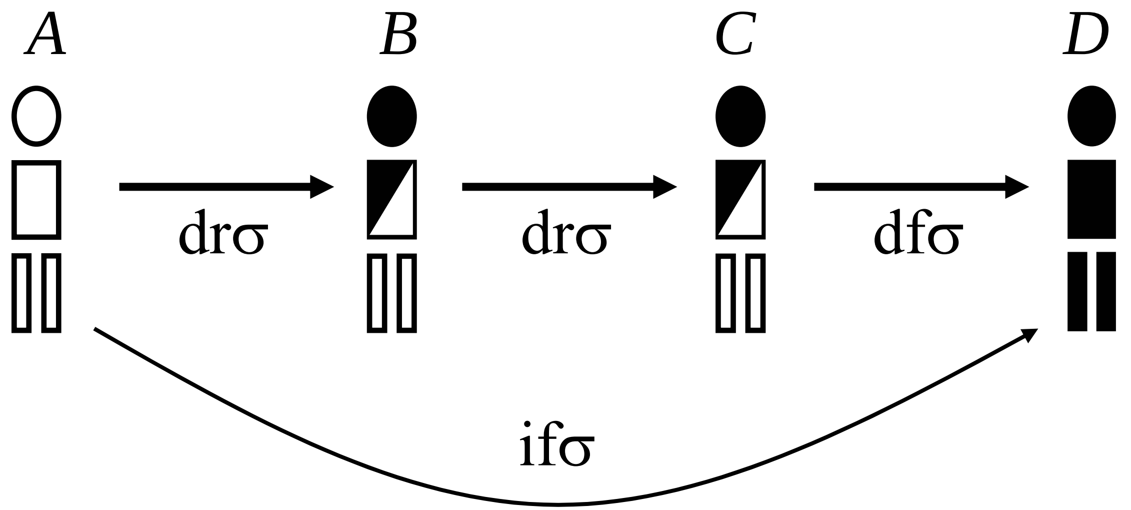

Trust transitivity

Trust is diluted in a transitive chain.

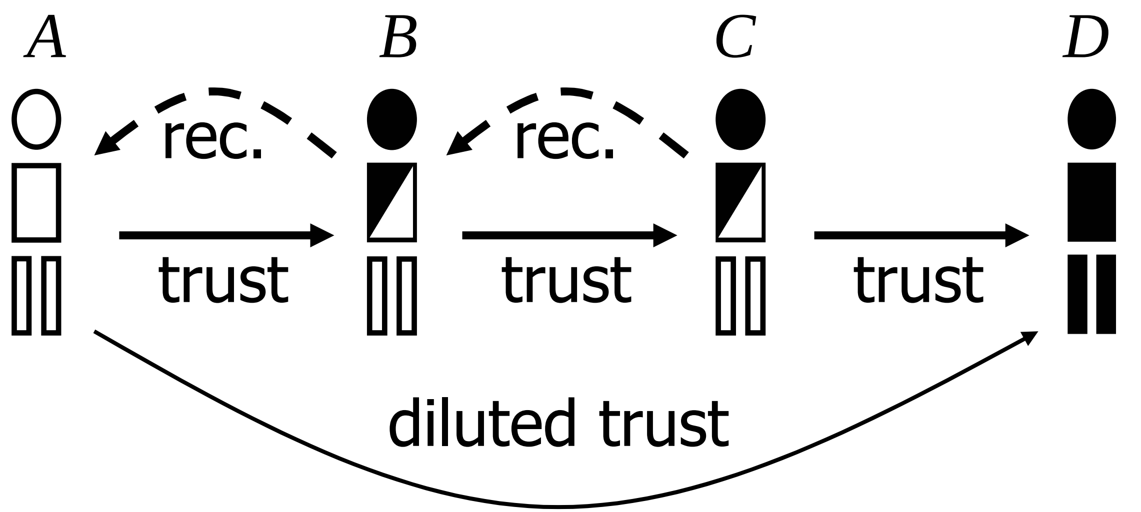

Computed with the discounting operator of subjective logic

Graph notation: $[A, D] = [A, B] : [B, C] : [C, D]$

Explicit notation: $[A, D, \text{if}\sigma] = [A, B, \text{dr}\sigma] : [B, C, \text{dr}\sigma] : [C, D, \text{df}\sigma]$

Trust fusion

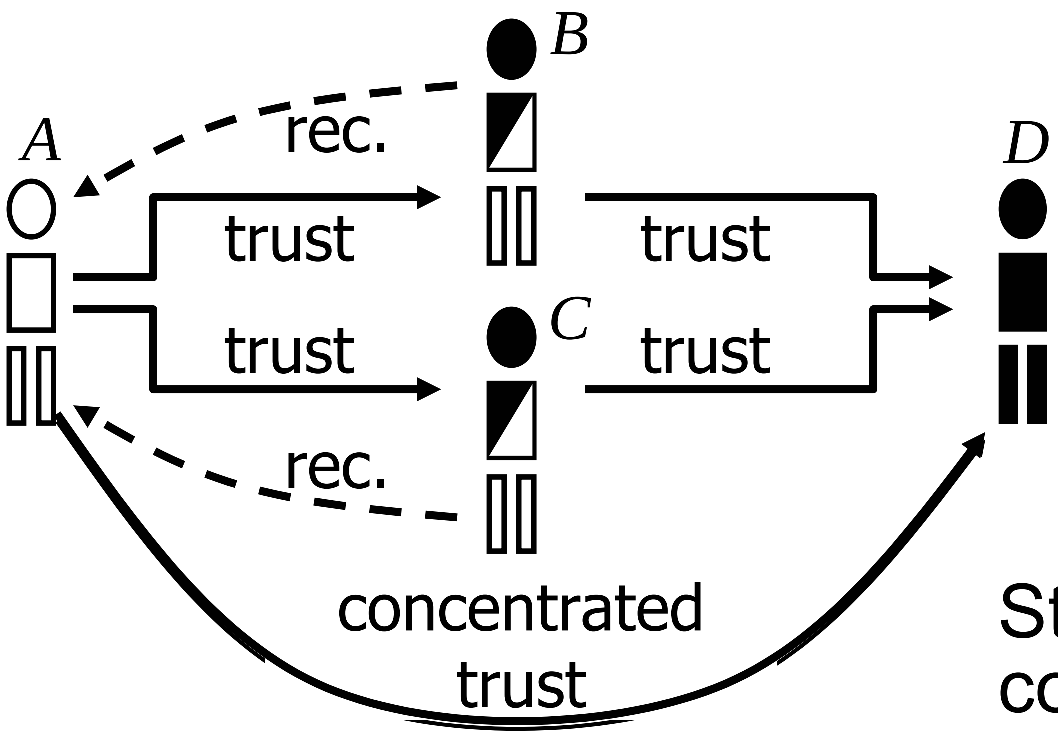

Strengthens trust confidence

Computed with the consensus operator

Graph notation: $[A, D] = ([A, B] : [B, D]) \Diamond ([A, C] : [C, D])$

Non-distributivity of serial and parallel trust

Discounting is non-distributive on consensus.

$$\omega_E^{A:((B:C),(B:D))} \neq \omega_E^{(A:B:C),(A:B:D)}$$

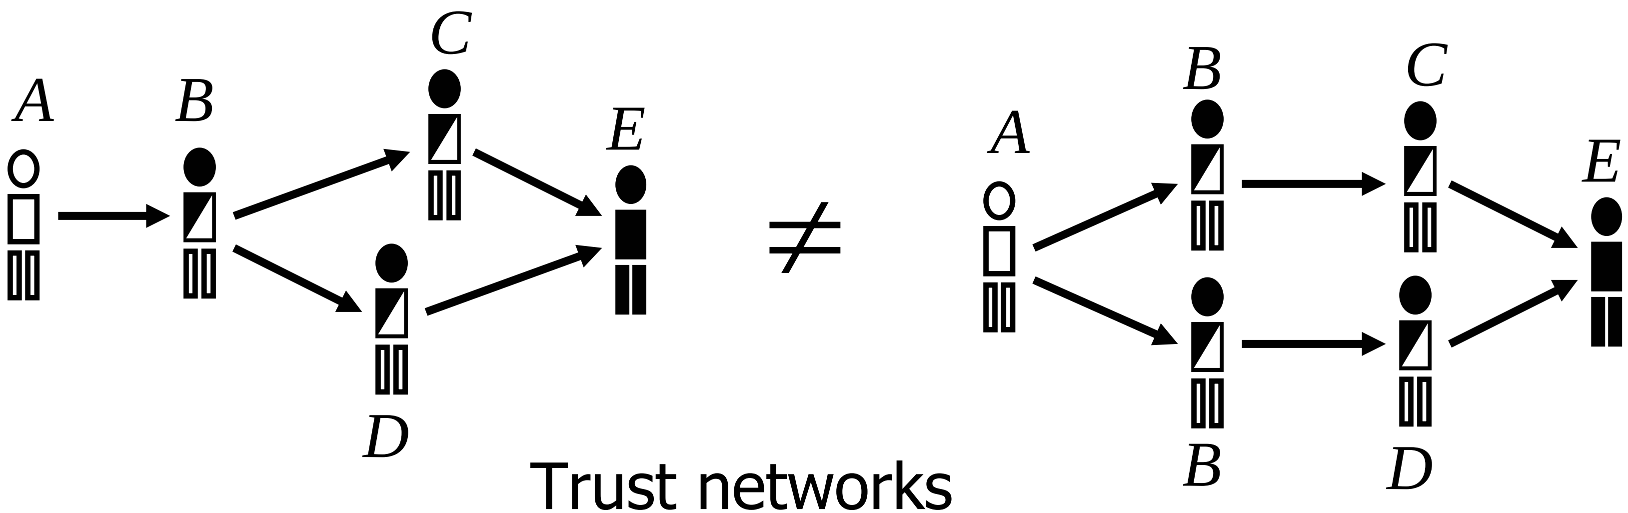

Trust networks

Indirect referral trust

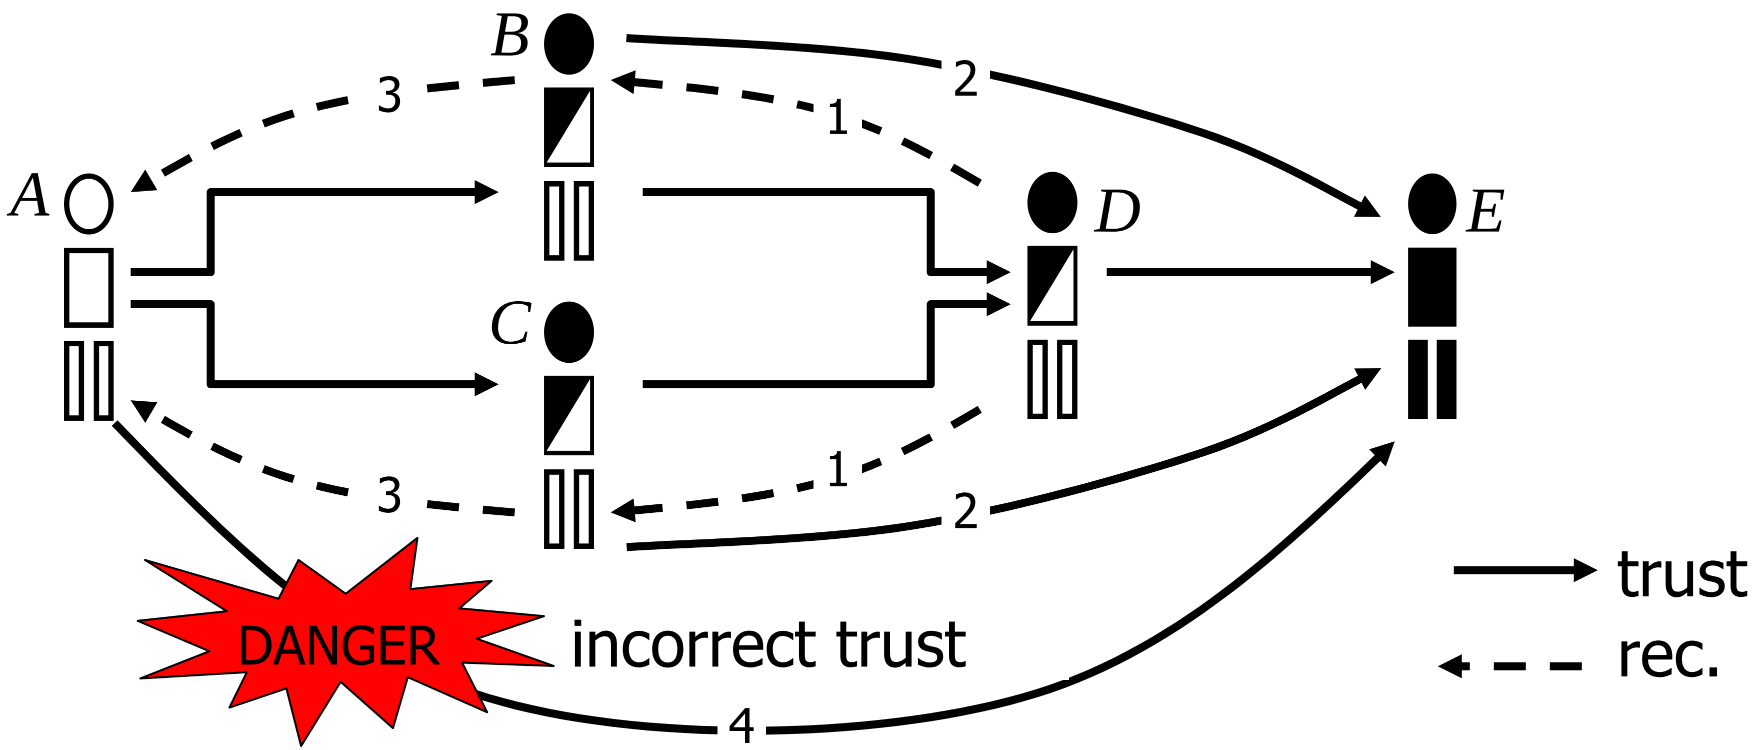

Perceived $[A, B] : [B, E] \lozenge [A, C] : [C, E]$ (OK)

Reality: $[A, B] : [B, D] : [D, E] \lozenge [A, C] : [C, D] : [D, E]$ (not OK)

Hidden and perceived topologies

Perceived topology:

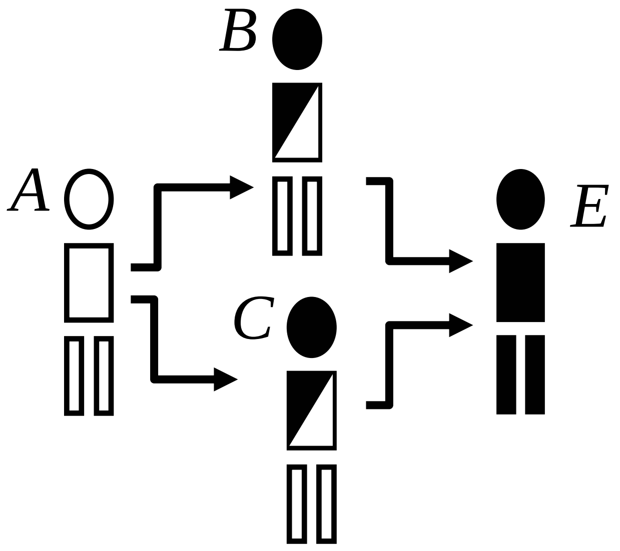

Hidden topology:

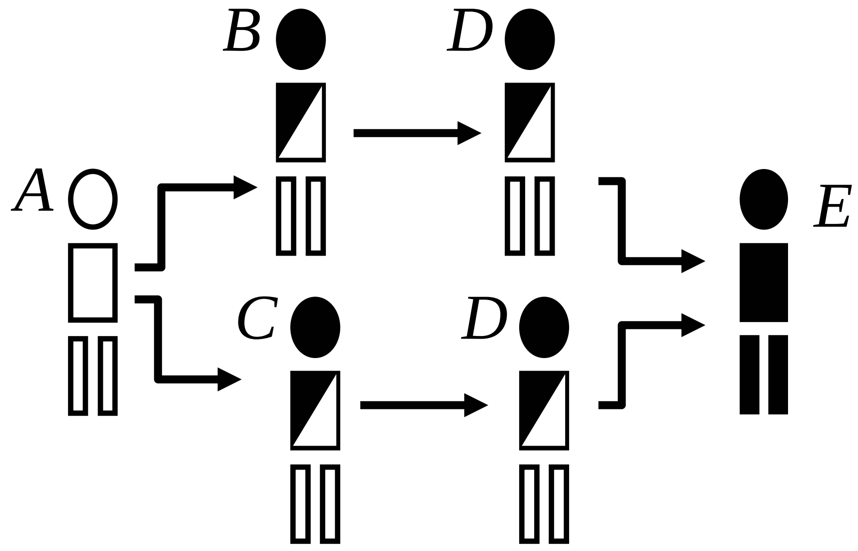

$$[A, B] : [B, E] \ \Diamond \ [A, C] : [C, E]$$
$$\neq [A, B] : [B, D] : \mathbf{[D, E]} \ \Diamond \ [A, C] : [C, D] : \mathbf{[D, E]}$$

(D, E) is taken into account twice

Indirect referral trust

(corrected)

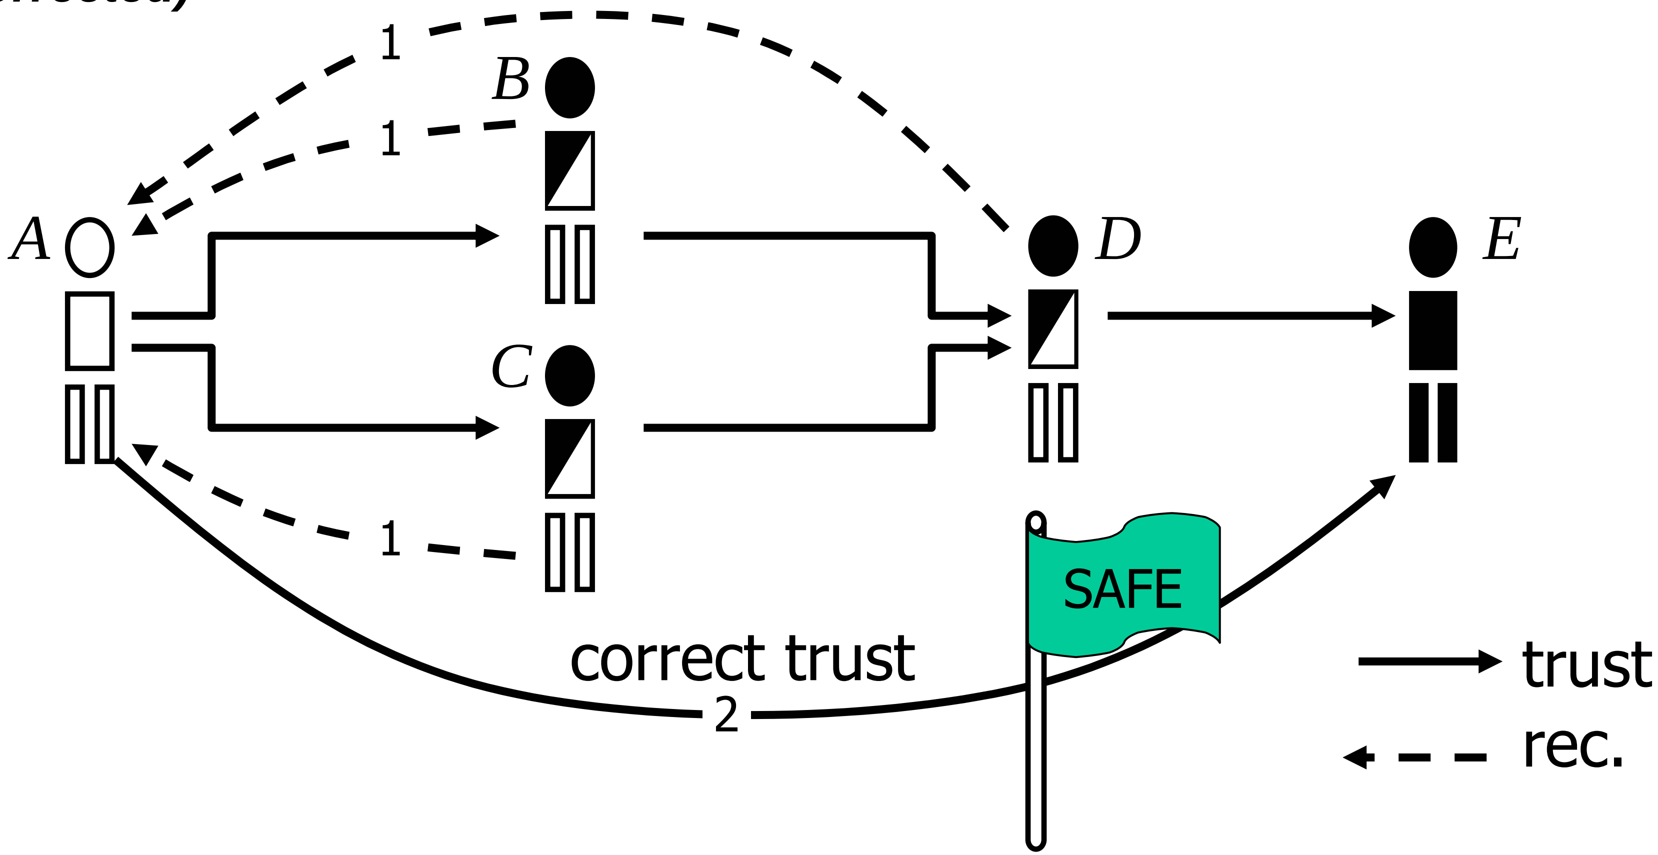

Perceived and real topology (OK):

$$([A, B] : [B, D] \Diamond [A, C] : [C, D]) : [D, E]$$

PKI and trust transitivity

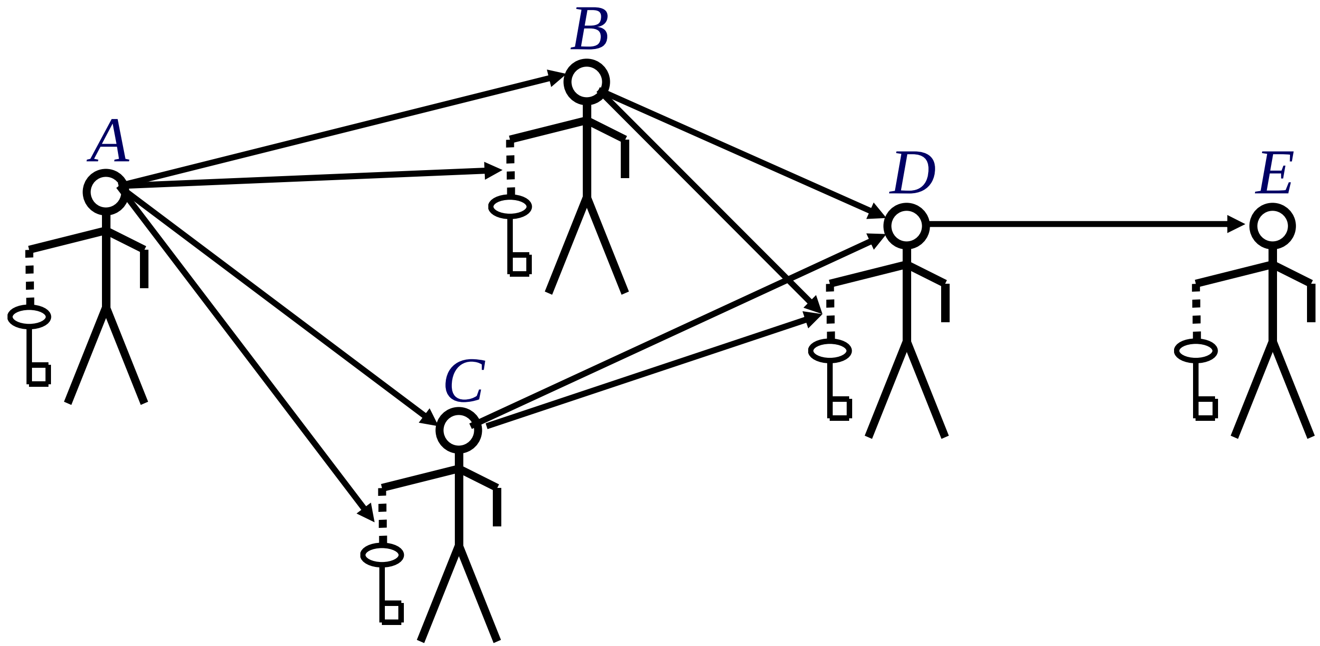

- Trust in public keys must be included in the analysis.
- Separate topology analysis for determining trust in each public key.

Principles for building trust and reputation systems

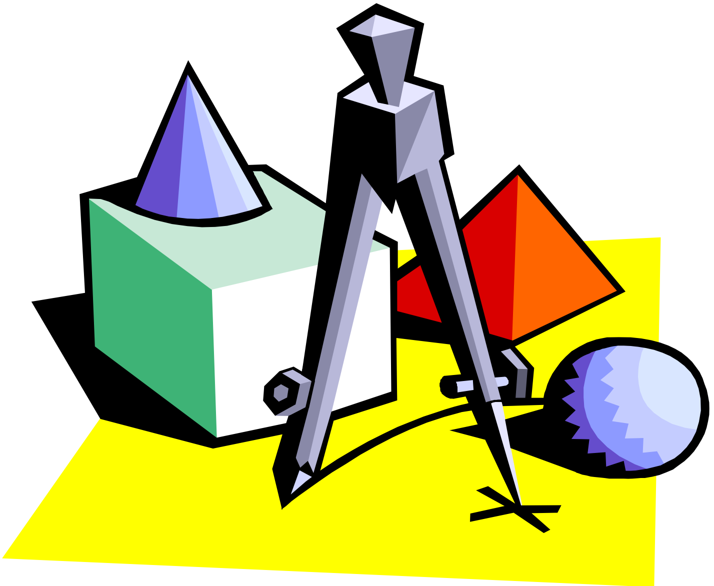

Difference between real and online world

	Availability and richness of evidence for trust	Efficiency in the communication and processing
Real world	Very good	Poor
IT / Online	Poor	Very good

- Communication of trust information often restricted to local community in the real world
- The online world currently provides very little reliable trust evidence

Basis for trust and rep. systems

- Focus on the trust evidence and on the methods for collecting this information
  - Find substitutes for traditional information used in physical world
  - Create new types of evidence
  - Application specific
- Exploit the efficiency of IT and the Internet for
  - Collection of information
  - Processing
  - Dissemination

Centralised reputation system

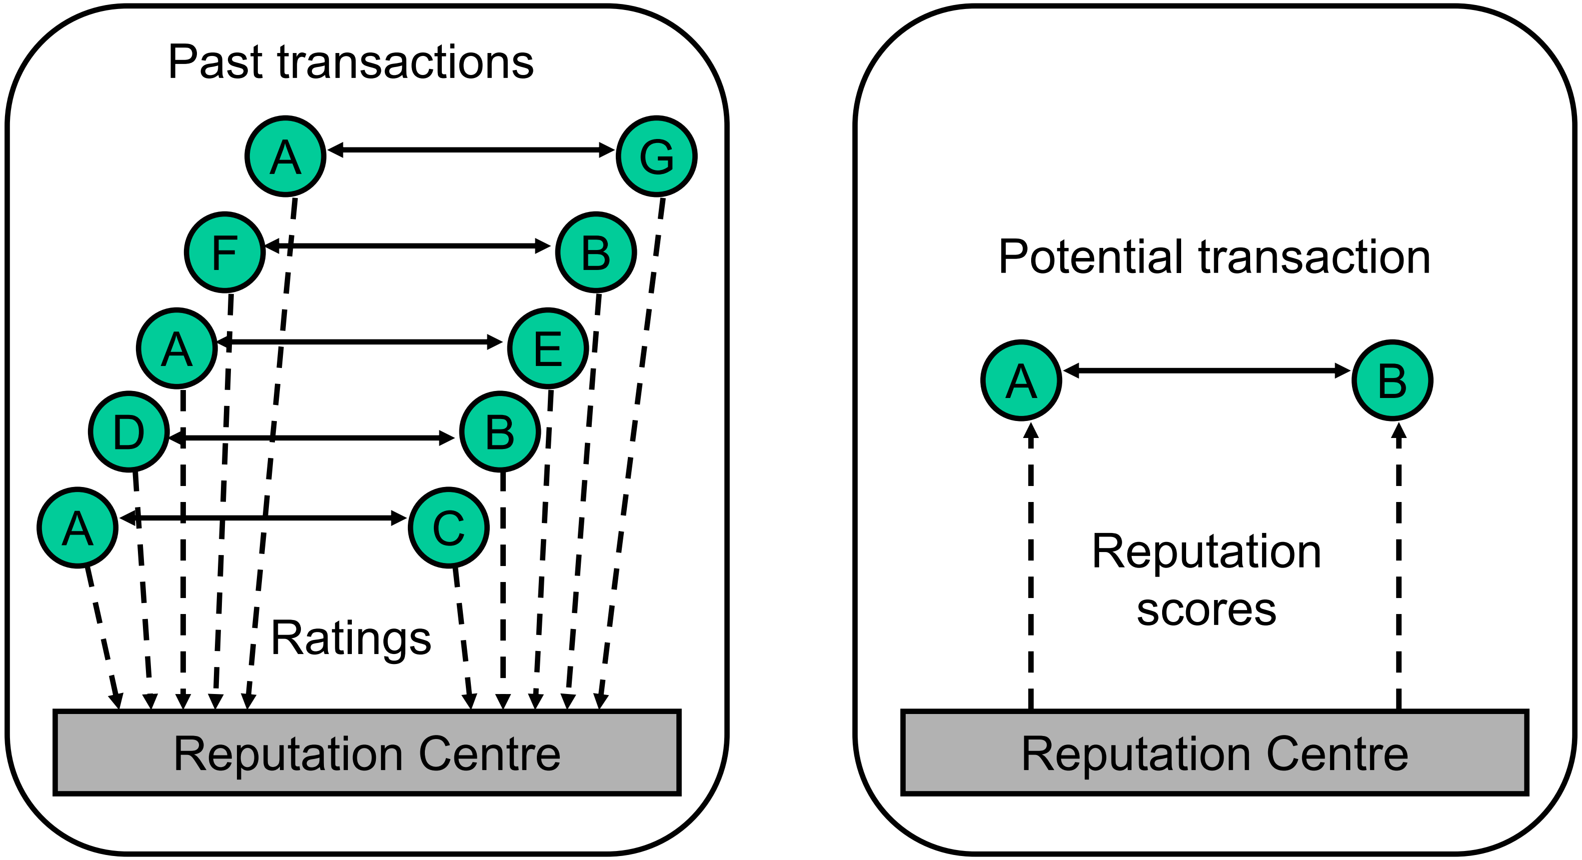

a) Past

b) Present

Distributed reputation system

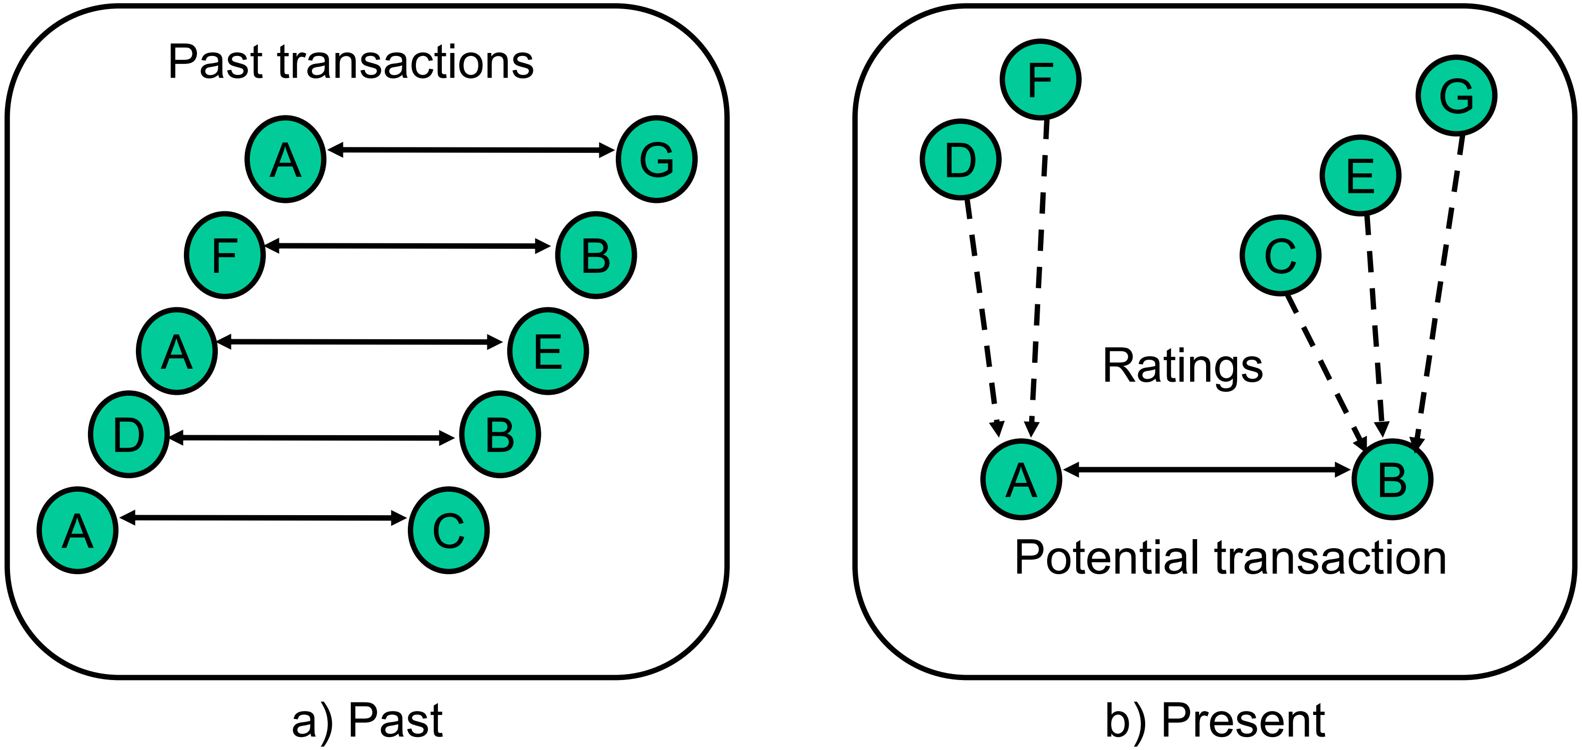

Applications of reputation systems

- e-Auctions
- P2P networks
- Software agent communities
- Contract negotiations
- Online markets: B2C, B2B, C2C
- Web service search and selection
- Information/intelligence gathering

Market Efficiency Experiment

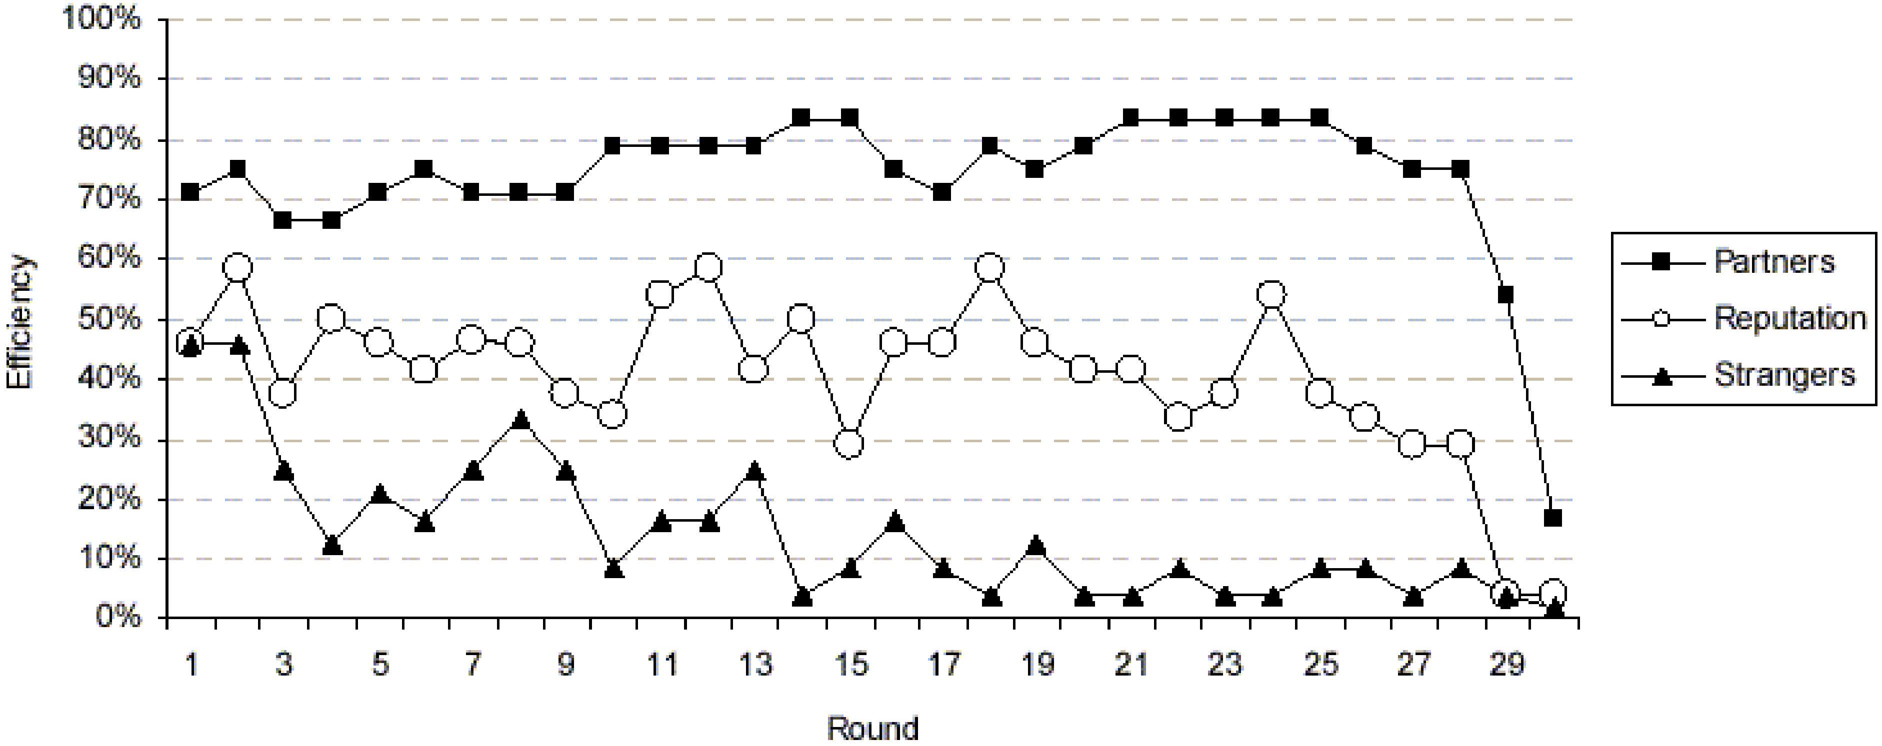

Source: Bolton,Katok,Ockenfels,2002

P2P networks

- P2P Networks: servent = server + client
- Search phase: discover resources
  - Centralised: e.g. Napster, with central directory
  - Pure distributed: Gnutella, Freenet
  - Semi-distributed: FastTrack, KaZaA, grokster, with distributed directory servers
- Download phase: get the resources
- Problems
  - Spreading malware
  - Free riding
  - Poisoning

Gnutella example

- Pure distributed search phase

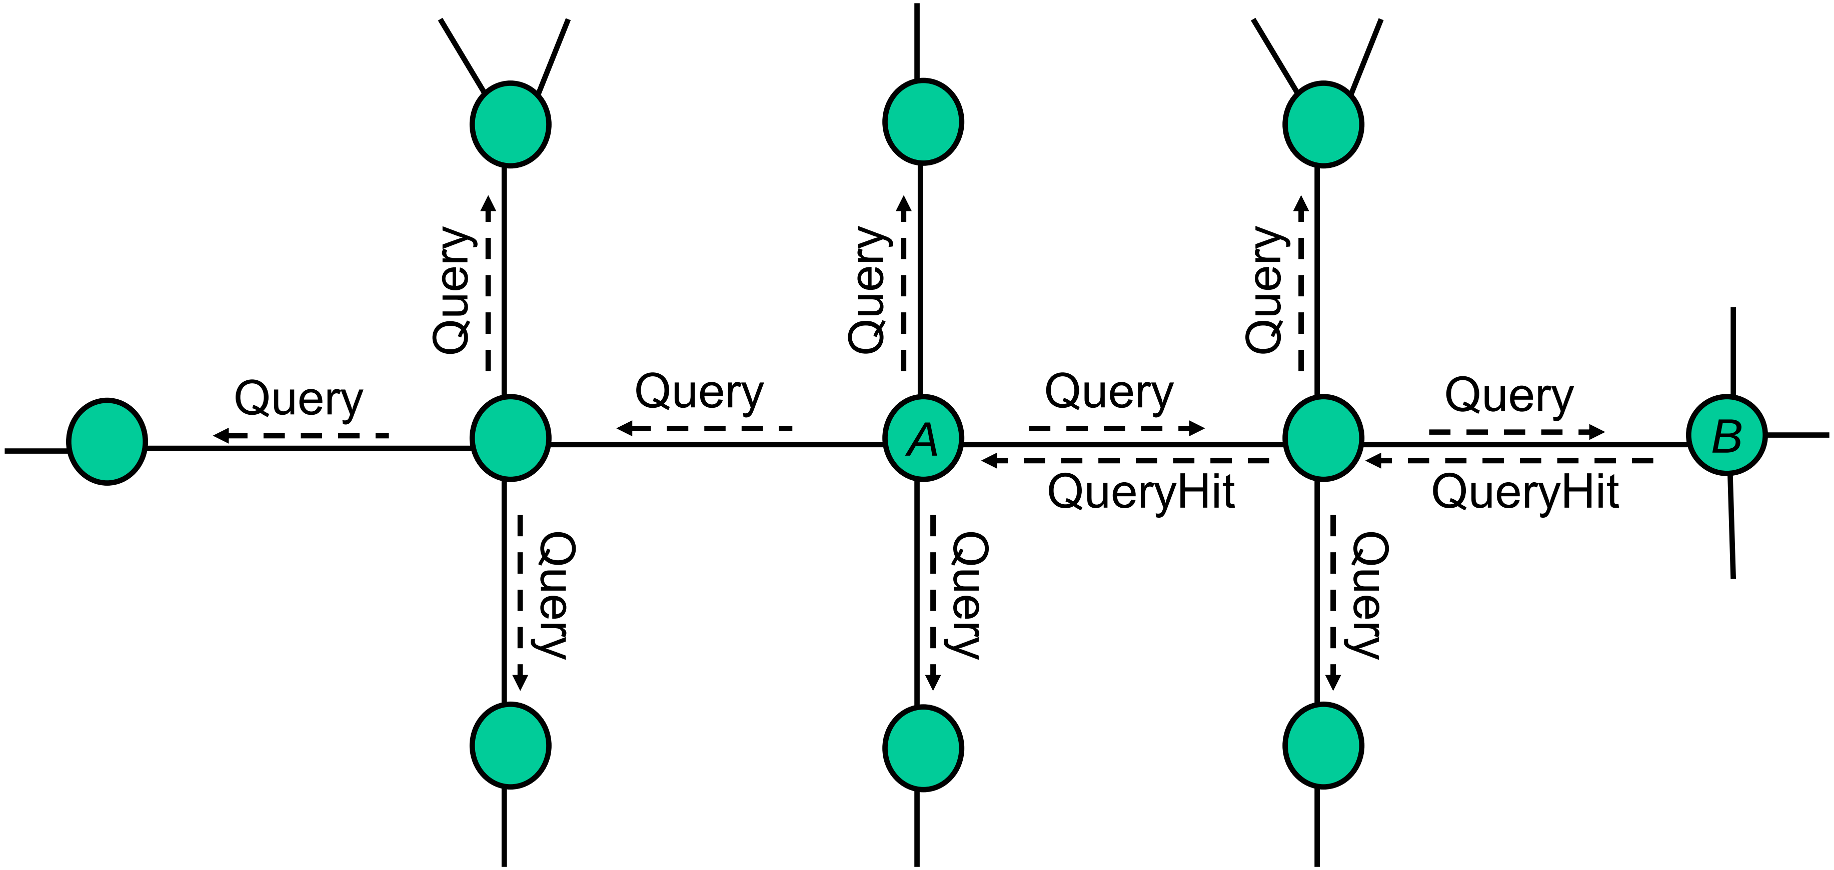

Reputation systems in P2P

- Purpose of Reputation systems in P2P
  1. Identify most reliable servents with best quality resources
  2. Determine which servents provide most reliable information w.r.t. 1.
- Reduces or eliminates existing problems
  - Many theoretical proposals
  - Few practical implementations

Reputation/trust system with Gnutella

- XRep proposed by Damiani *et al.*

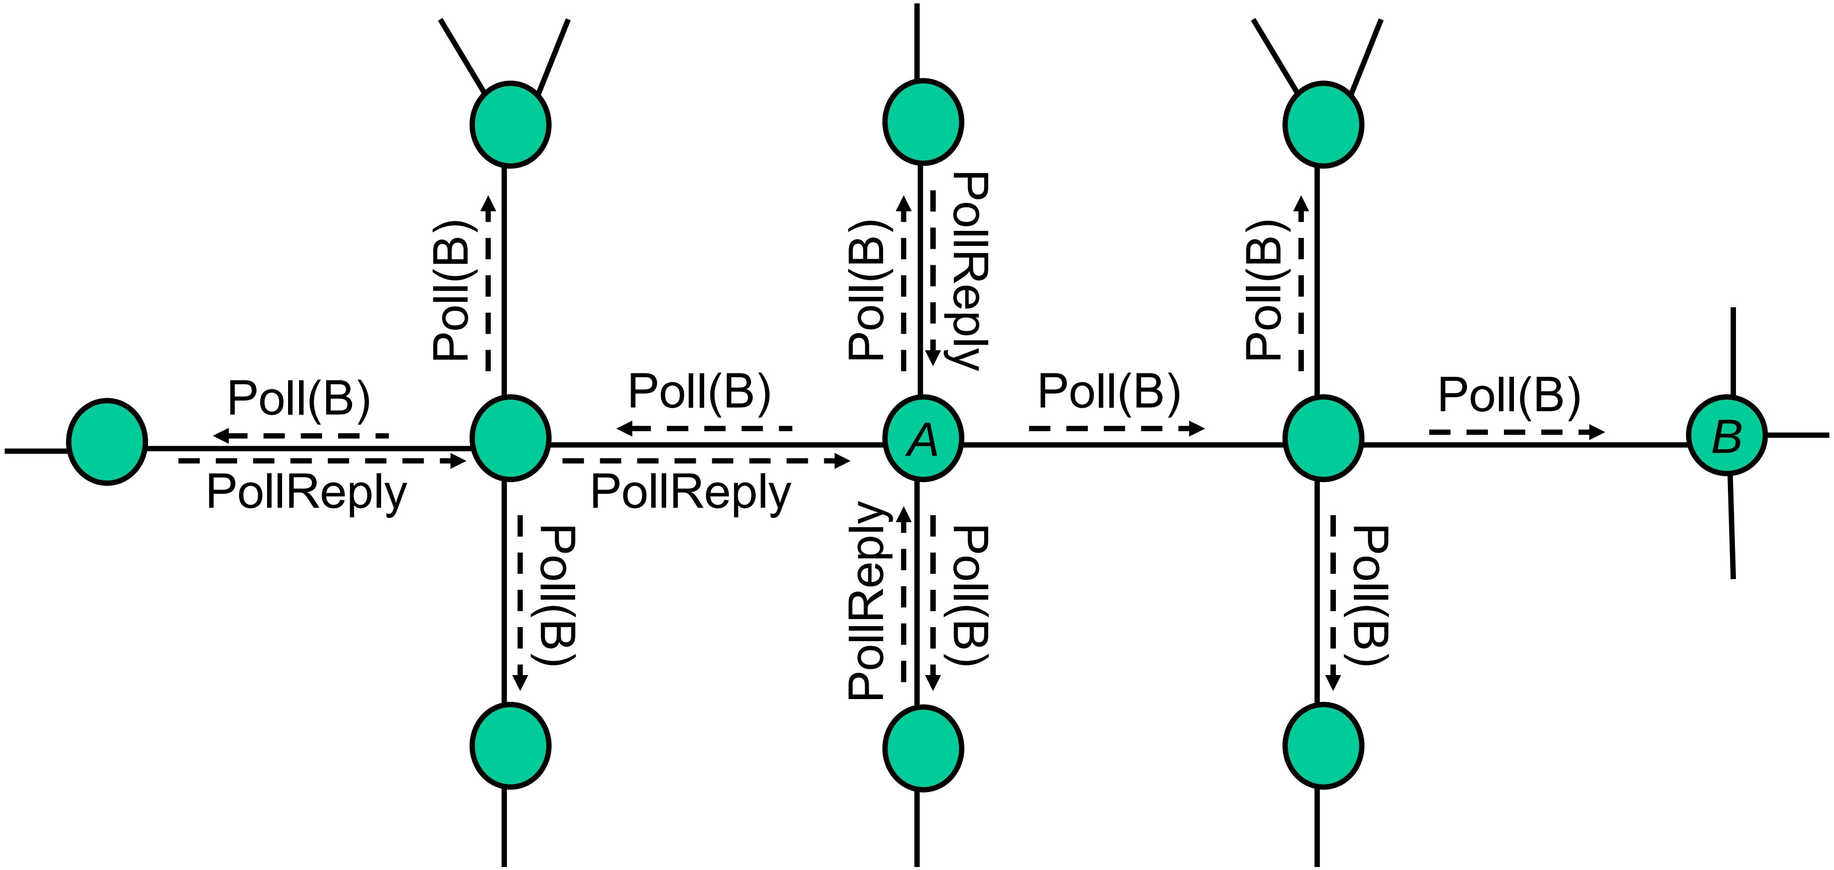

Trust and reputation computation engines

- Summation or average
- Bayesian models
- Discrete models
- Belief models
- Fuzzy models
- Flow models

Summation and average

- Summation
  - Reputation score = Σ(positive) - Σ(negative)
  - E.g. eBay
- Average
  - Reputation score = Σ(ratings)/N(ratings)
  - E.g. Epinions
- Can be combined with sliding time windows
- Simple to understand
- Can give false impression of reputation

Bayesian Reputation Systems

- Theoretically sound rating algorithm.
- Binomial and multinomial models.
- Rating possibilities:
  - any range,
  - combination,
  - discounting,
  - longevity,
  - weight ~ transaction value.

Binomial model

- Based on the Beta PDF

$$f(p \mid \alpha, \beta) = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} p^{\alpha-1}(1-p)^{\beta-1} \qquad 0 \le p \le 0 \qquad \alpha,\beta > 0$$

- Probability expectation: $\mathrm{E}(p) = \dfrac{\alpha}{\alpha+\beta}$

- The Beta PDF naturally expresses the probability of binary events.

Beta PDF of Binary Events

Define: $\begin{cases} \alpha = r + 1 \\ \beta = s + 1 \end{cases}$ where: $r, s \geq 0$

r: positive observations

s: negative observations

Example: uniform density

$r = 0, \; s = 0, \quad E(p) = 0.5$

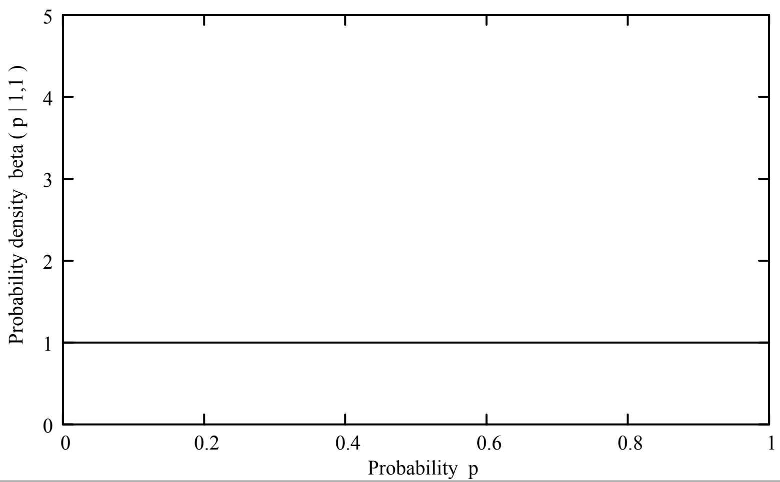

Example:

$r = 7, \; s = 1, \quad E(p) = 0.8$

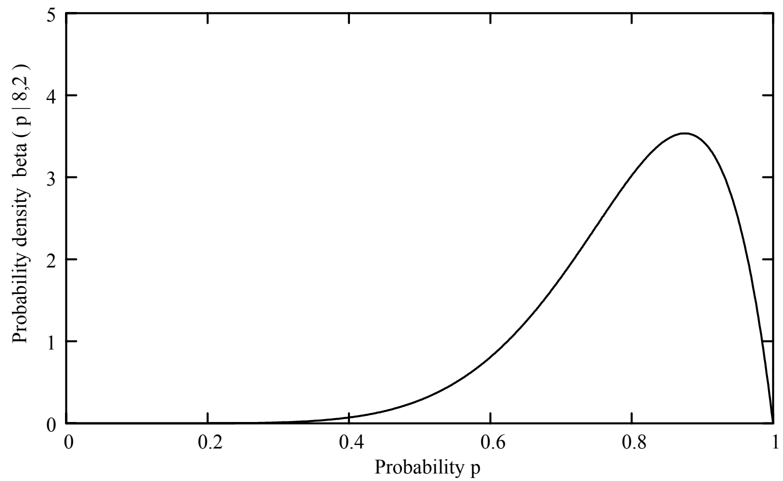

Binomial reputation score

- Based on probability expectation of Beta-PDF.

$$Sc(Z) = \frac{r_{\text{base}} + \sum r(Z)}{r_{\text{base}} + s_{\text{base}} + \sum r(Z) + \sum s(Z)}$$, where: $S(Z) \in [0.1]$

- $Sc(Z)$: reputation score of Z
- $\Sigma r(Z)$: positive evidence sum.
- $\Sigma s(Z)$: negative evidence sum.
- r_{base} , s_{base} : default base rate parameters.

Limitation: Unable to reflect polarised ratings!

Computing binomial reputation over time with longevity factor

- R_i : accumulated positive evidence at time i
- S_i : accumulated negative evidence at time i
- r : positive evidence during 1 time period
- s : negative evidence during 1 time period
- λ : longevity factor in range [0,1]
- $R_{i+1} = \lambda \cdot R_i + r$: Recursive updating algorithm
- $S_{i+1} = \lambda \cdot S_i + s$: Recursive updating algorithm
- $Sc_i(Z) = \dfrac{r_{\text{base}} + R_i(Z)}{r_{\text{base}} + s_{\text{base}} + R_i(Z) + S_i(Z)}$: Score at time period i
- Typically, $r_{\text{base}} = 1$, $s_{\text{base}} = 1$

Score evolution with different longevity

Period 1-25: Positive rating, $r = 1,\ s = 0$

Period 26-50: Negative rating, $r = 0,\ s = 1$

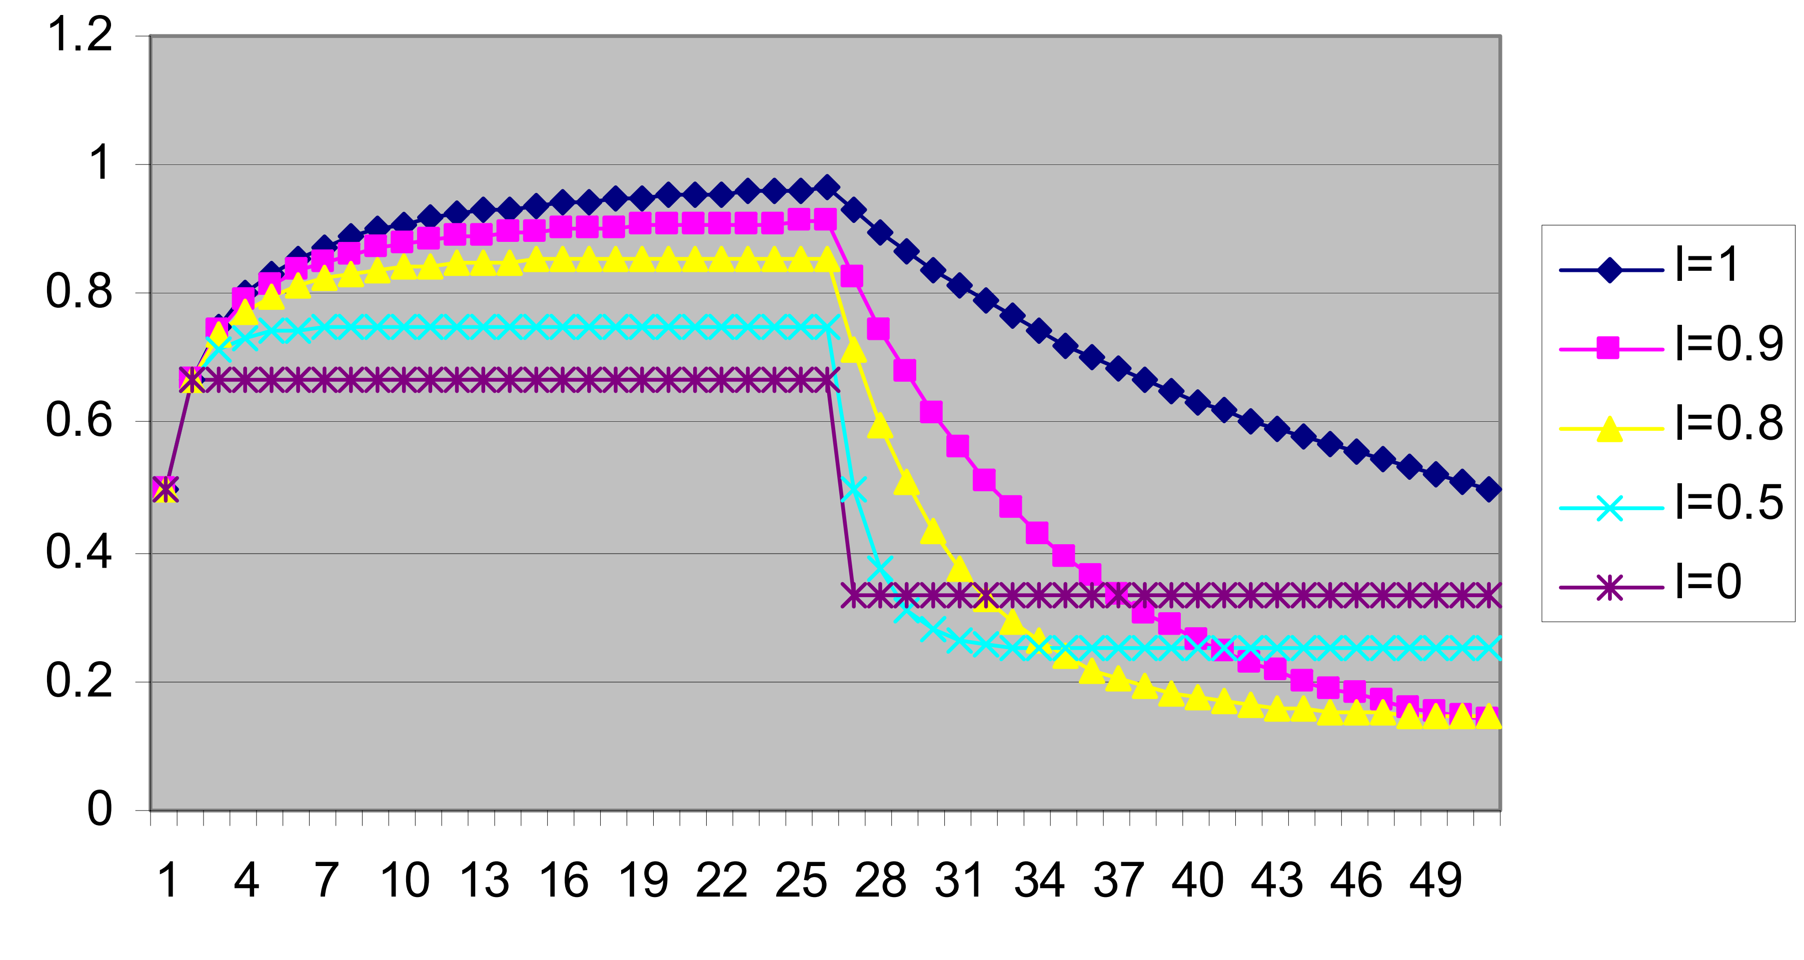

e-Market Simulation

- Buyers:
  - Vary risk aversion
  - Provide ratings
- Sellers:
  - Vary honesty and price
  - Have reputation score
- Sellers and buyers seek to maximise own profit

Seller honesty with fixed reputation

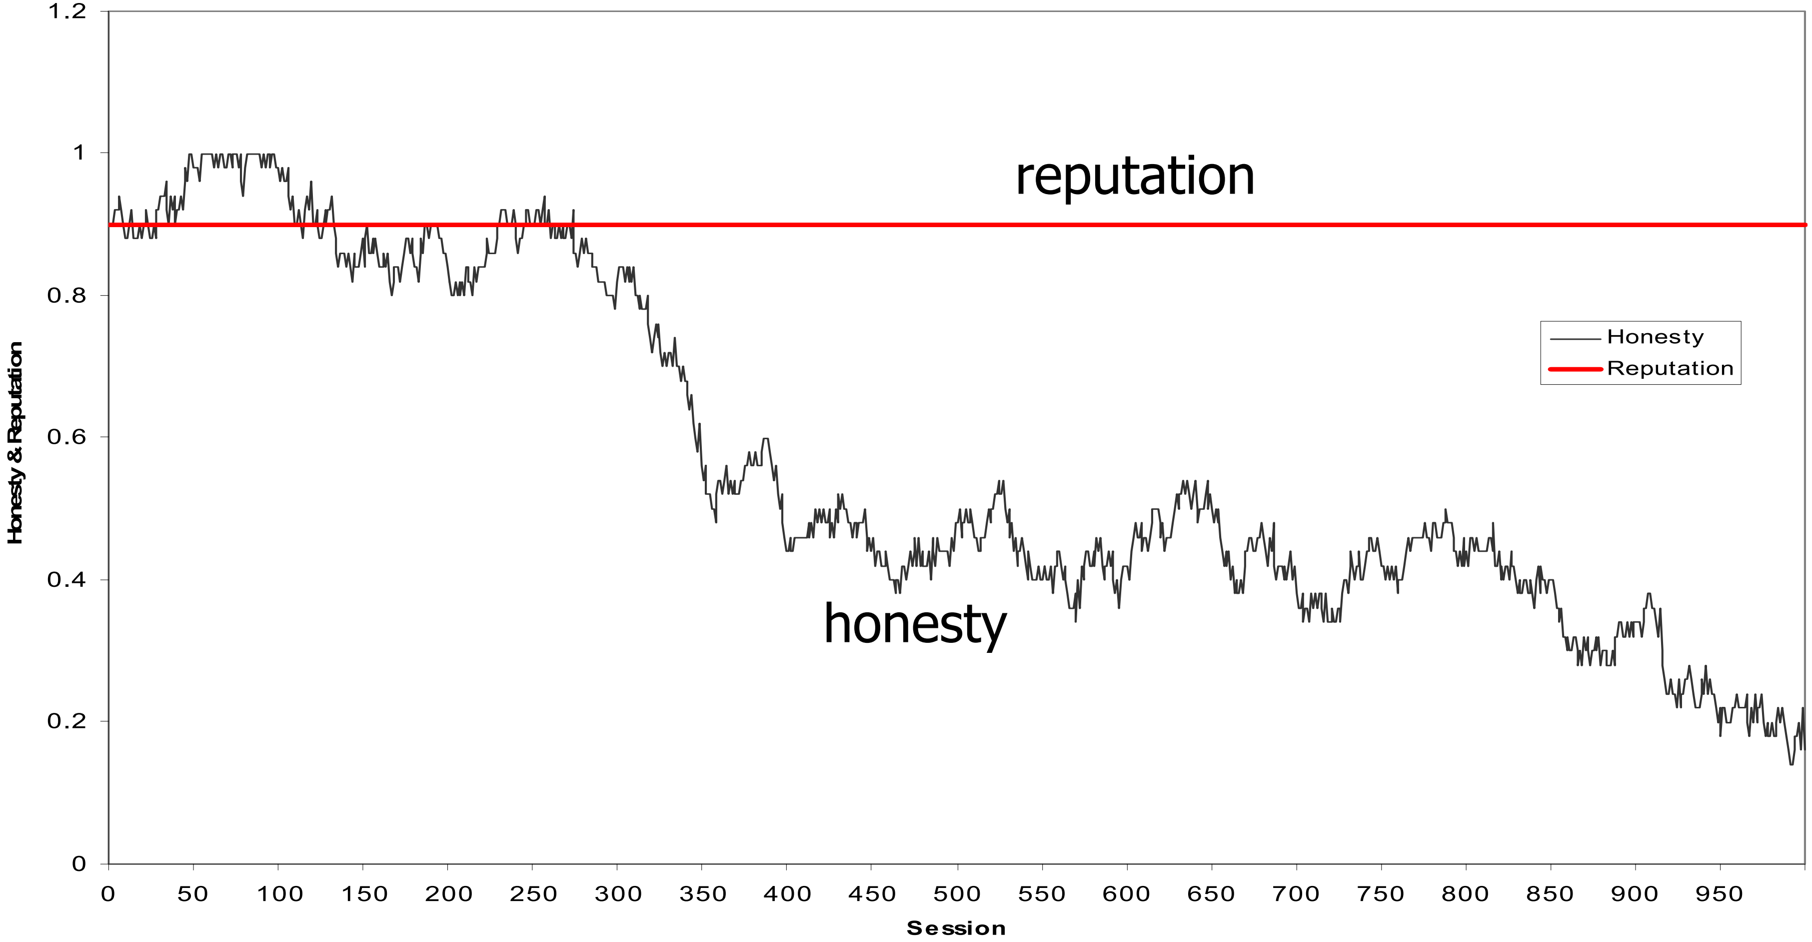

Seller honesty and reputation with λ=1

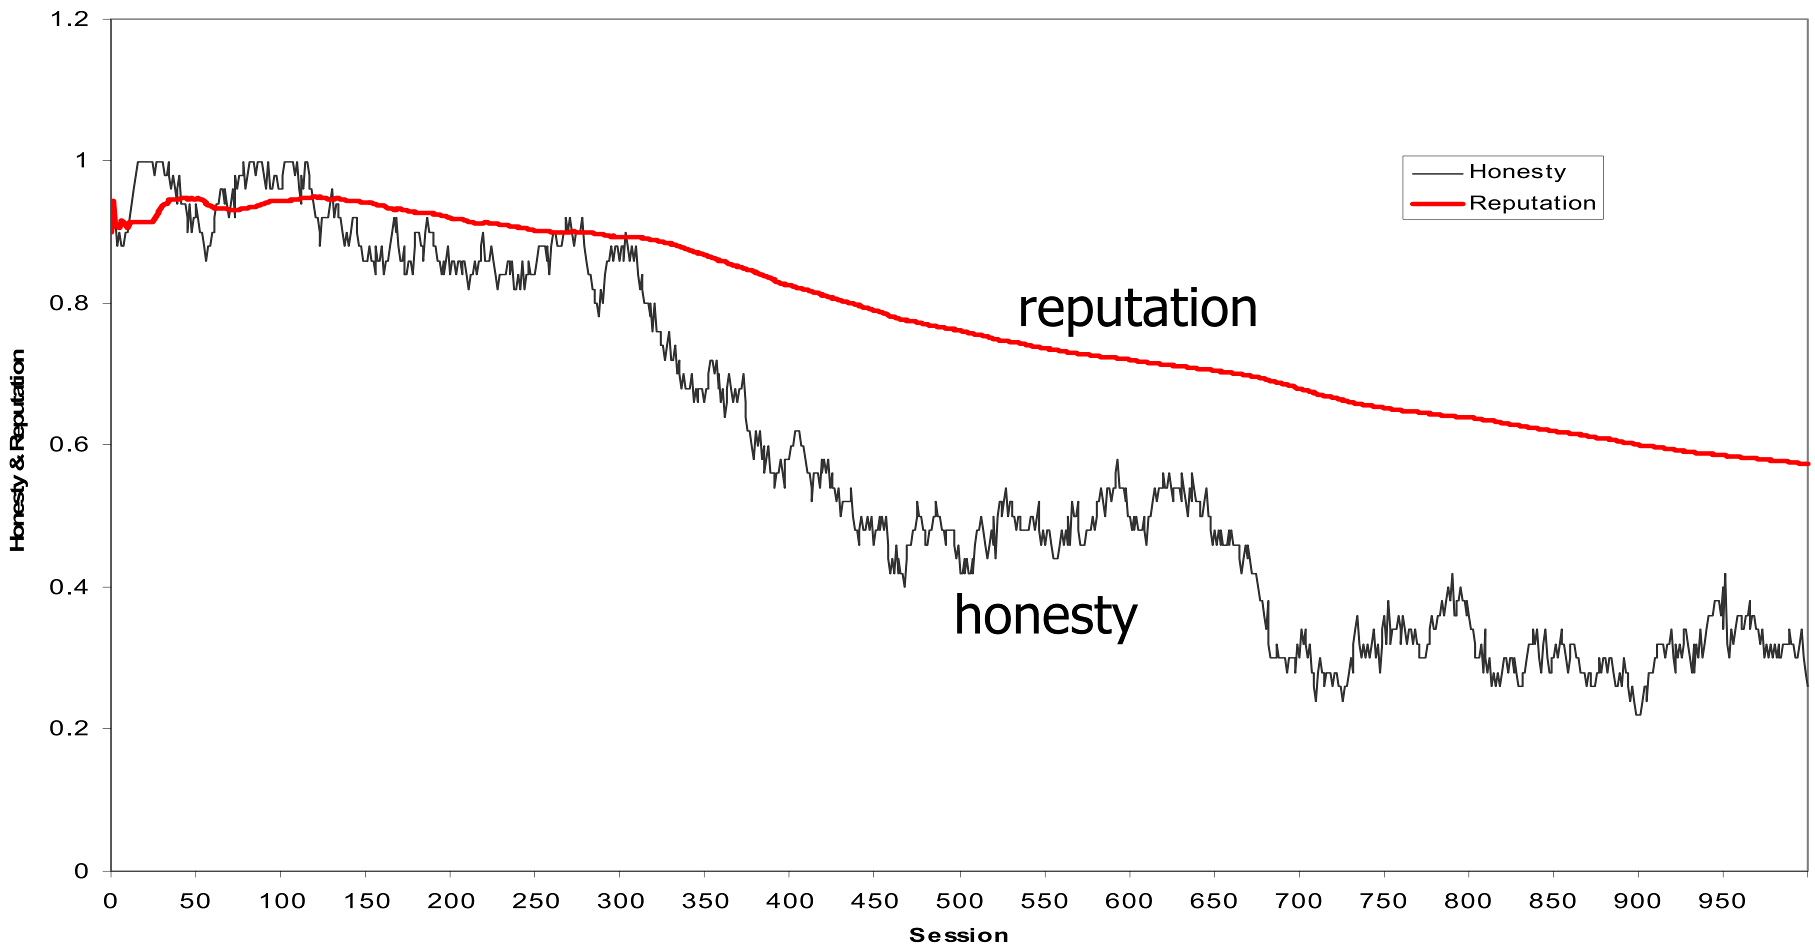

Seller honesty and reputation with λ=0.99

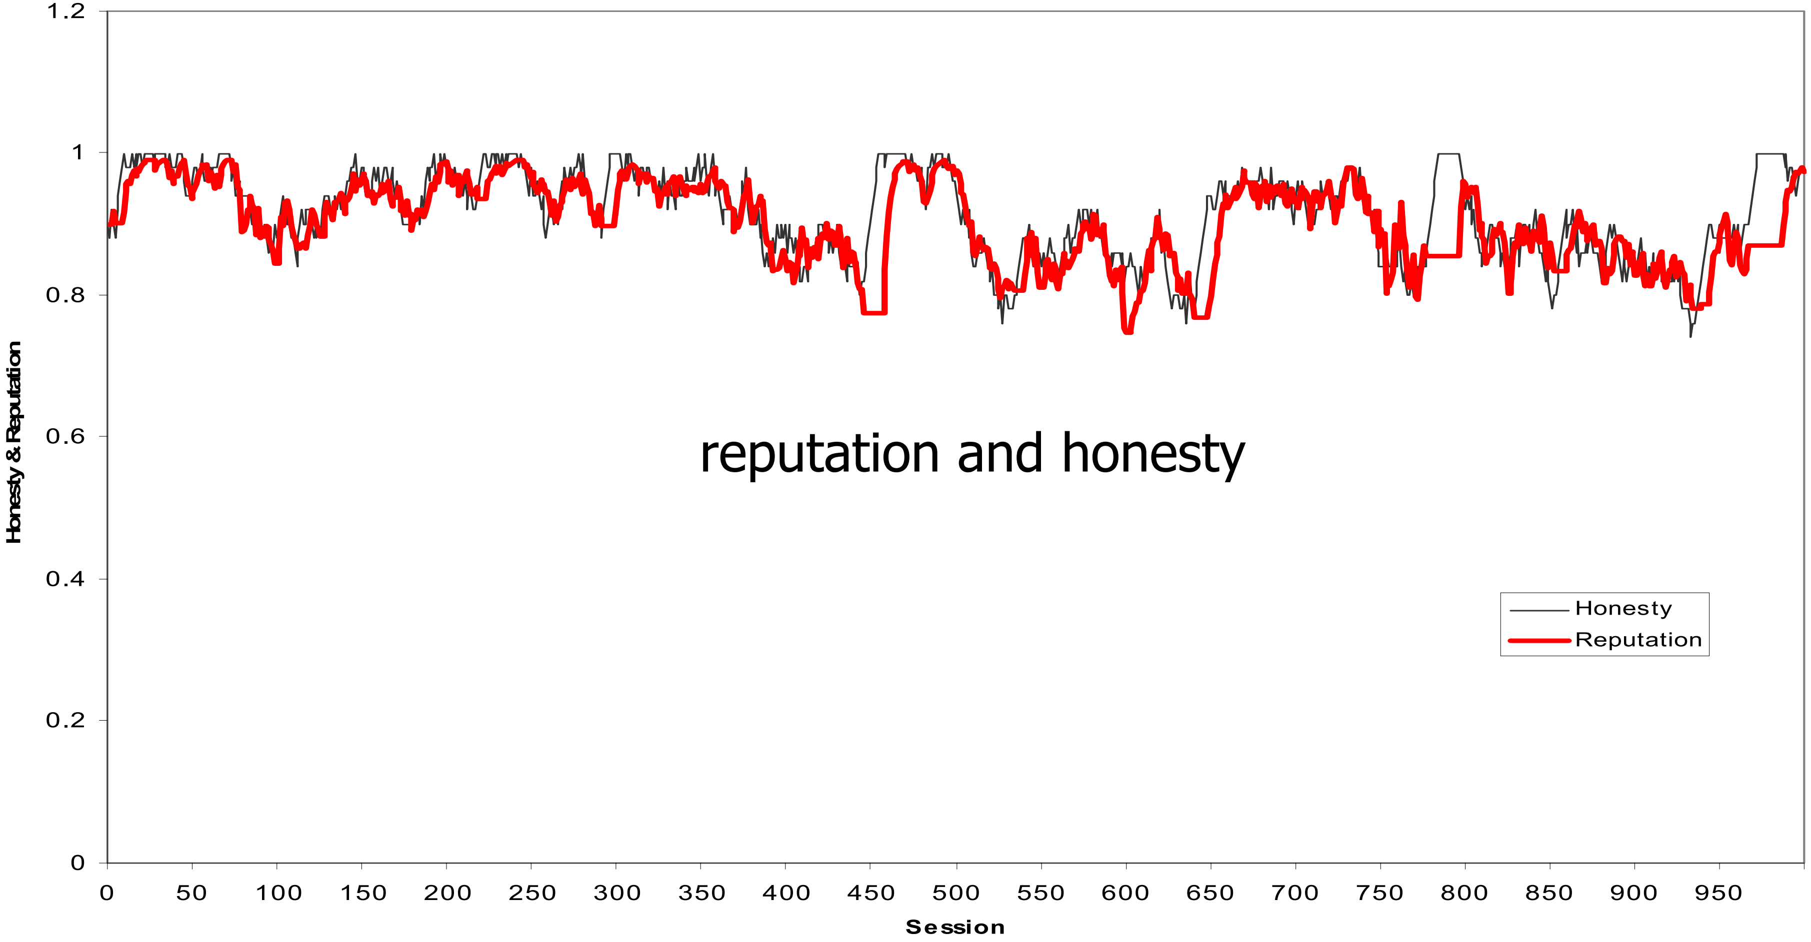

Observations from simulation

- A market without reputation system will degenerate
- A market with a reputation system that never forgets will degenerate
- A market with a reputation system that gradually forgets old behaviour can have stable quality.

Trust and Reputation Systems Part 2

Audun Josang
QUT, Australia
a.josang@qut.edu.au
http://www.fit.qut.edu.au/~josang/

Answer to question about convergence

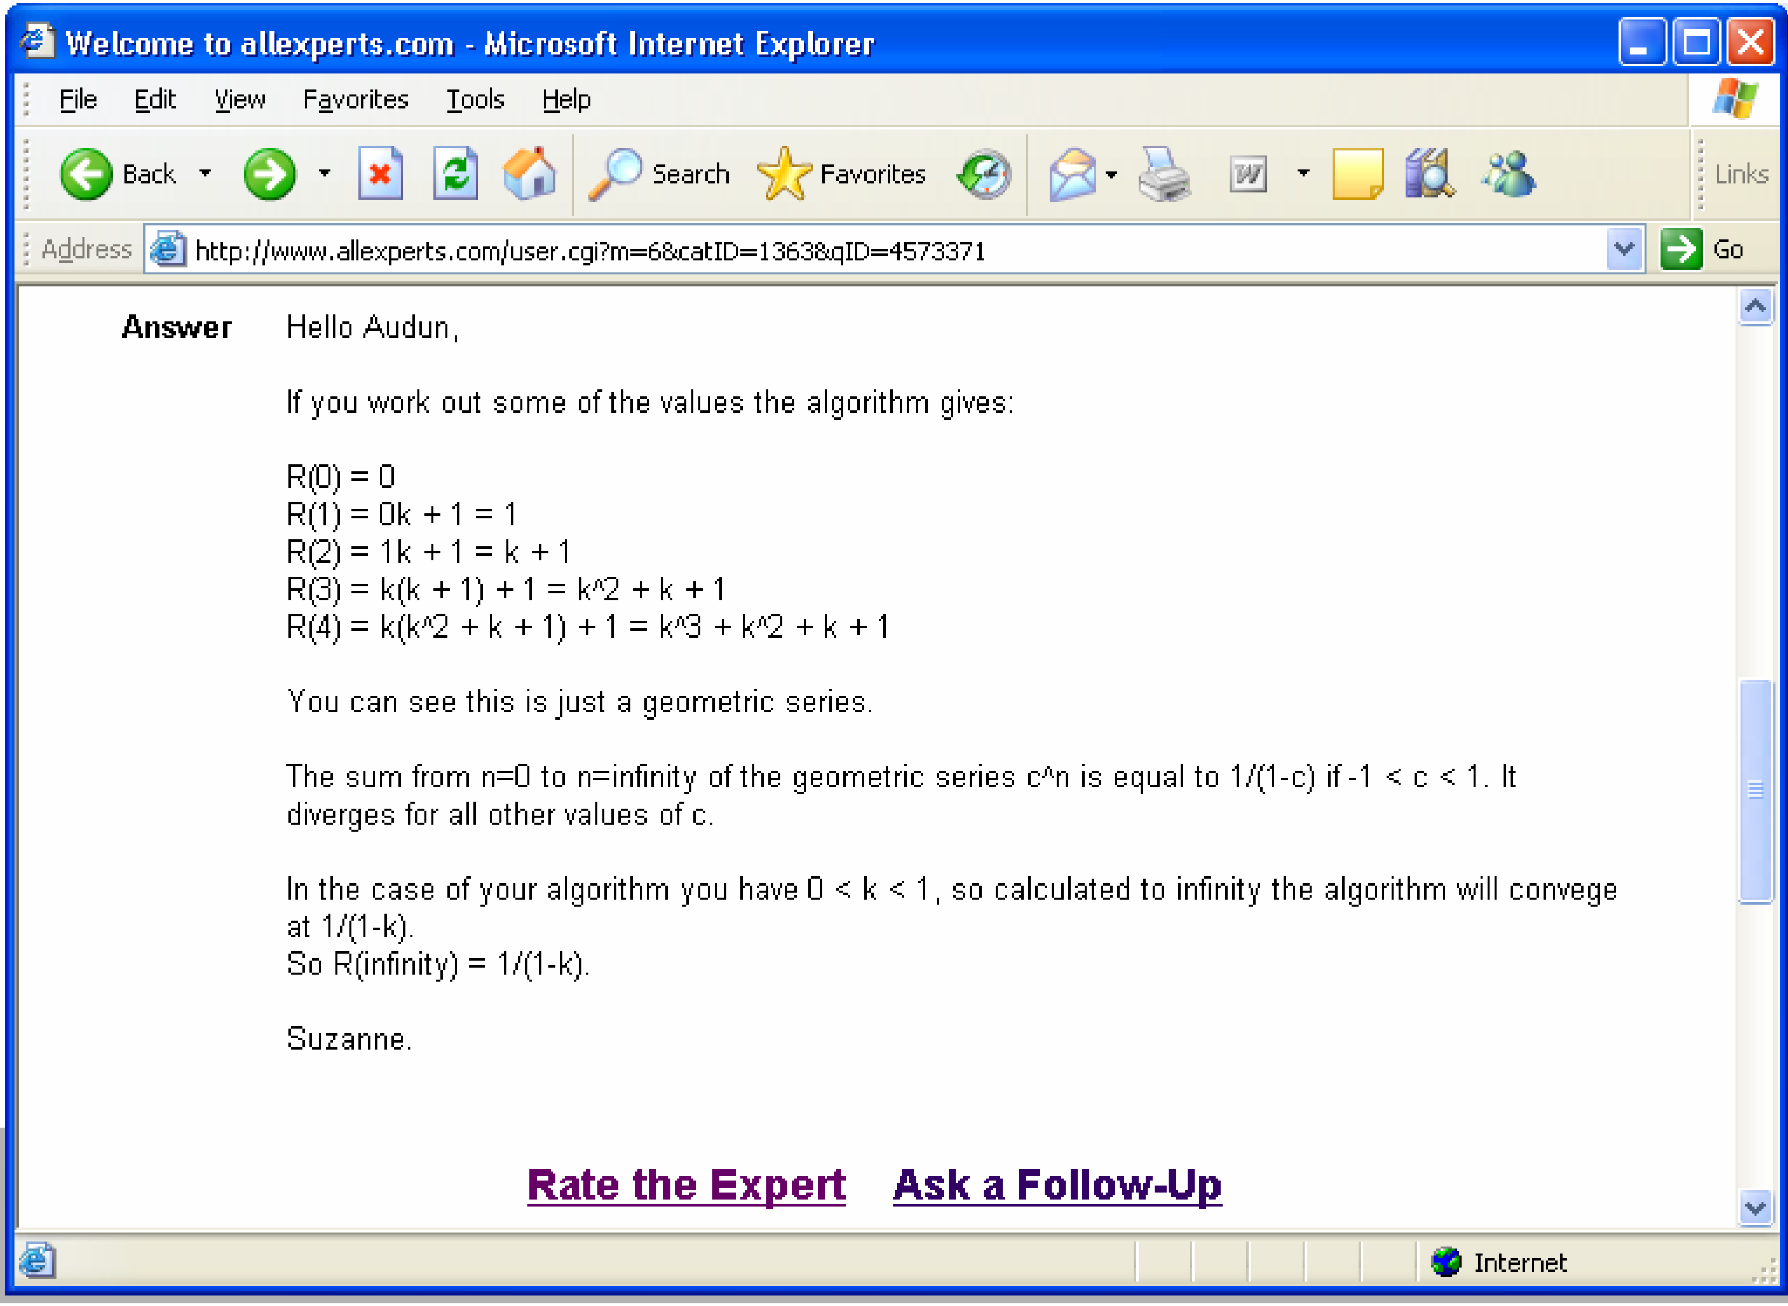

Answer Hello Audun,

If you work out some of the values the algorithm gives:

R(0) = 0
R(1) = 0k + 1 = 1
R(2) = 1k + 1 = k + 1
R(3) = k(k + 1) + 1 = k^2 + k + 1
R(4) = k(k^2 + k + 1) + 1 = k^3 + k^2 + k + 1

You can see this is just a geometric series.

The sum from n=0 to n=infinity of the geometric series c^n is equal to 1/(1-c) if -1 < c < 1. It diverges for all other values of c.

In the case of your algorithm you have 0 < k < 1, so calculated to infinity the algorithm will convege at 1/(1-k).
So R(infinity) = 1/(1-k).

Suzanne.

Rate the Expert **Ask a Follow-Up**

Convergence values

- For an infinite series of positive ratings $r = 1, s = 0$
  - $R_\infty = 1/(1-\lambda)$
  - $S_\infty = 0$
  - Score converges to $Sc(Z) = \dfrac{2-\lambda}{3-2\lambda}$ (with $r_{\text{base}} = s_{\text{base}} = 1$)
- For an infinite series of negative ratings $r = 0, s = 1$
  - $R_\infty = 0$
  - $S_\infty = 1/(1-\lambda)$
  - Score converges to $Sc(Z) = \dfrac{1-\lambda}{3-2\lambda}$ (with $r_{\text{base}} = s_{\text{base}} = 1$)

Rating Suzanne

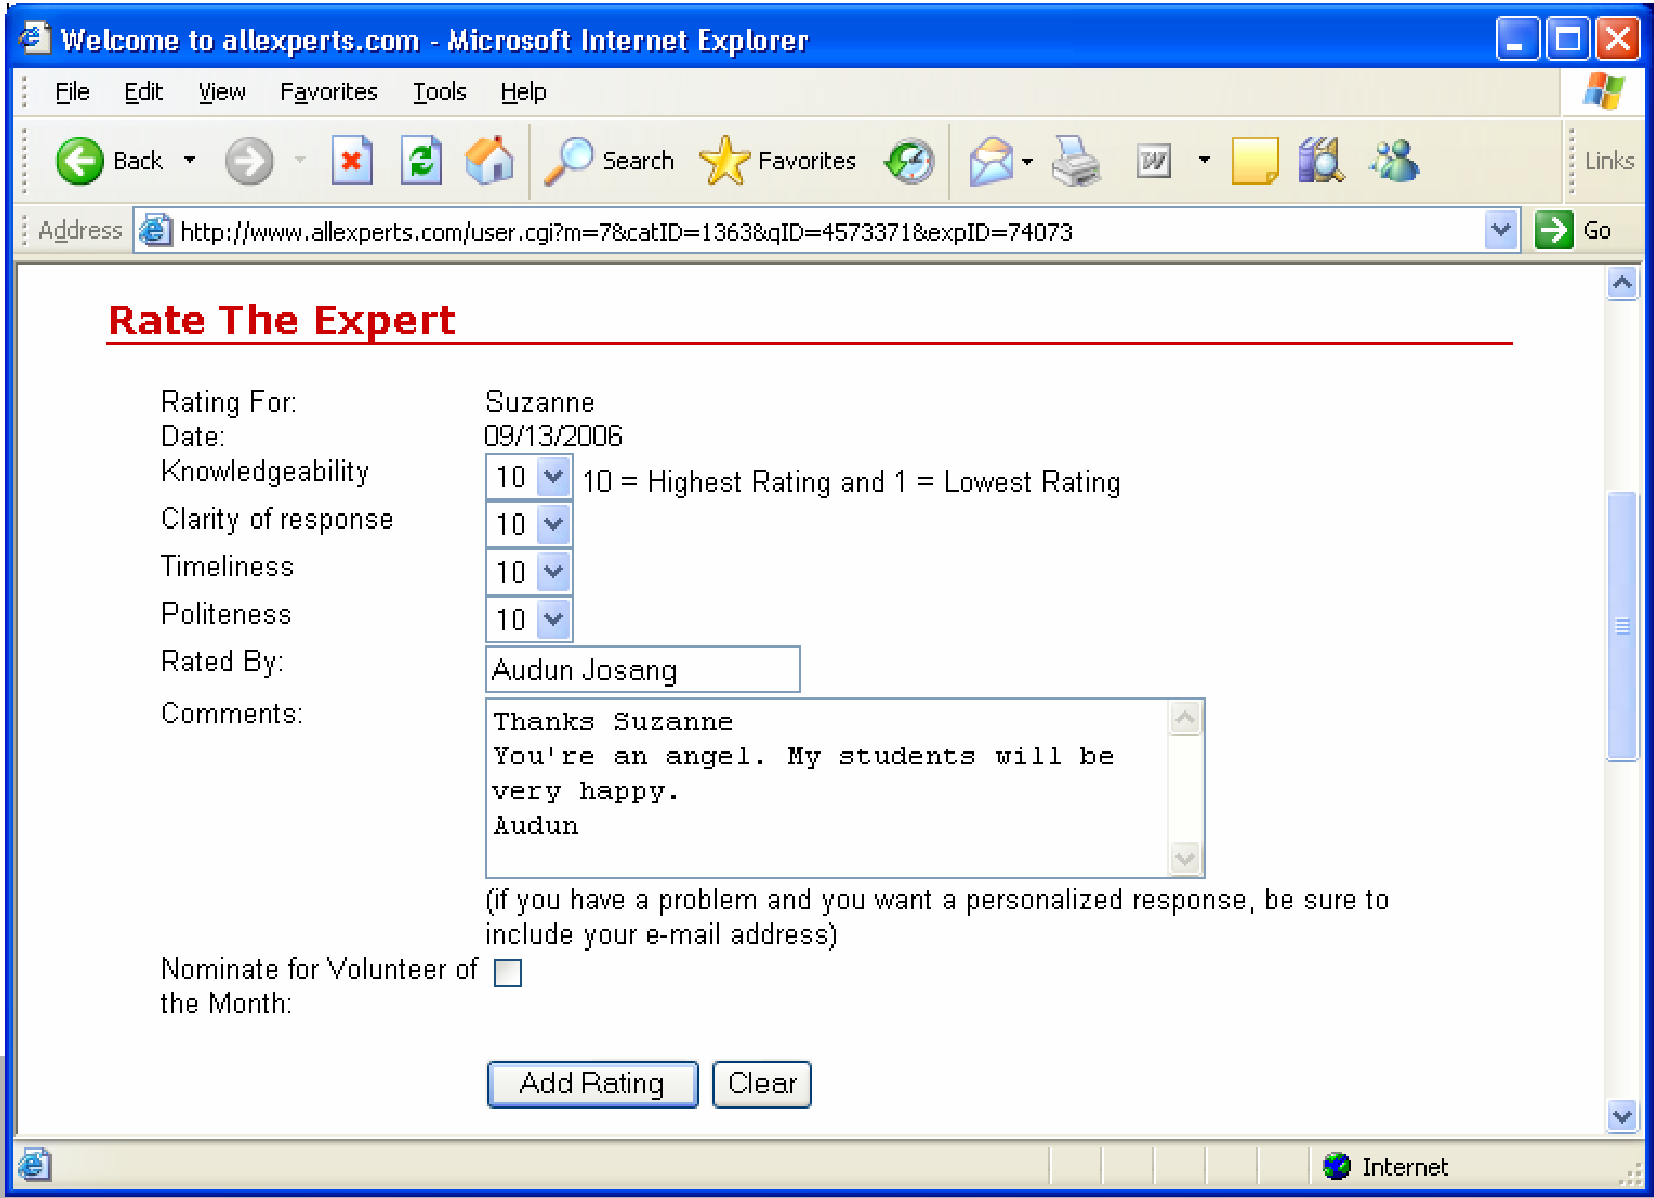

Suzanne’s reputation score

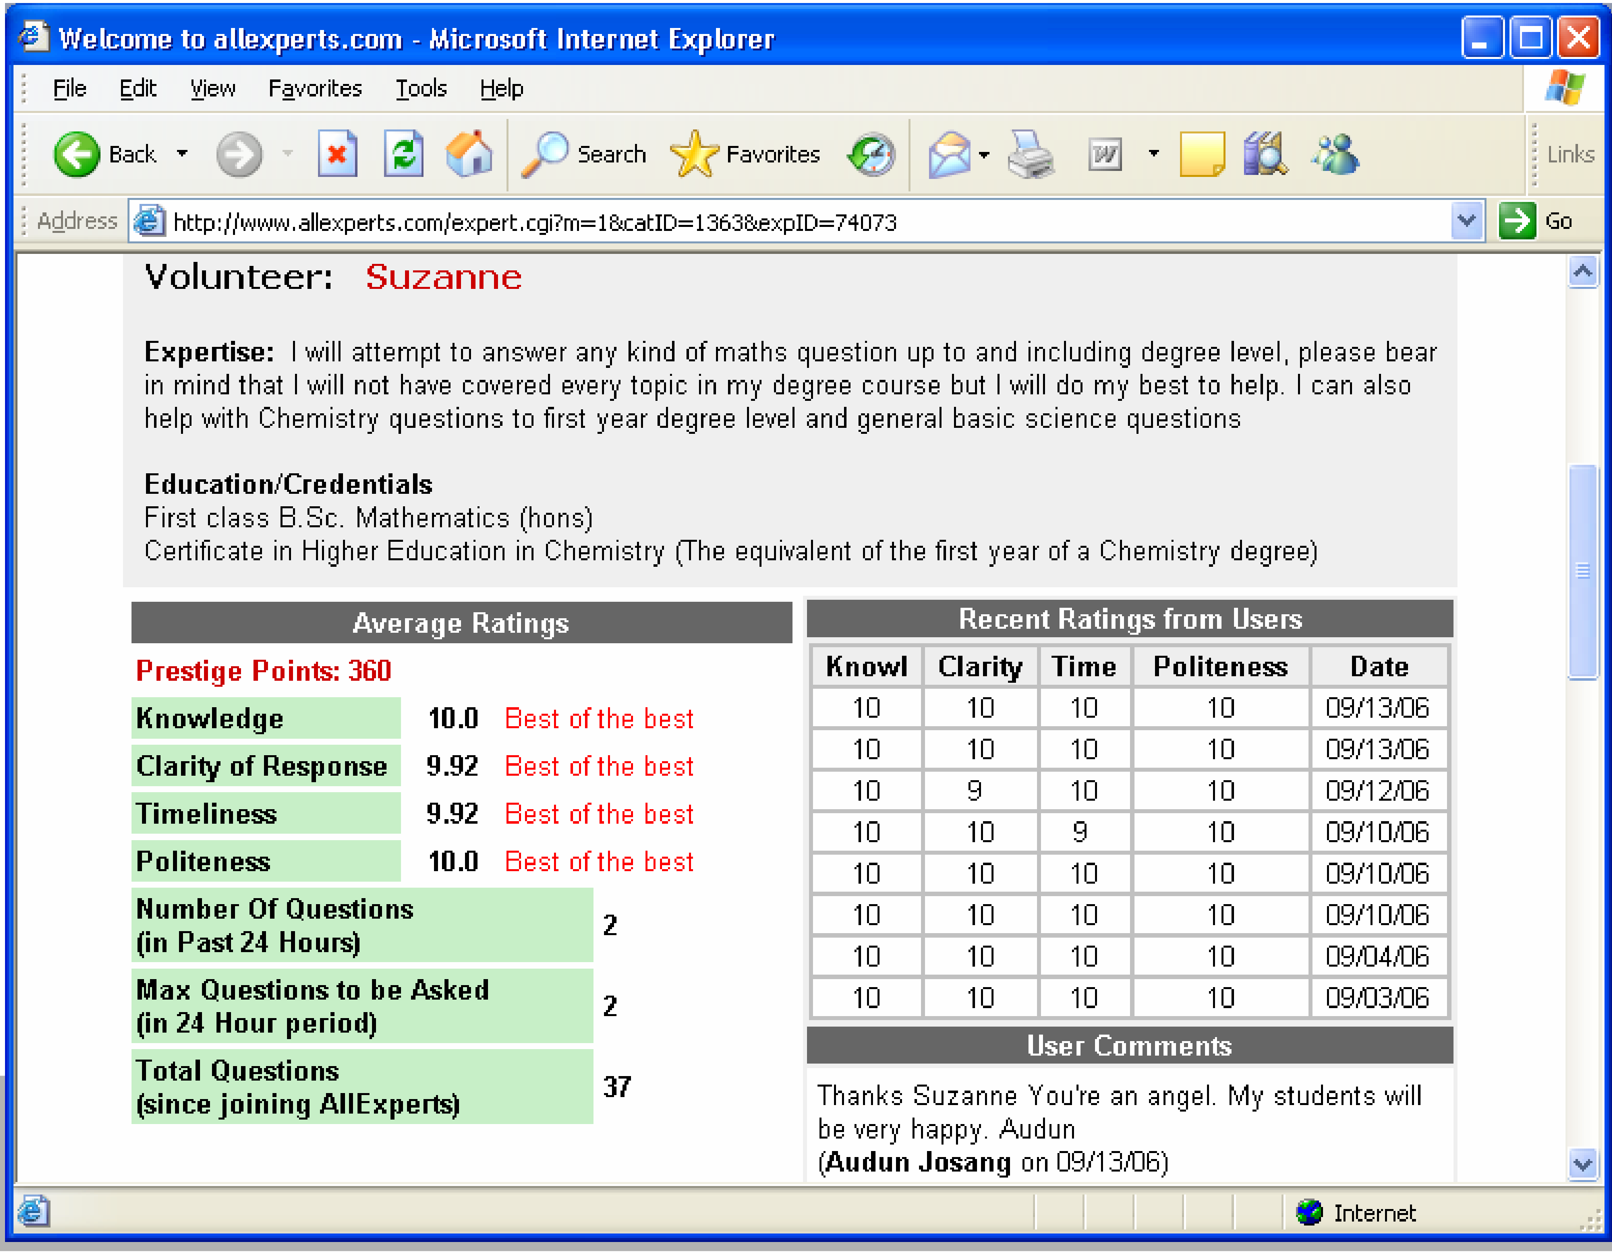

Welcome to allexperts.com - Microsoft Internet Explorer

http://www.allexperts.com/expert.cgi?m=1&catID=1363&expID=74073

Volunteer: Suzanne

Expertise: I will attempt to answer any kind of maths question up to and including degree level, please bear in mind that I will not have covered every topic in my degree course but I will do my best to help. I can also help with Chemistry questions to first year degree level and general basic science questions

Education/Credentials
First class B.Sc. Mathematics (hons)
Certificate in Higher Education in Chemistry (The equivalent of the first year of a Chemistry degree)

Average Ratings

Prestige Points: 360

Knowledge	10.0	Best of the best
Clarity of Response	9.92	Best of the best
Timeliness	9.92	Best of the best
Politeness	10.0	Best of the best
Number Of Questions (in Past 24 Hours)	2	
Max Questions to be Asked (in 24 Hour period)	2	
Total Questions (since joining AllExperts)	37	

Recent Ratings from Users

Knowl	Clarity	Time	Politeness	Date
10	10	10	10	09/13/06
10	10	10	10	09/13/06
10	9	10	10	09/12/06
10	10	9	10	09/10/06
10	10	10	10	09/10/06
10	10	10	10	09/10/06
10	10	10	10	09/04/06
10	10	10	10	09/03/06

User Comments

Thanks Suzanne You're an angel. My students will be very happy. Audun
(**Audun Josang** on 09/13/06)

Multinomial Bayesian reputation

- Problems with binomial reputation systems
  - Can only take binary ratings (positive, negative)
  - Can not represent polarised ratings
- Multinomial reputation systems
  - Can have any number of rating levels
  - Can represent polarised ratings

Multinomial reputation example

- Example from Microsoft
- Reflects polarised ratings

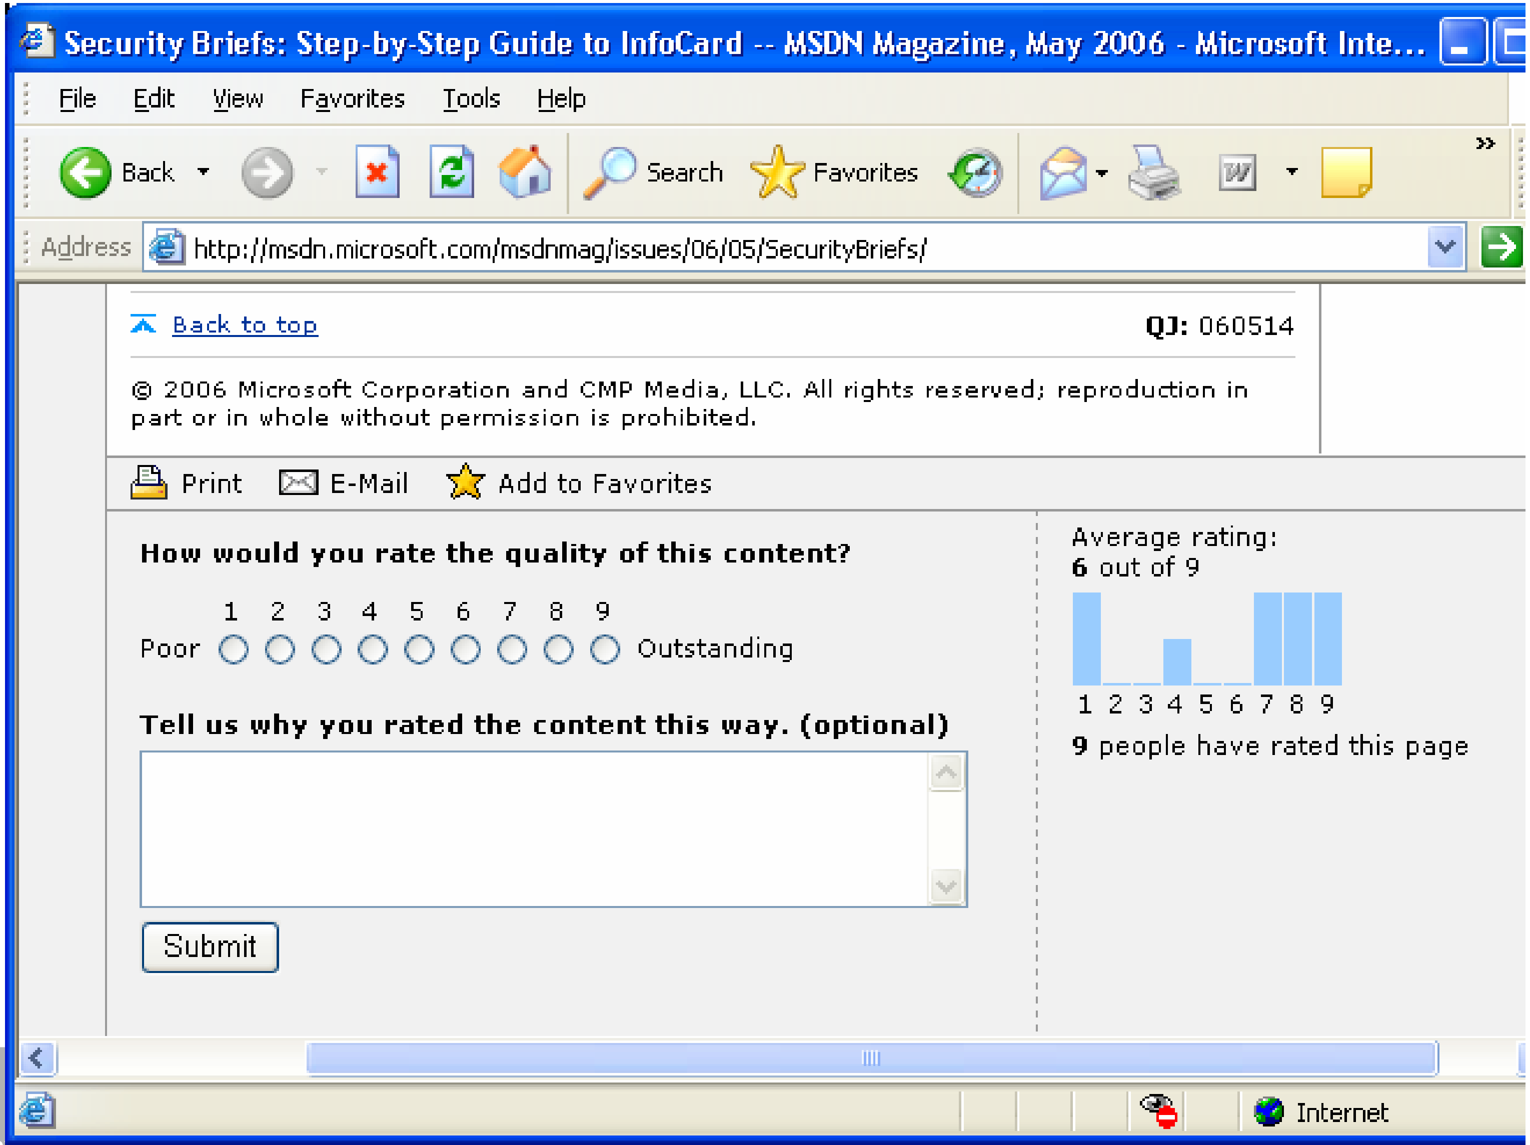

Multinomial Bayesian model

- Dirichlet PDF (probability density function)

$$f(\vec{p} \mid \vec{r}, \vec{a}) = \frac{\Gamma\left(\sum_{j=1}^{l} \left(r(x_j) + Ca(x_j)\right)\right)}{\prod_{j=1}^{l} \Gamma\left(r(x_j) + Ca(x_j)\right)} \prod_{j=1}^{l} p(x_j)^{r(x_j)+Ca(x_j)-1}$$

$\vec{p}$: multinomial probability vector

$\vec{r}$: multinomial evidence vector

$\vec{a}$: multinomial base rate vector

x_j : event

$C = 2$

Example: ternary state space

Example:

Urn with balls of 3 different colours

- t1 = θ1 = Red
- t2 = θ2 = Yellow
- t3 = θ3 = Black

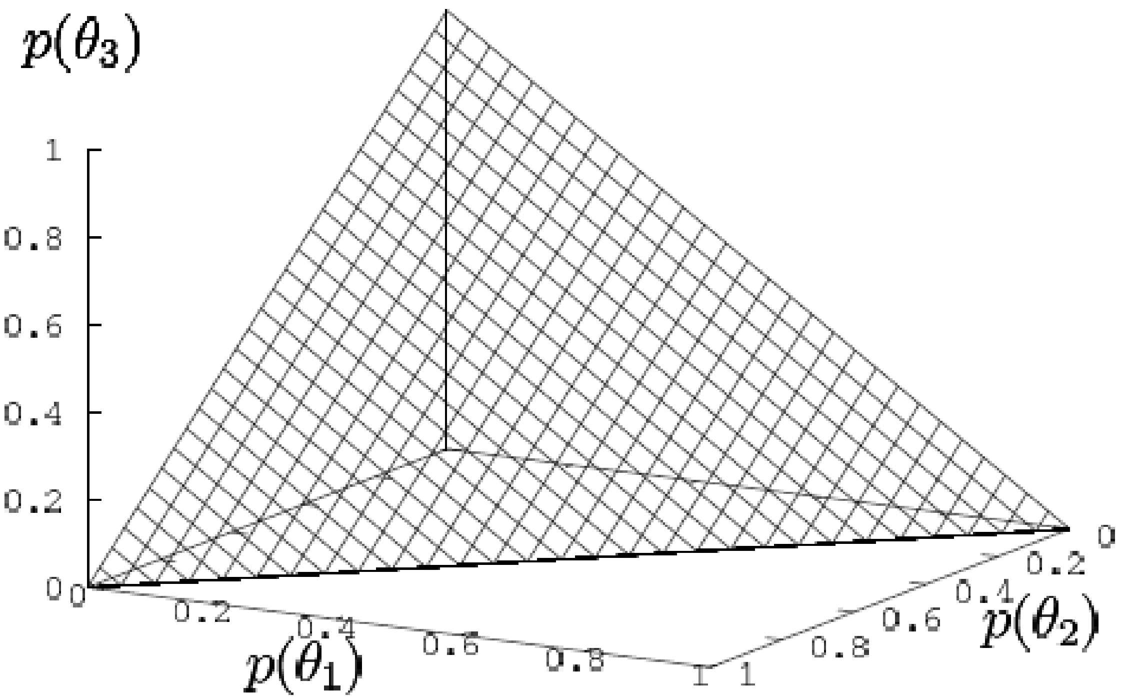

- Additivity requires: $p(t_1) + p(t_2) + p(t_3) = 1$

Prior ternary Dirichlet PDF

Example:

Urn with balls of 3 different colours. Ternary *a priori* probability density.

– t1: Red
– t2: Yellow
– t3: Black

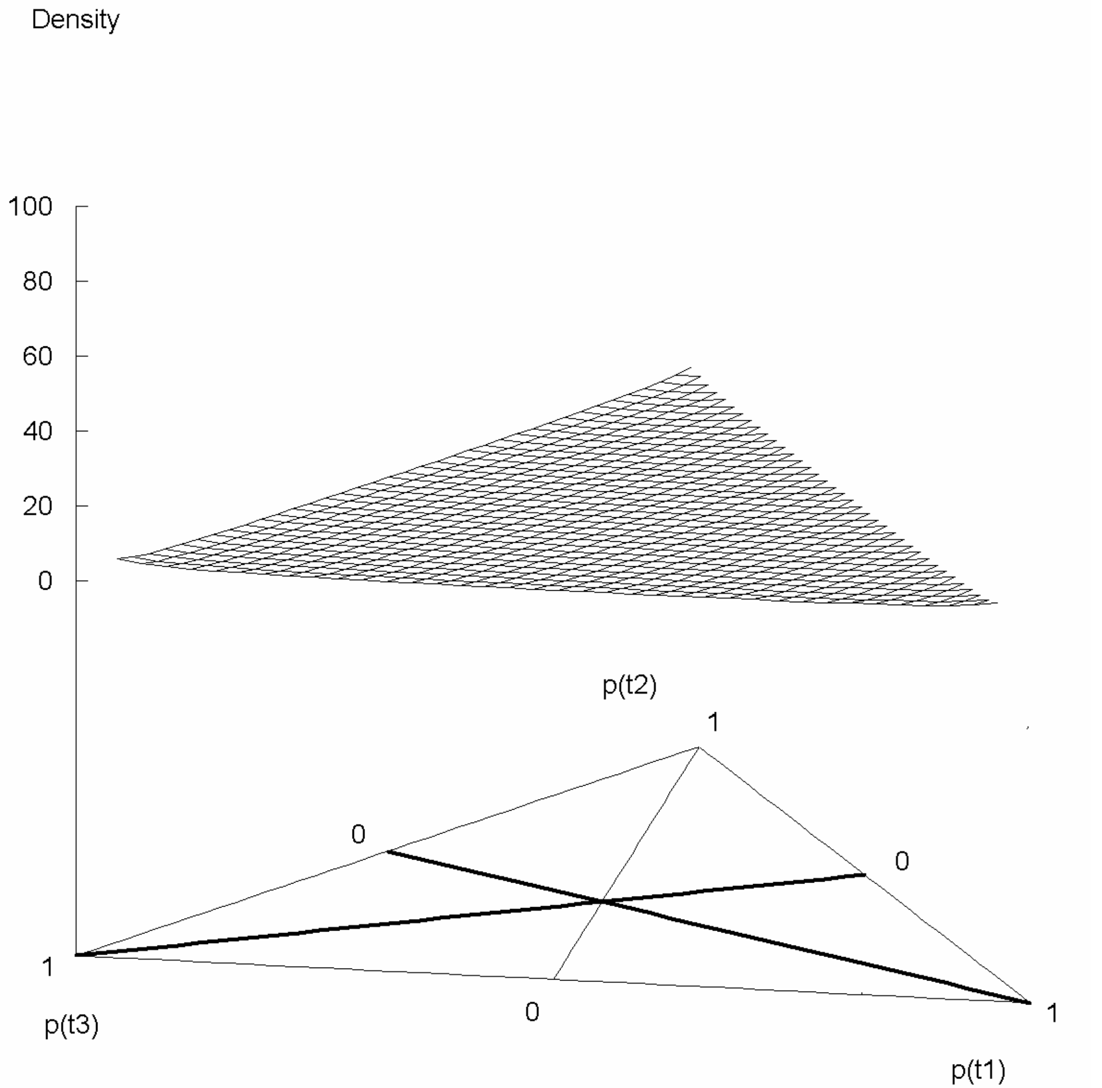

Example posterior ternary Dirichlet PDF

A posteriori probability density after picking:

- 6 red balls (t1)
- 1 yellow ball (t2)
- 1 black ball (t3)

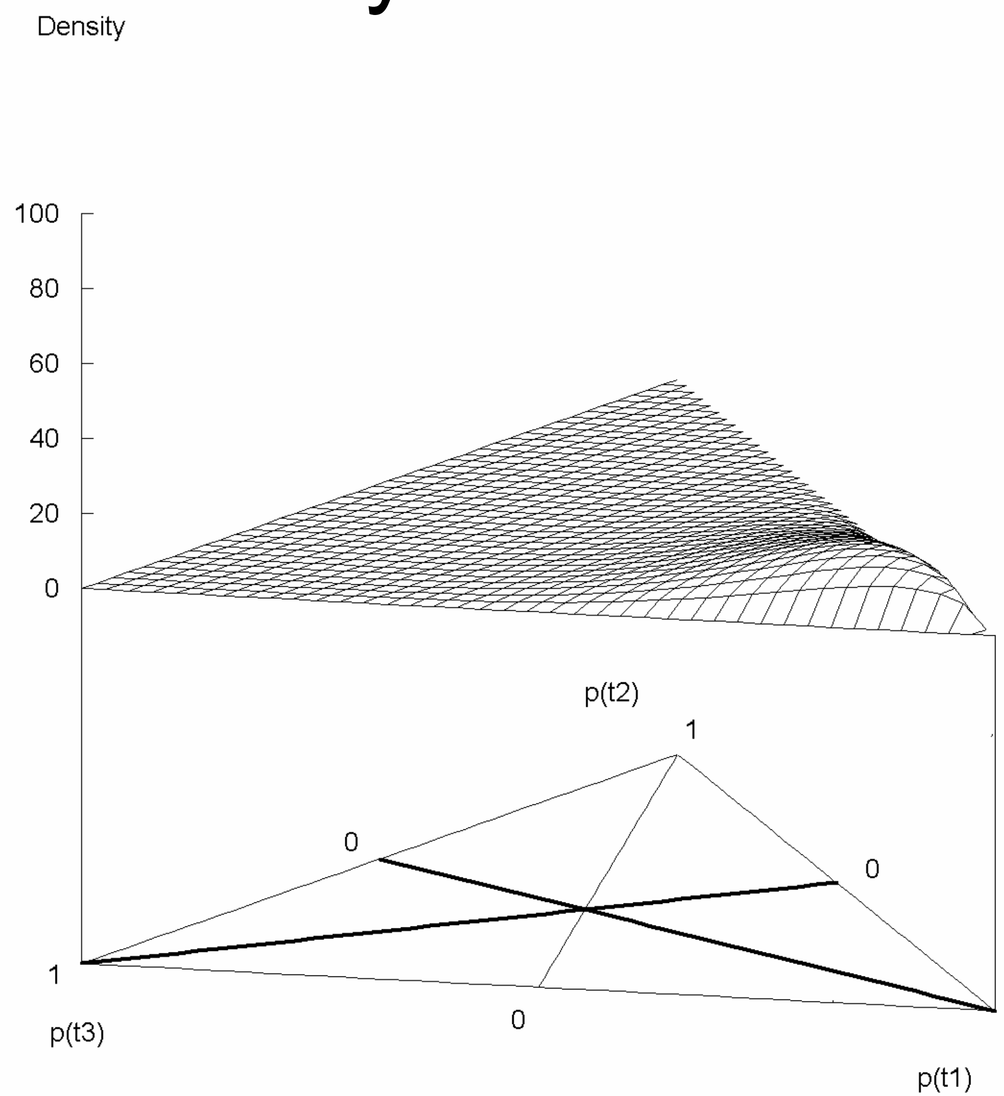

Example posterior ternary Dirichlet PDF

A posteriori probability density after picking:

- 20 red balls (t1)
- 20 yellow balls (t2)
- 20 black balls (t3)

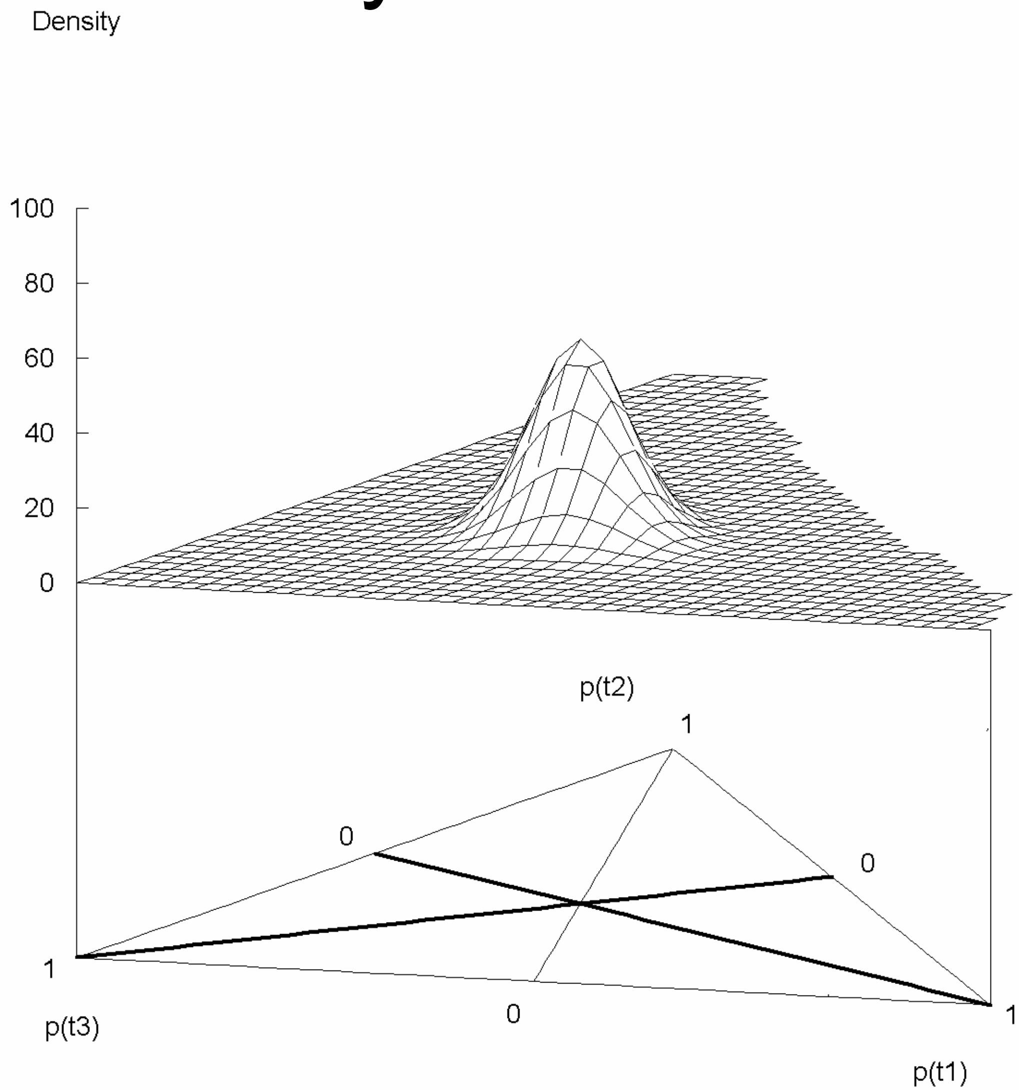

Example posterior ternary Dirichlet PDF

A posteriori probability density after picking:

- 20 red balls (t1)
- 20 yellow balls (t2)
- 50 black balls (t3)

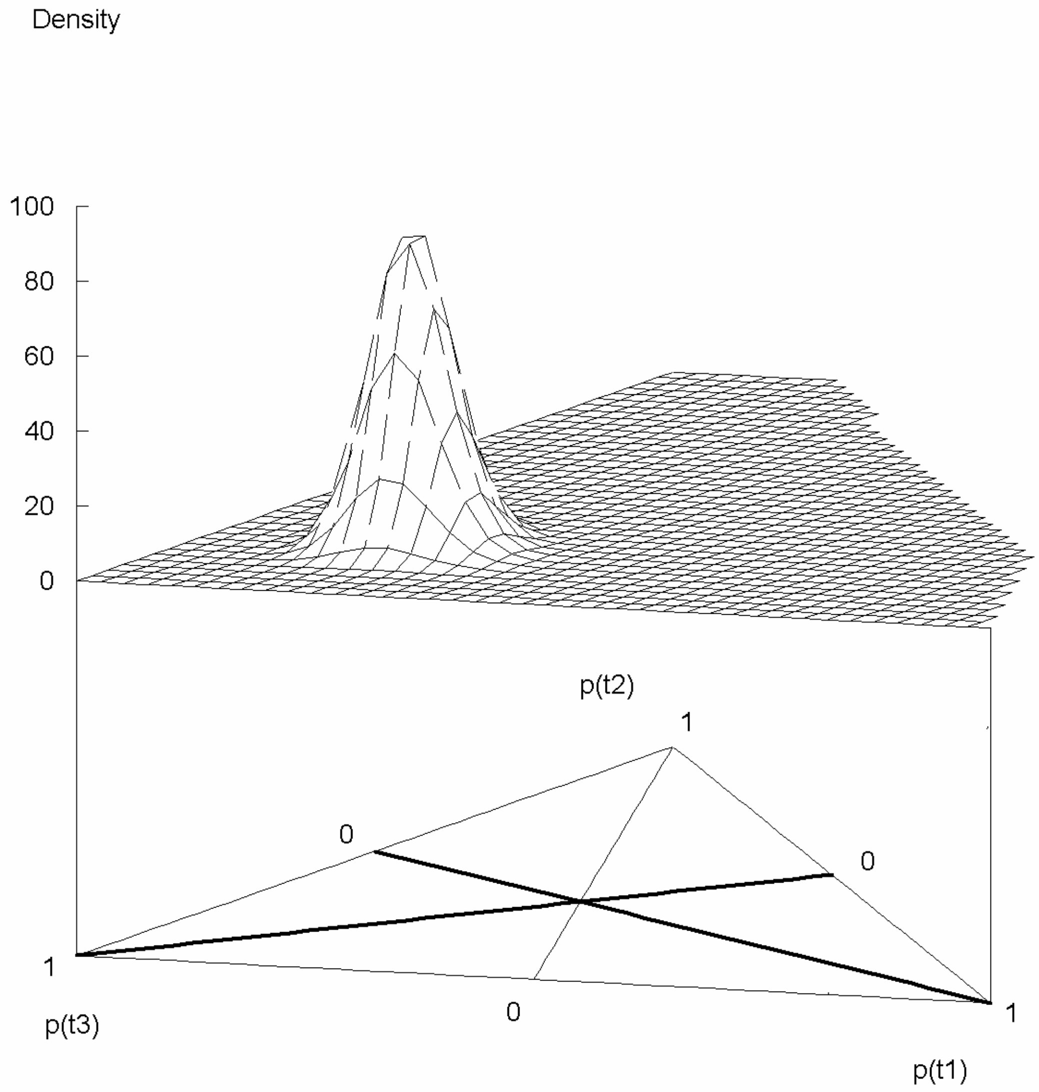

Score of ordered set of outcomes

- Density functions do not naturally represent the reputation score on an ordered set of rating levels
  - e.g. Set of {Mediocre, Bad, Average, Good, Excellent}
- Rating levels can be represented as a set of different outcomes
- Probability expectation of each rating level can be represented separately

Multinomial reputation score

- The multinomial reputation score can be defined equal to the Dirichlet-PDF probability expectation

$$\mathrm{E}(p(L_j) \mid \vec{r}, \vec{a}) = \frac{r(L_j) + C \cdot a(L_j)}{C + \sum_{j=1}^{l} r(L_j)}$$ Proba. expect.

$\vec{r}$: Multinomial evidence vector

$\vec{a}$: Multinomial base rate vector

$C = 2$

l : Number of rating levels

L_j : particular rating level

$Sc(L_j) = \mathrm{E}(p(L_j) \mid \vec{r}, \vec{a})$: Multinomial reputation score

Initial reputation score

Example with $l = 5$ discrete rating levels:

1) mediocre, 2) bad, 3) average, 4) good, 5) excellent

Initial uniform reputation score before any ratings have been received.

Base rate $a(x_i) = 0.2$

Can represent polarised ratings!

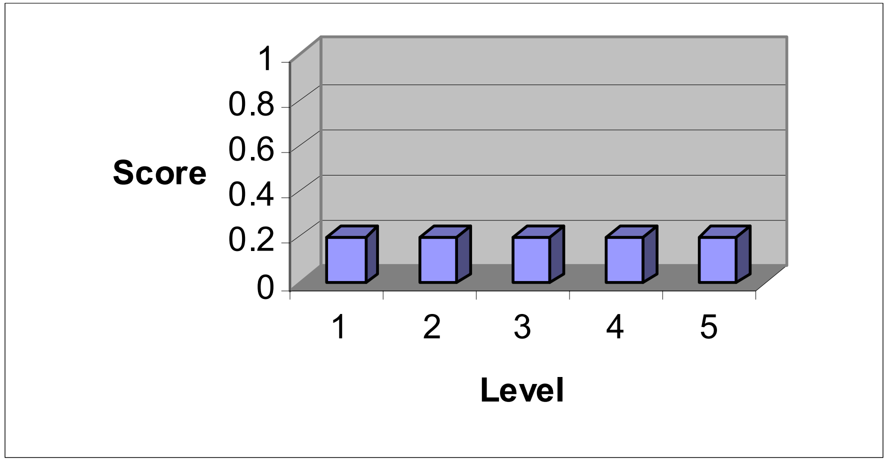

Reputation score of polarise ratings

As before, 5 discrete levels:

1) very bad, 2) bad, 3) average, 4) good, 5) very good

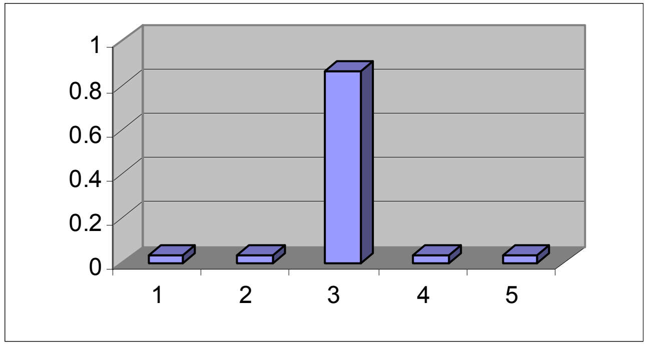

Non-polarised reputation score after 10 average ratings

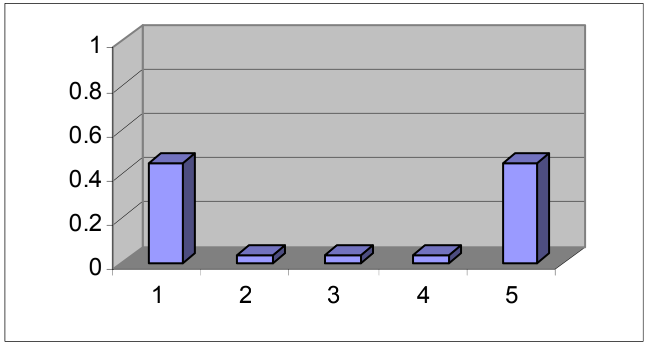

Polarised reputation score after 5 very bad and 5 very good ratings

Computing multinomial reputation over time with fixed base rate

- $\vec{R}_i$: accumulated evidence at time i
- $\vec{r}$: evidence collected during 1 time period.
- λ : longevity factor
- $\vec{R}_{i+1} = \lambda \cdot \vec{R}_i + \vec{r}$: Recursive updating algorithm
- $Sc_i(L_j) = \mathrm{E}_i(p(L_j) \mid \vec{R}_i, \vec{a})$: Score at time period i

Score evolution over time with fixed base rate

Five discrete rating levels:

1. Mediocre
2. Bad,
3. Average,
4. Good,
5. Excellent

Longevity $\lambda = 0.9$

Base rate $a(x) = 0.2$

Periods 1-5: Mediocre

Periods 6-10: Excellent

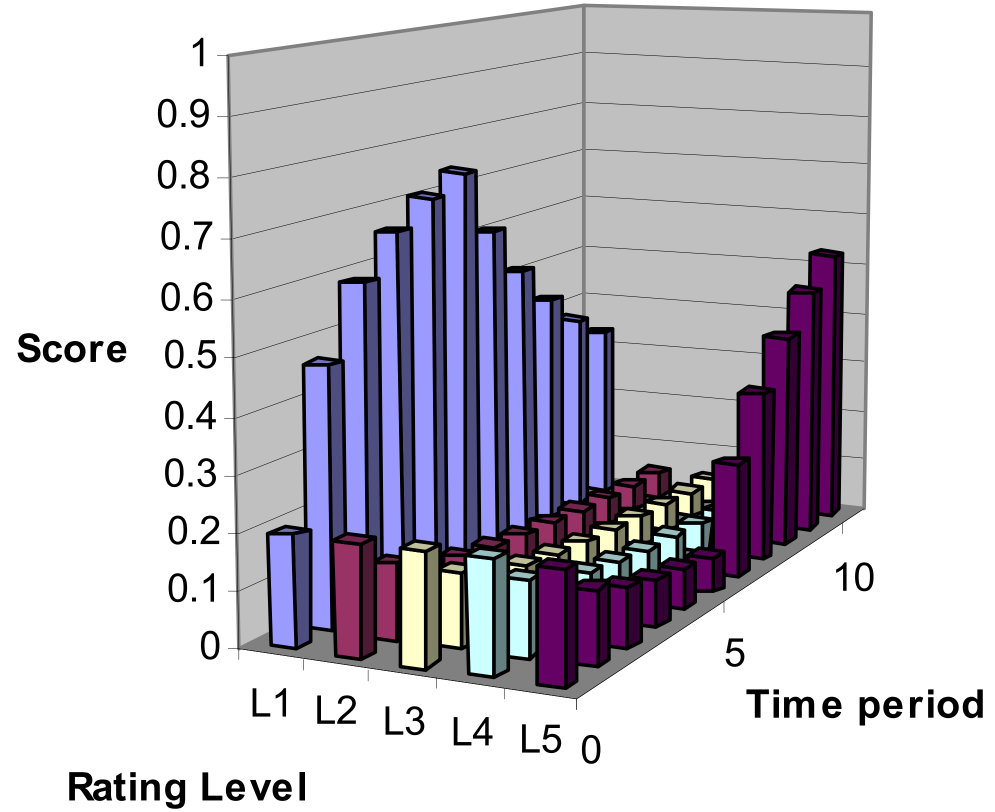

Score evolution over time with fixed base rate

Longevity λ= 0.9

Base rate $a(x)$ = 0.2

Periods 1-10: Mediocre

Periods 11-50: Excellent

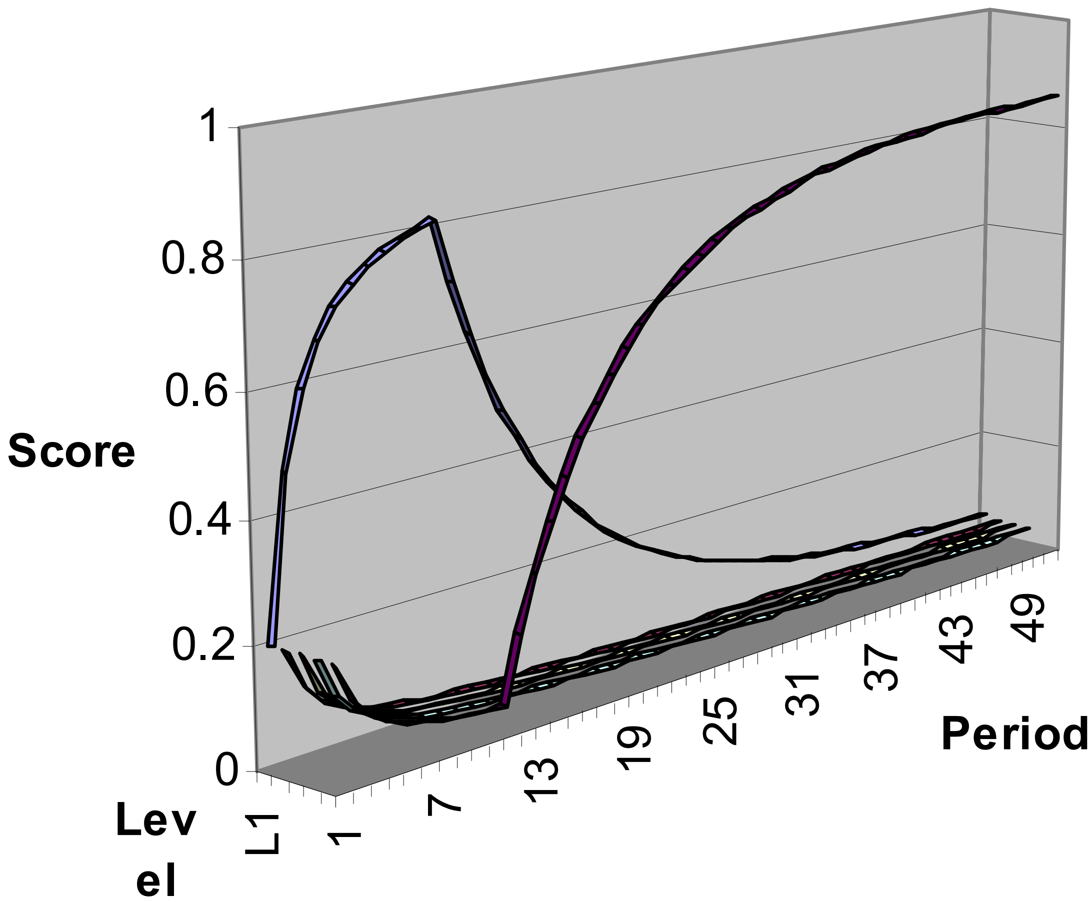

The max and min reputation score is determined by the longevity factor λ

Dynamic base rate as function of average reputation score

- New members should get a base rate equal to the average reputation score of the community
- Same for existing members
- Let M denote the whole community
- $\vec{a}_{i+1} = \vec{\mathrm{E}}_i(M)$: Dynamic base rate at time period $i+1$
- Dynamic base rate is thus updated each period
- $Sc_i(L_j) = \mathrm{E}_i(p(L_j) \mid \vec{R}_i, \vec{a}_i)$: Score with dynamic base rate
- Max and min reputation score become independent of the longevity factor λ

Score evolution over time with dynamic base rate

Longevity λ= 0.9

Base rate $a_{i+1}(Lj) = E_i(Lj)$

Periods 1-10: Mediocre

Periods 11-50: Excellent

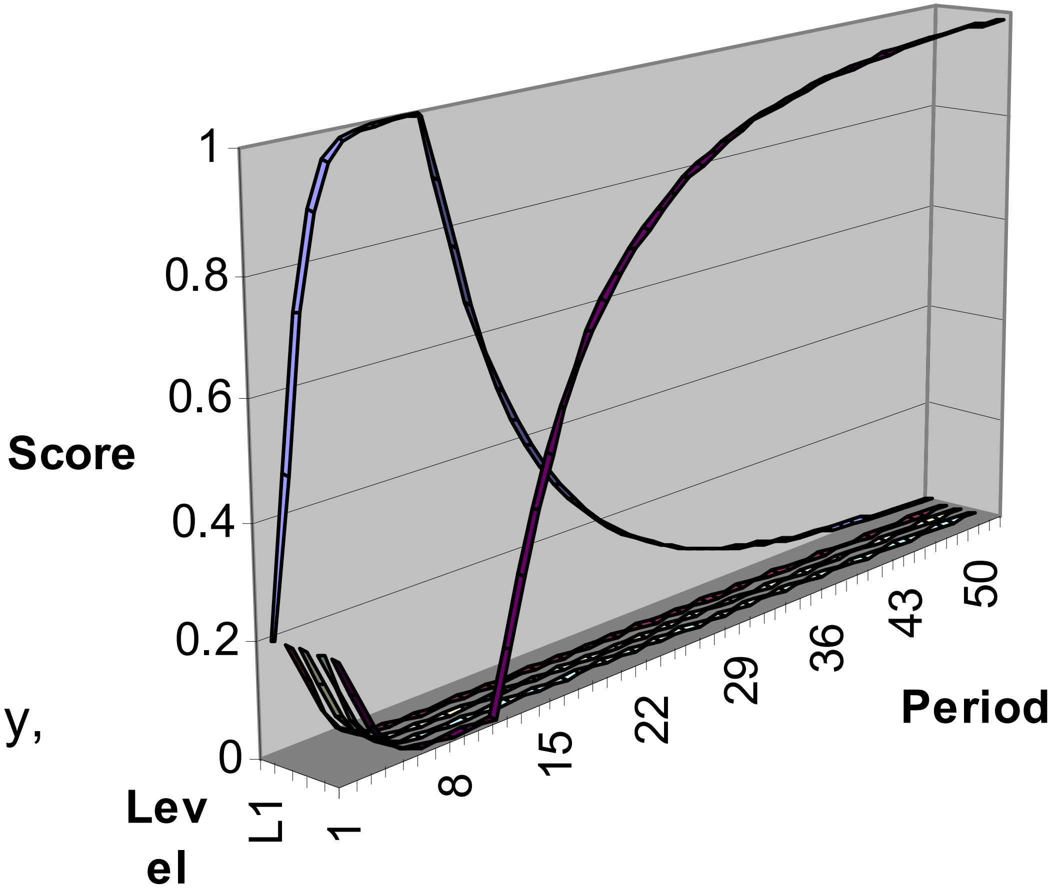

The max and min reputation scores are 0 and 1 respectively, and are independent of the longevity factor λ.

Point Estimate Reputation Score

- Sometimes useful to have a single-valued score
- Translate multinomial score to point-estimate score
- l : number of different rating levels
- j : particular rating level

- $v(L_j) = \frac{j\text{-}1}{l\text{-}1}$: Point value for each rating level

- $\sigma = \sum_{j=1}^{l} v(L_j) \cdot Sc(L_j)$: Point estimate

Multinomial score and point estimate with dynamic base rate

- Level values:
  - $v(L_1)=0$
  - $v(L_2)=0.25$
  - $v(L_3)=0.5$
  - $v(L_4)=0.75$
  - $v(L_5)=1$

- s=σ= point estimate
- Longevity λ= 0.9
- Base rate $a_{i+1}(L_j)= \mathrm{E}_i(L_j)$
- Periods 1-10: Mediocre
- Periods 11-50: Excellent

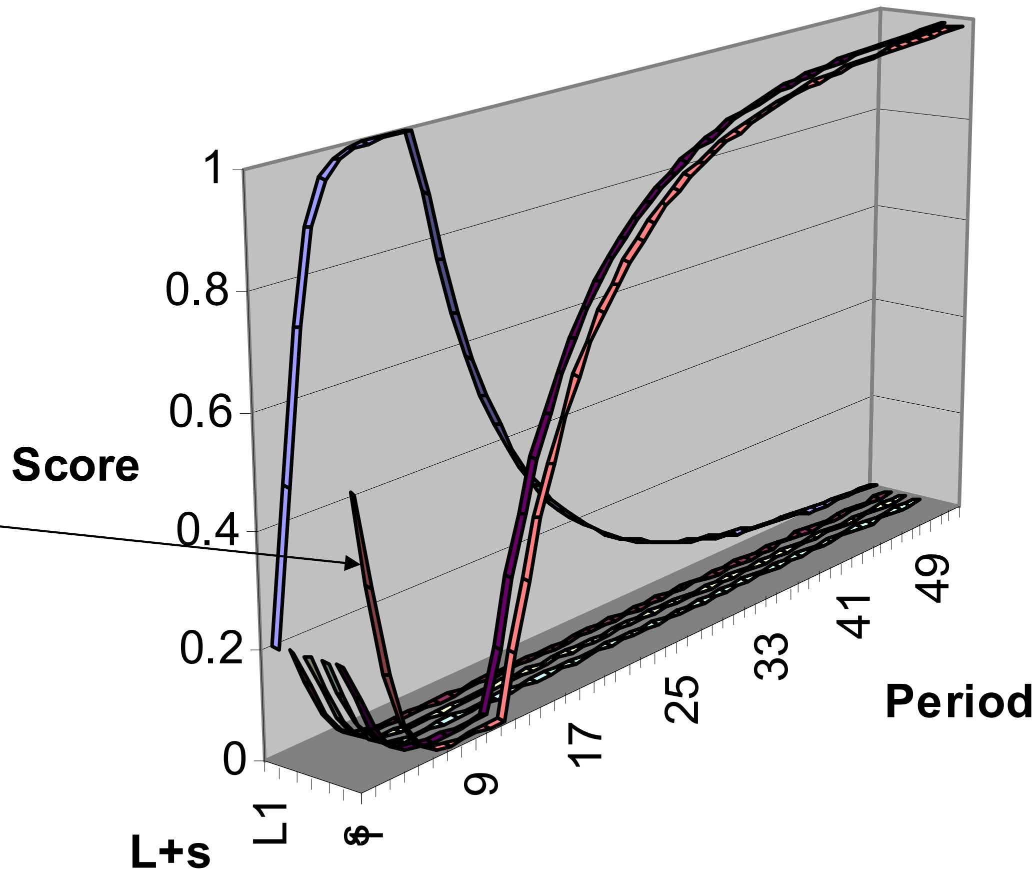

Score and point estimate with 5 consecutive uniform rating periods

Longevity λ= 0.9

Base rate $a_{i+1}(Lj)= E_i(Lj)$

Periods 1-10: Mediocre

Periods 11-20: Bad

Periods 21-30: Medium

Periods 31-40: Good

Periods 41-50: Excellent

•s=σ= point estimate

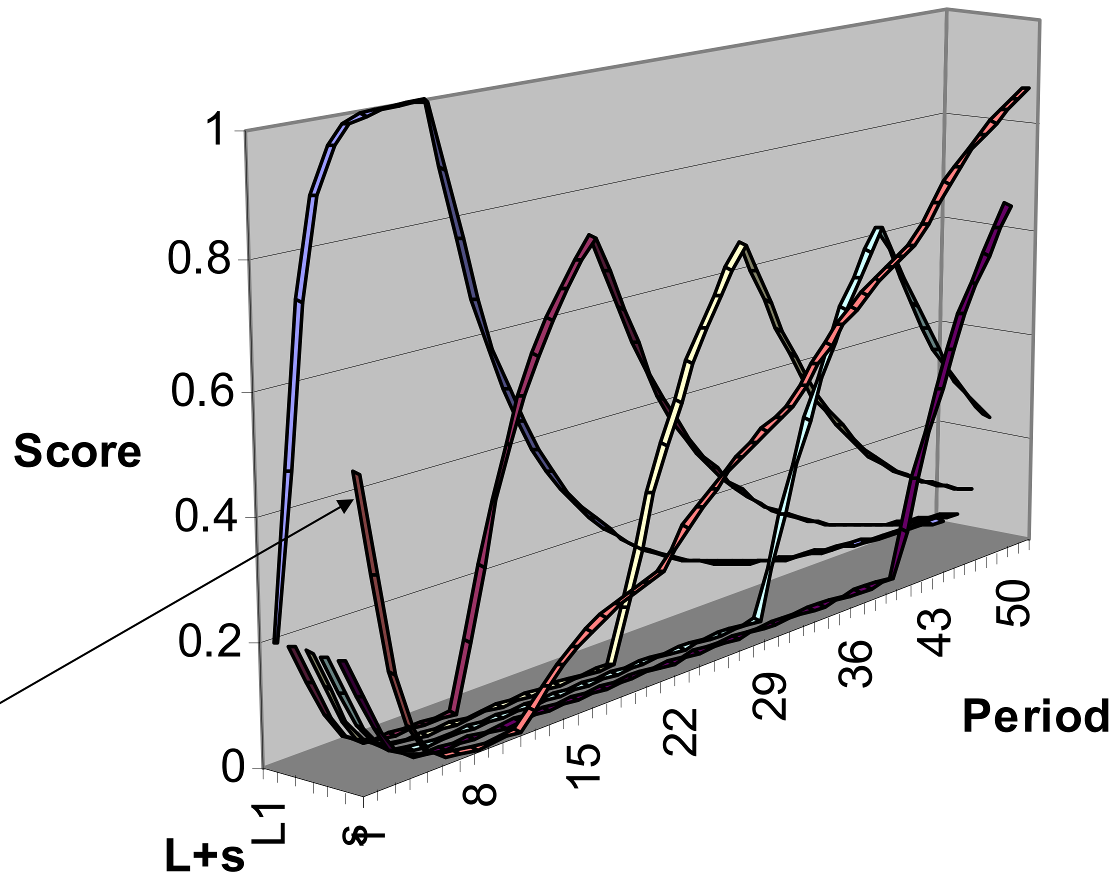

Discrete models

- Discrete measures
  - *"Very trustworthy", "trustworthy", "untrustworthy"*
- Computation
  - Heuristic formula, or lookup tables
- Simple to understand
- Qualitative
- Theoretically misguided

Belief models

- Assumes a trust scope σ
- Two semantic variants of each trust scope
  - Fuctional: Trust x for scope σ
    (e.g. "to be a good mechanic")
  - Referral: Trust x to refer or recommend someone/thing for scope σ
    (e.g. "to be a good at recommending mechanics)
- Two topological types
  - Direct: Trust as a result of direct experience
  - Indirect: Trust as a result of second hand evidence

Computing Trust with Subjective Logic

- Generalization of binary logic and probability calculus.
- Trust represented as binomial opinion: $\omega_x^A = (b, d, u, a)$
  - b: belief
  - d: disbelief
  - u: uncertainty
  - a: base rate

  in range [0,1]
- Where: $b + d + u = 1$
- Expectation value: $\mathrm{E}(\omega) = b + au$
- Explicit belief ownership.

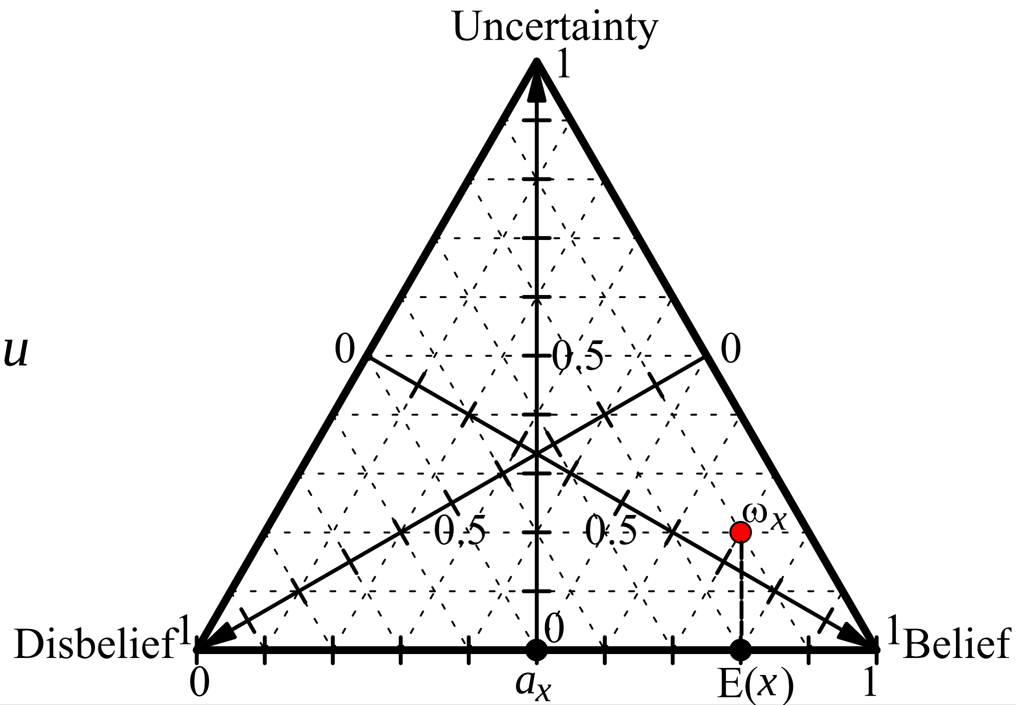

Subjective logic operators 1

Opinion operator name	Opinion operator symbol	Logic operator symbol	Logic operator name
Addition	$+$	$\cup$	UNION
Subtraction	$-$	$\backslash$	DIFFERENCE
Complement	$\neg$	$\overline{x}$	NOT
Expectation	E(x)	n.a.	n.a.
Multiplication	$\cdot$	$\wedge$	AND
Division	$/$	$\overline{\wedge}$	UN-AND
Comultiplication	$\sqcup$	$\vee$	OR
Codivision	$\overline{\sqcup}$	$\overline{\vee}$	UN-OR

Subjective logic operators 2

Opinion operator name	Opinion operator symbol	Logic operator symbol	Logic operator name
Discounting	$\otimes$	:	TRANSITIVITY
Consensus	$\oplus$	$\Diamond$	FUSION
Conditional deduction	$\circledcirc$	$\Vert$	DEDUCTION (Modus Ponens)
Conditional abduction	$\overline{\circledcirc}$	$\overline{\Vert}$	ABDUCTION (Modus Tollens)

Example belief model

Simple Trust Network Demo

Four entities, labelled A, B, C and D have opinios about each other represented as points in triangles. Entity A is trying to form an opinion about D, and receives opinions from B and C as to the trustworthiness of D. Furthermore, A has his own opinions about the trustworthiness of B and C.

Left-click and drag opinion points to set opinion values. Entity A combines these opinions using the Subjective Logic Operators to derive his own opinion about D, as shown by the bottom opinion triangle. In detail, entity A *discounts* B's opinion about D by his opinion about B, and does similarly for C. Finally, he combines the two discounted opinions using the *consensus* operator in order to determine his opinion about D. Right-click on the opinion triangles to see the exact values of each opinion. Opinion values can also be visualised using three-coloured rectangles.

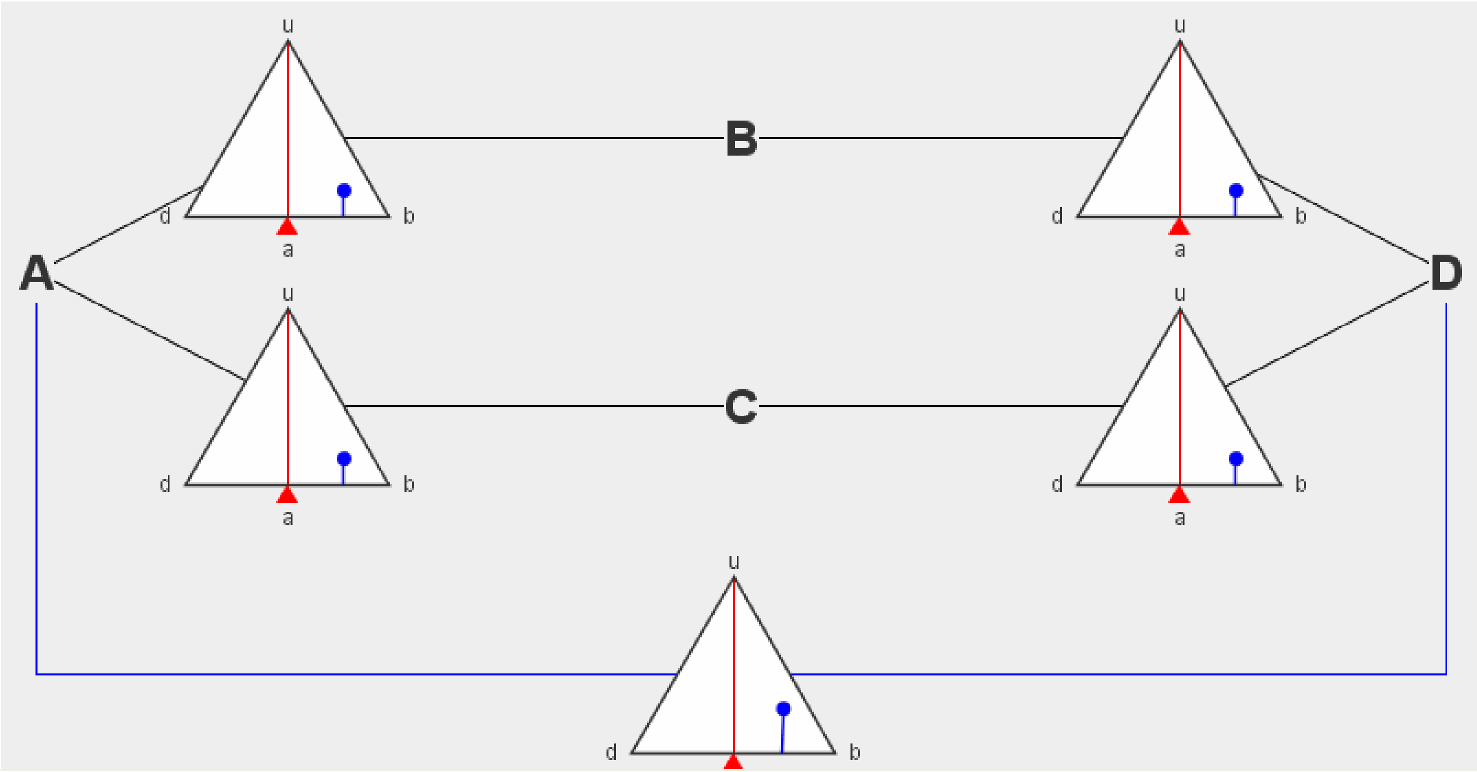

Simplifying complex trust networks

- Trust graphs can contain dependent paths, e.g.:

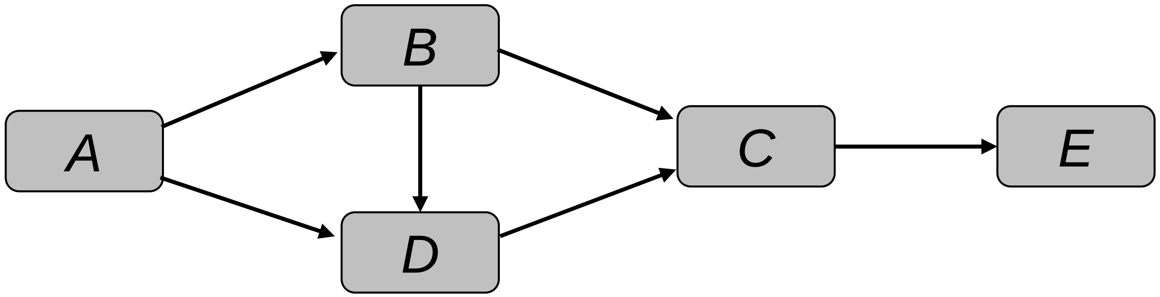

- One path can be removed to produce e.g.:

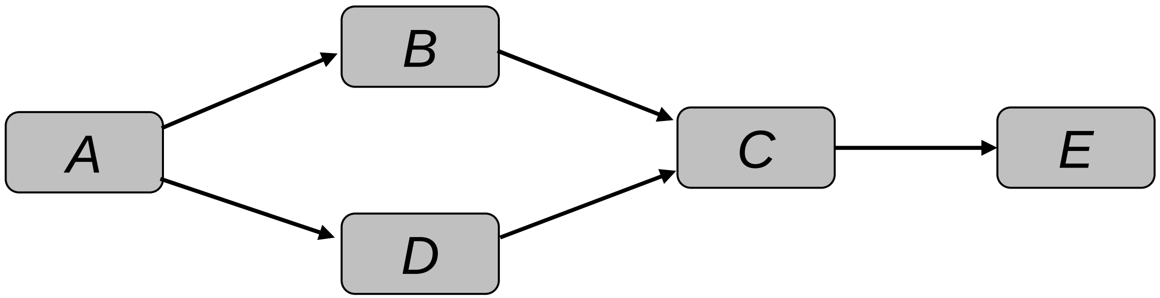

Building series-parallel graphs

- Trust graph analysis with subjective logic requires independent paths
  - called series-parallel graphs
- Constructed with series compositions and parallel compositions

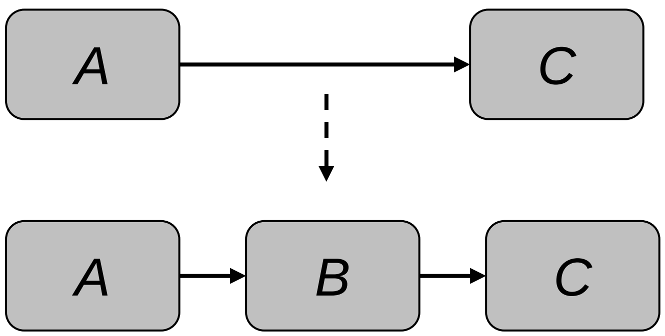

a) Directed series composition

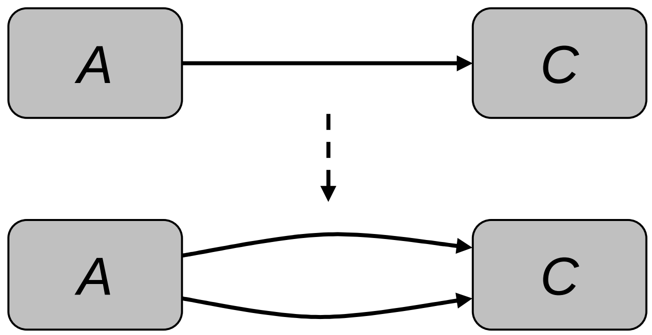

b) Directed parallel composition

Method for building independent graphs

1. Determine all possible paths from relying party to trusted party through initial graph
2. Rank all paths on trust confidence/certainty
3. Build series-parallel graph by
   – including paths one-by-one according to rank
   – rejecting paths that can not be included with series-parallel composition

Resulting graph contains no dependent paths

– can be directly analysed with subjective logic

Series-parallel trust graph examples

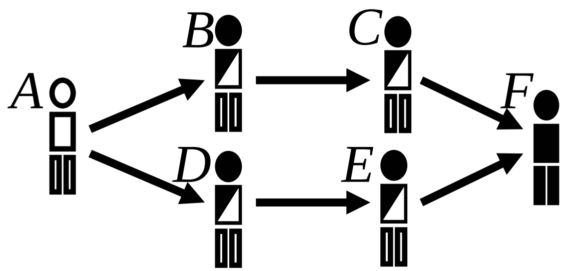

[A,F] = ([A,B]:[B,C]:[C,F]) ◊([A,D]:[D,E]:[E,F]) graph expr.

$$\omega_F^A = (\omega_B^A \otimes \omega_C^B \otimes \omega_F^C) \oplus (\omega_D^A \otimes \omega_E^D \otimes \omega_F^E)$$ SL expr.

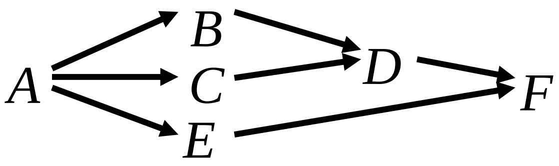

[A,F] = ((([A,B]:[B,D]) ◊ ([A,C]:[C,D])) : [D,F])) ◊ ([A,E]:[E,F])

$$\omega_F^A = (((\omega_B^A \otimes \omega_D^B) \oplus (\omega_C^A \otimes \omega_D^C)) \otimes \omega_F^D) \oplus (\omega_E^A \otimes \omega_F^E)$$

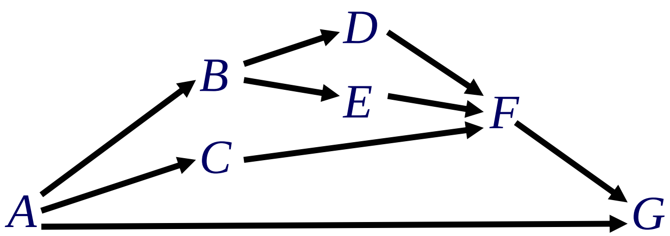

[A,G] = ((([A,B] : ([B,D]:[D,F]) ◊ ([B,E]:[E,F])) ◊ ([A,C]:[C,F])) : [F,G]) ◊ [A,G]

$$\omega_G^A = (((\omega_B^A \otimes (\omega_D^B \otimes \omega_F^D) \oplus (\omega_E^B \otimes \omega_F^E)) \oplus (\omega_C^A \otimes \omega_F^C)) \otimes \omega_G^F) \oplus \omega_G^A$$

Flow models

- Transitive iteration through graph
- Loops and arbitrarily long paths
- Source of trust can be distributed
  - evenly, e.g. early version of PageRank
  - discretely, e.g. current PageRank, EigenTrust
- Sum of trust over all parties can be
  - constant, e.g. PageRank, so one party's increase comes at the cost of another party's decrease
  - function network size, e.g. EigenTrust

Proposed and Commercial trust and reputation systens

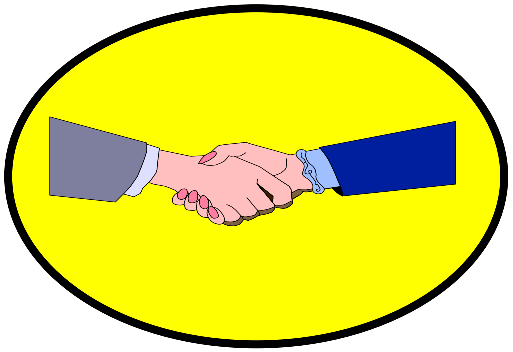

EigenTrust algorithm

- Decision support for P2P networks
- Individual experience recorded
- Based on
  - Normalised local trust scores made public
  - Iterative transitivity to compute global trust
- No negative ratings
- Sum of trust scores in community increases with number of members
- EigenTrust is a reputation system

EigenTrust visualisation

- Local satisfaction score: s_{ij} = sat(i,j) – unsat(i,j)
- Normalised local trust score: $c_{ij} = \max(s_{ij},0)/\Sigma\max(s_i,0)$
- Iterative computation of trust score: $t_{ik} = \Sigma c_{ij} \cdot c_{jk}$
- Iterative vector $t_i = C^n \cdot c_i$ converges to Eigenvector of C.
  - where C is the matrix of all local trust scores

Example with connectivity = 2 so that c_{ij} = 0.5

Local trust scores

+ transitive trust and iterations

Converges to global scores:

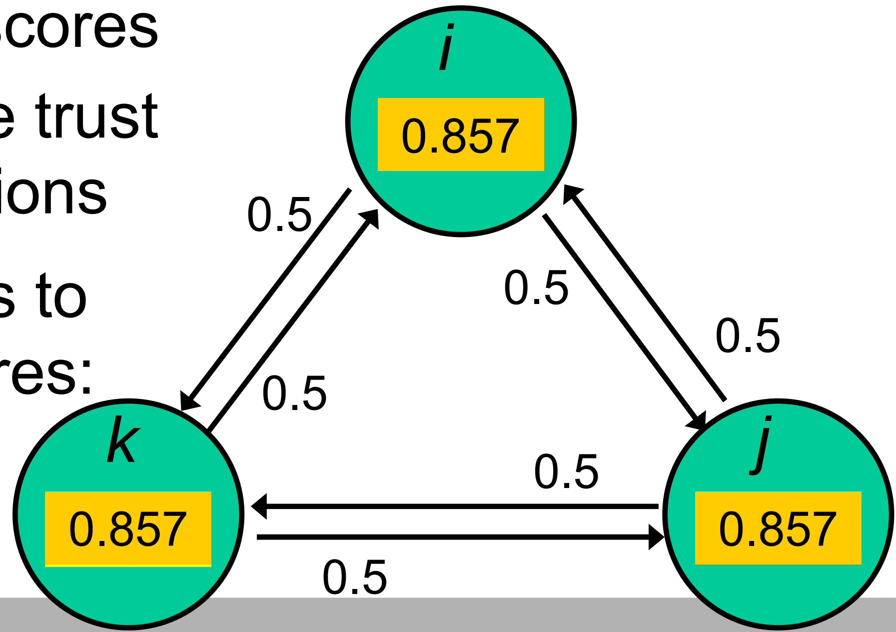

Google's PageRank

- Purpose to provide quality search results
- Based on:
  - Number of incoming links, weighted by the
  - PageRank of the sites behind incoming links
- Hyperlinks interpreted as positive ratings.
- No negative ratings.
- Random surfer model.
- PageRank is a reputation system

PageRank visualisation

- $R(A) = (1-d)/N(Web) + d \cdot \Sigma R(prev(A))/N(next(prev(A)))$
- Damping factor $d \approx 0.85$
- $\Sigma R(A) \approx 1$, i.e. $R(A)$ is the probability of the random surfer
- $PageRank(A) = I + \log_{\approx 10} R(A)$, where $I \approx 11$

Example

with N(Web)=3

Initial rank R

+ imported rank

and iterations

Converges to:

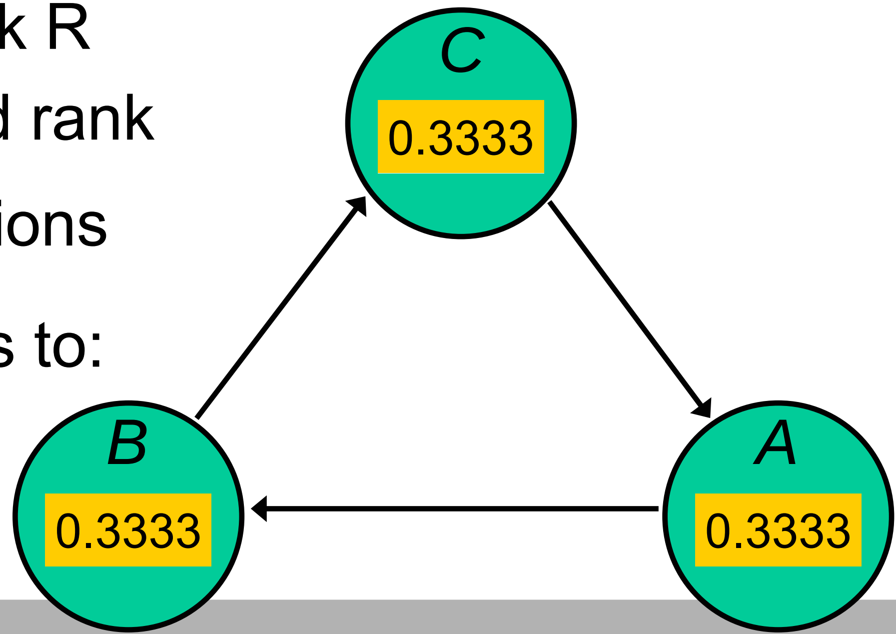

Link spam and “nofollow”

- Survival of e-commerce sites depends on rank
- Possible to increase rank with link spam
  - consists of putting URLs to own Web site in wikis (publicly editable Web sites) and in postings to public discussion groups
- The “nofollow” tag, introduced in 2005, instructs Web crawlers not to follow a link

```
<a href=http://some-spammer-website.com
rel="nofollow">Link</a>
```

- Wikis and discussion groups now enforce that all URLs have “nofollow”, thereby solving the link spam problem

Negative side-effects of “nofollow”

- Outgoing URLs causes rank leak
- Many webmasters misuse “nofollow” to avoid leaking Web ranking
- Undermines basis for original PageRank algorithm
- Alternative info sources required for ranking Web pages
  - Toolbars
  - Reputation systems

The future of Web page ranking

- 1990s: No ranking, random order (Altavista)
  - Boolean selection criteria possible
- 1998: Random surfer model
  - Based on PageRank algorithm
- 2005: Intentional surfer model
  - Based on Toolbar feedback
- 2008: Critical surfer model
  - Based on reputation systems

Random surfer model

- Assumes a monkey that randomly clicks on Web links.
- The monkey is the random surfer.
- Ranking = probability of monkey accessing a given page
- PageRank algorithm is the basis for this model

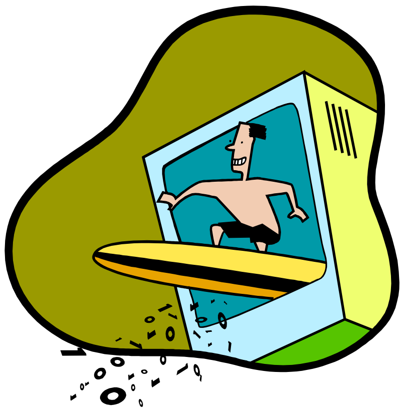

Intentional surfer model

- Assumes people who actually surf the Web
- Ranking = probability of people accessing a given page
- Difficult to obtain global information about how often a page is actually accessed.
- Browser toolbar provides source of info

Browser toolbar architecture

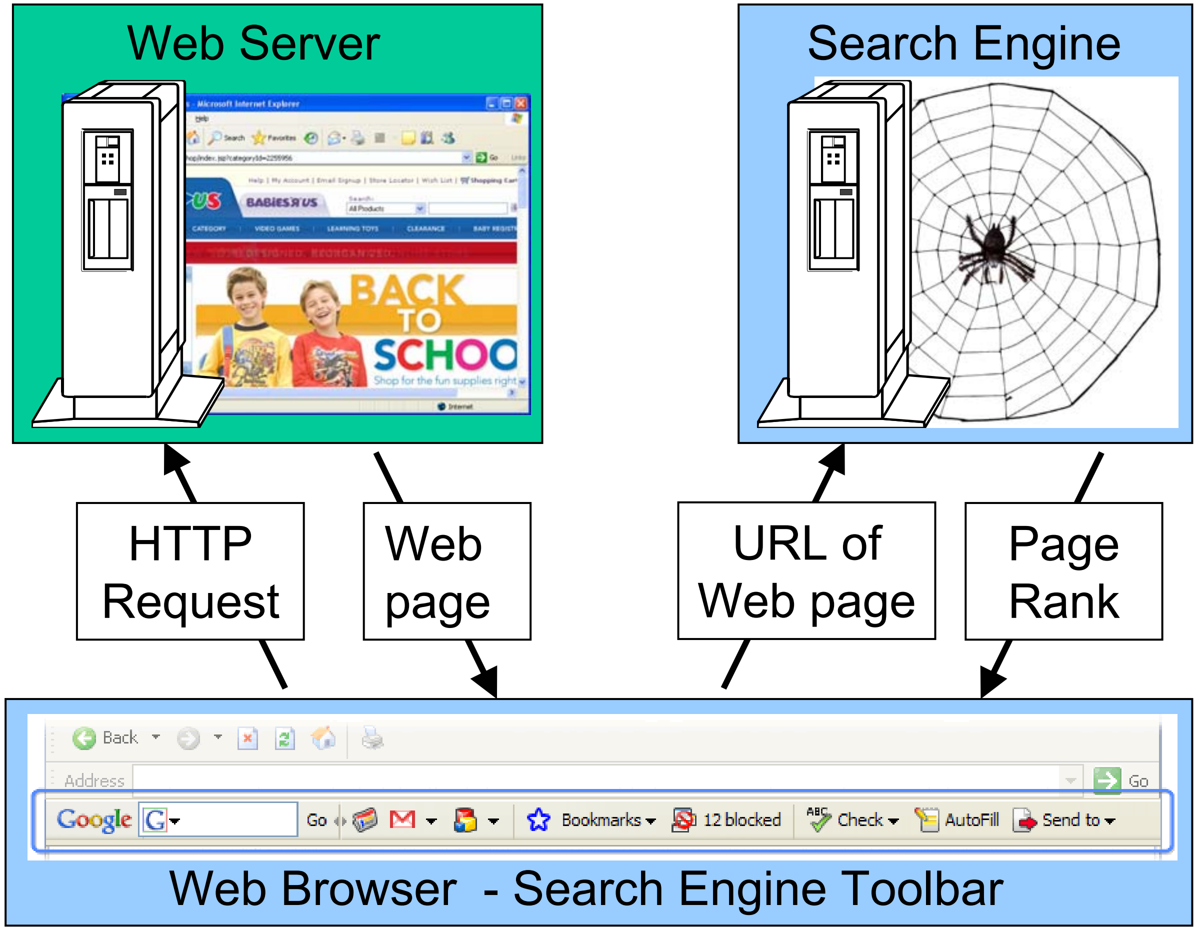

Evidence from toolbars and spyware

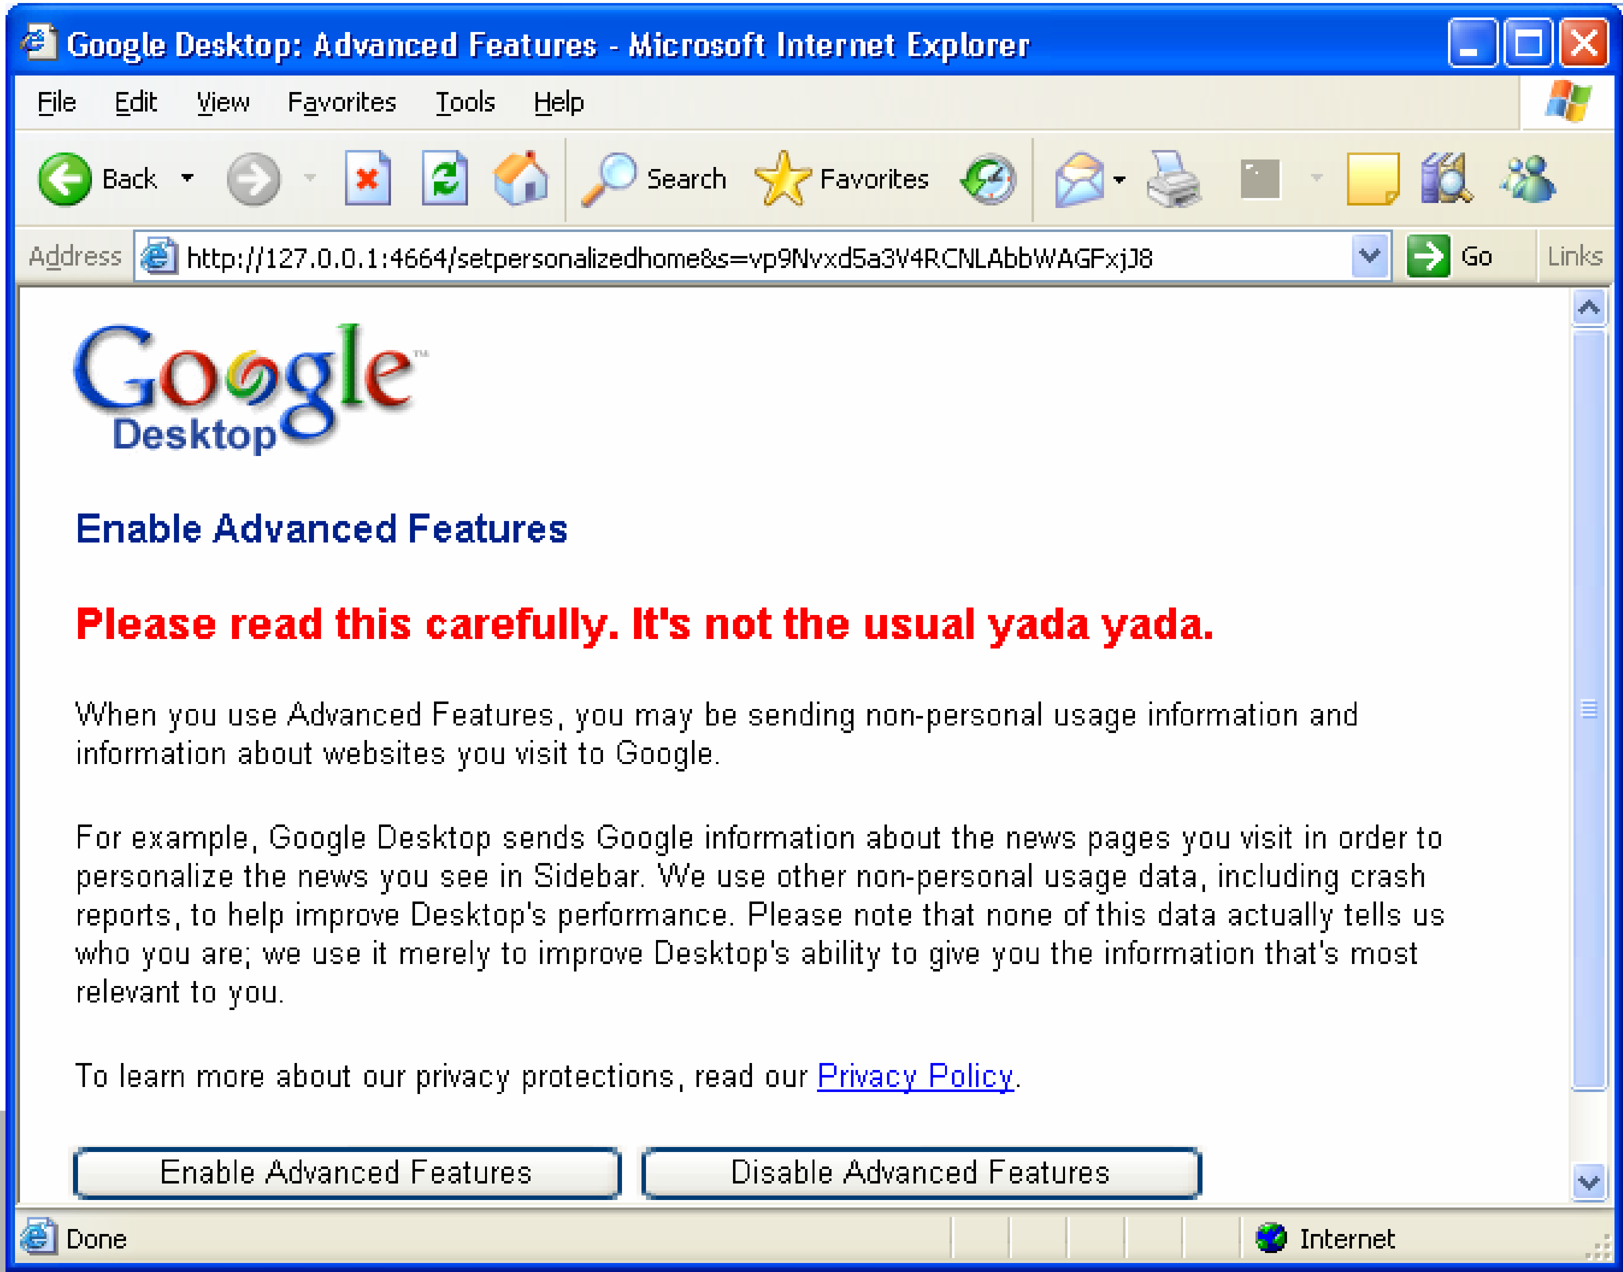

Critical surfer model

- People sometimes access a Web site even though they don't approve of its content
  - e.g. IT security researcher investigating phishing sites
- Critical surfer model depends on people rating Web pages
- Ranking = probability of people accessing a given page, weighted by its reputation score

Critical surfer model implementation

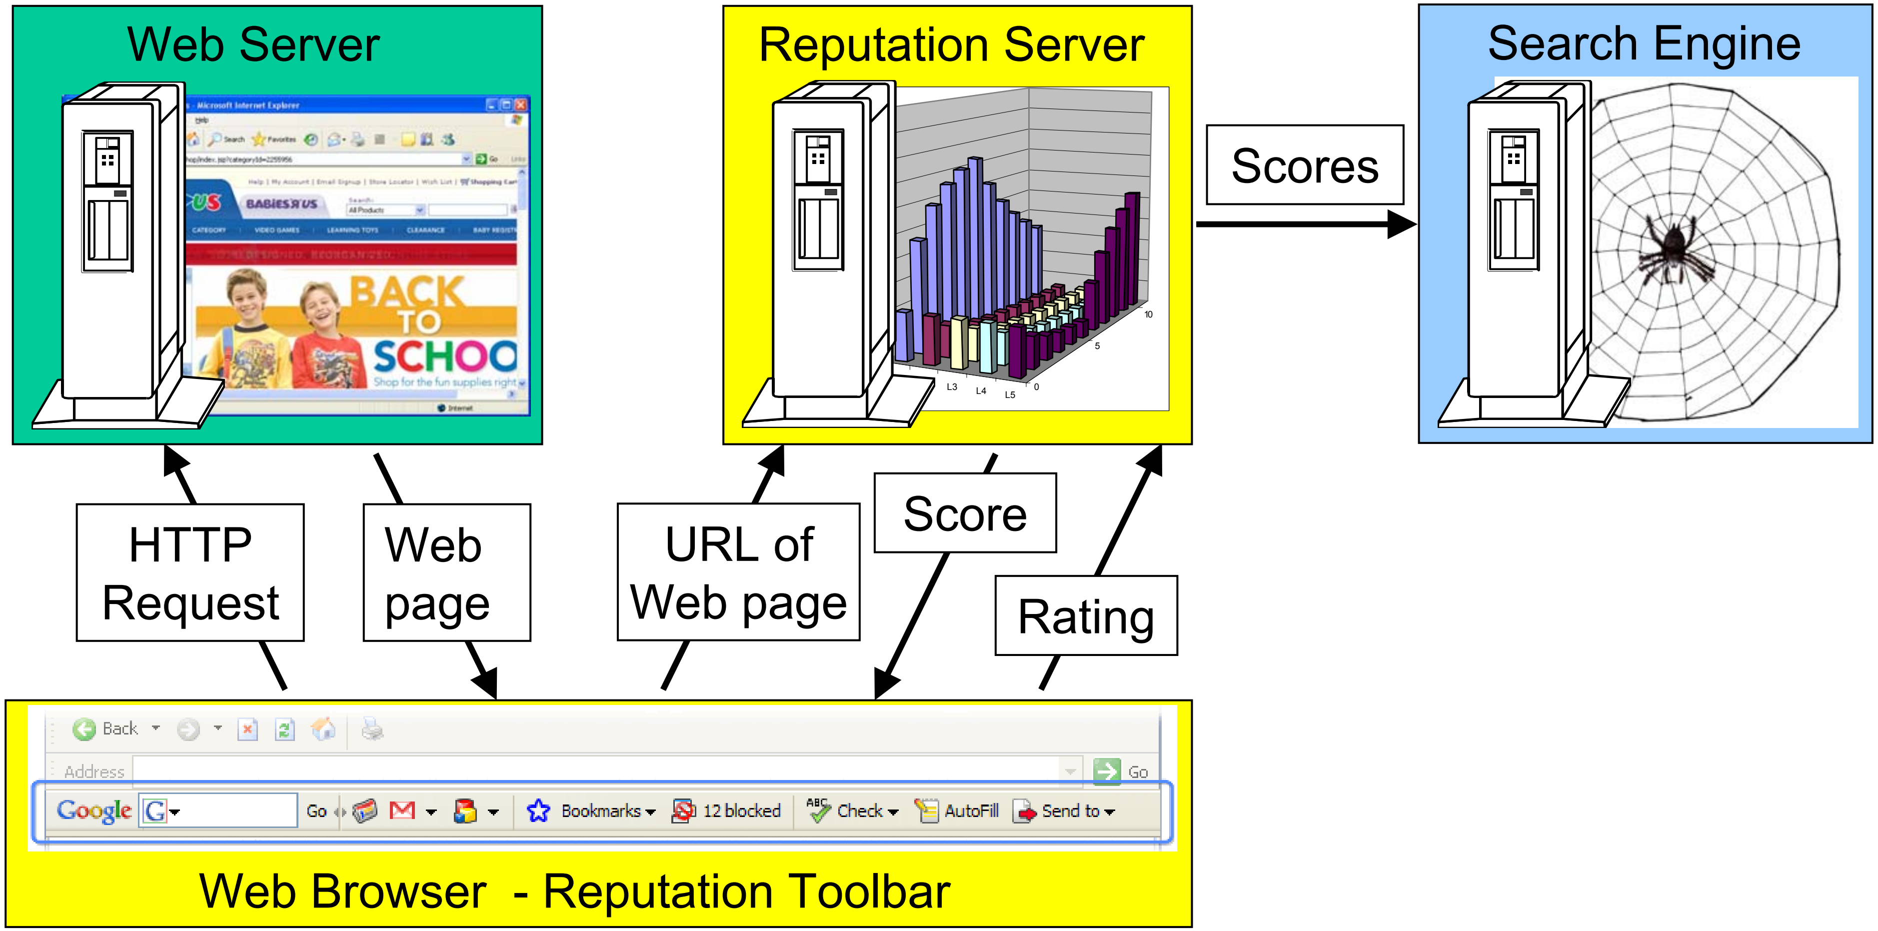

Web Sites with reputation systems

- Auction sites:
  - www.ebay.com
  - auctions.yahoo.com
- Expert sites
  - www.expertcentral.com
  - www.askme.com
- Product review sites
  - www.epinions.com
  - www.amazon.com
- e-commerce
  - www.bizrate.com
  - www.virtualratings.com
- Article postings
  - www.slashdot.com
  - www.everything2.org
- Education
  - us.ratemyteachers.com
  - www.virtualratings.com
- Entertainment
  - www.citysearch.com
  - www.imdb.com
  - radio.weblogs.com

The eBay Feedback Forum

- Centralised reputation system
- Ratings:
  - Buyers and sellers rate each other, 50% - 60% times
  - positive, negative, neutral, + short comment
- Score = Σ positive - Σ negative
- Time windows
- Surprisingly positive ratings, only 1% negative
- Correlation between seller and buyer ratings
- Many empirical studies
- Purpose: to control the quality of market

Example eBay member’s profile

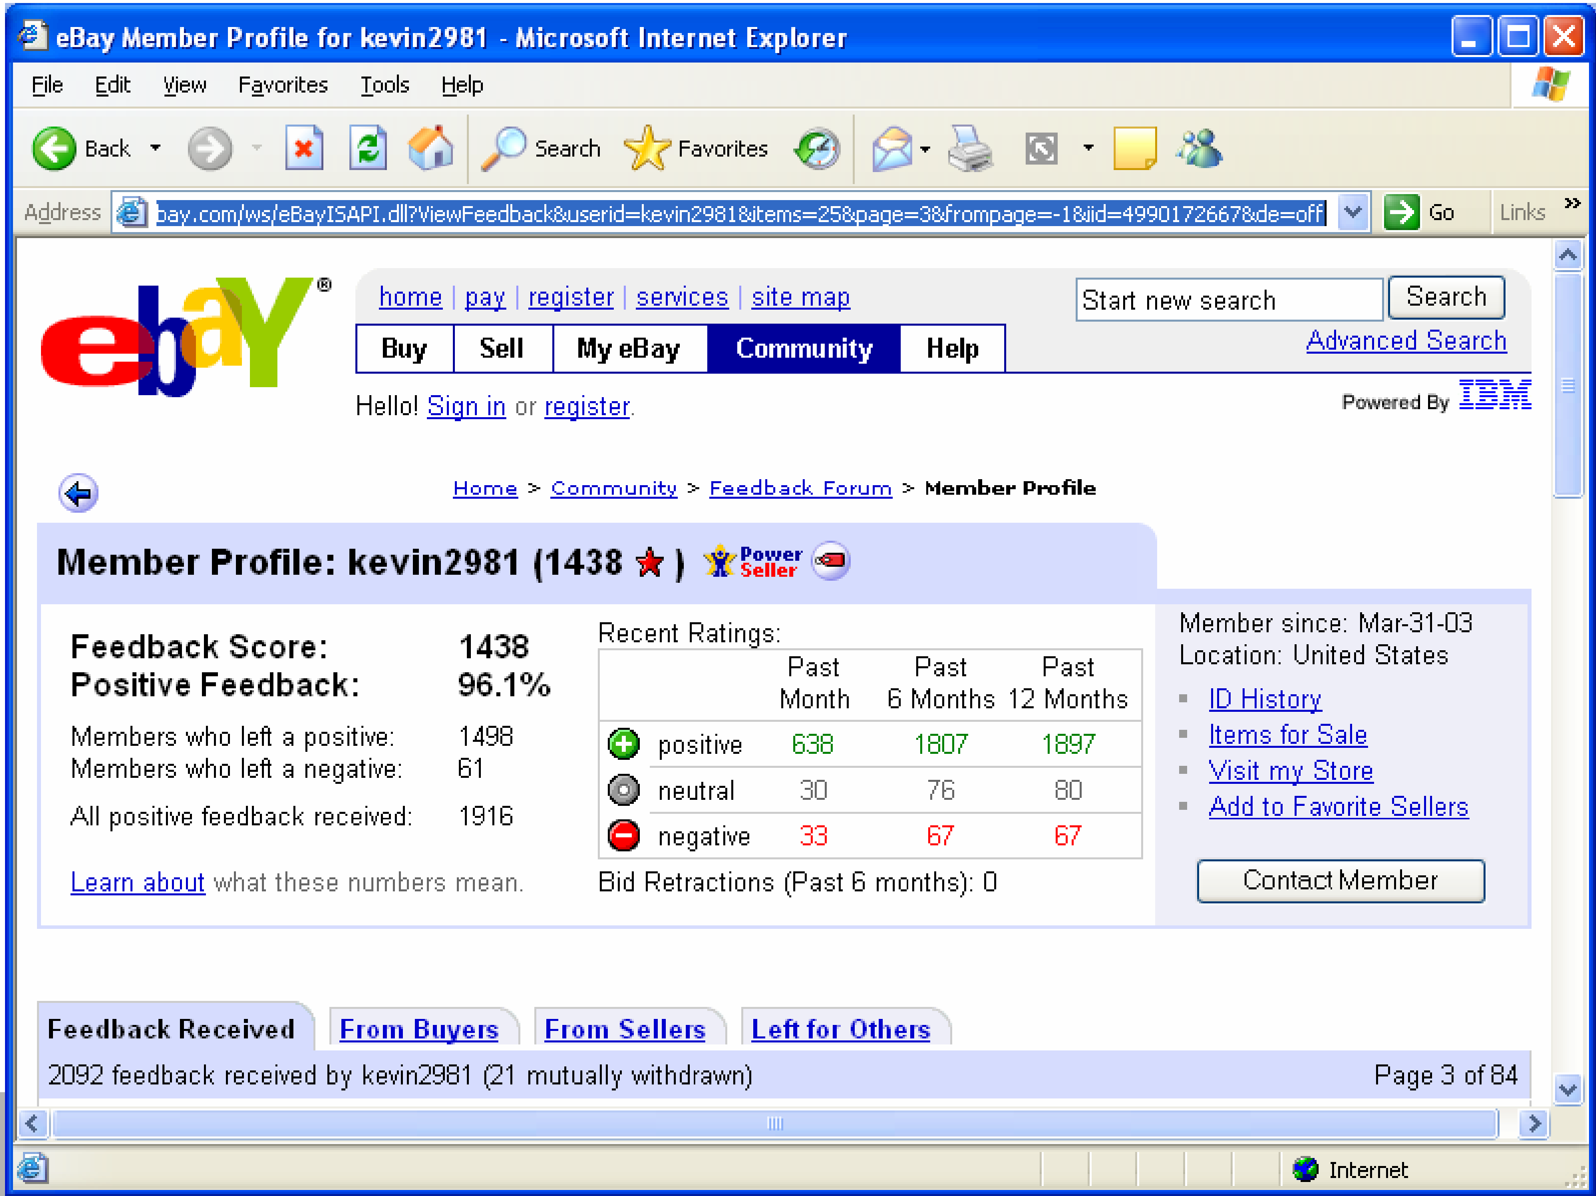

Example eBay feedback comments

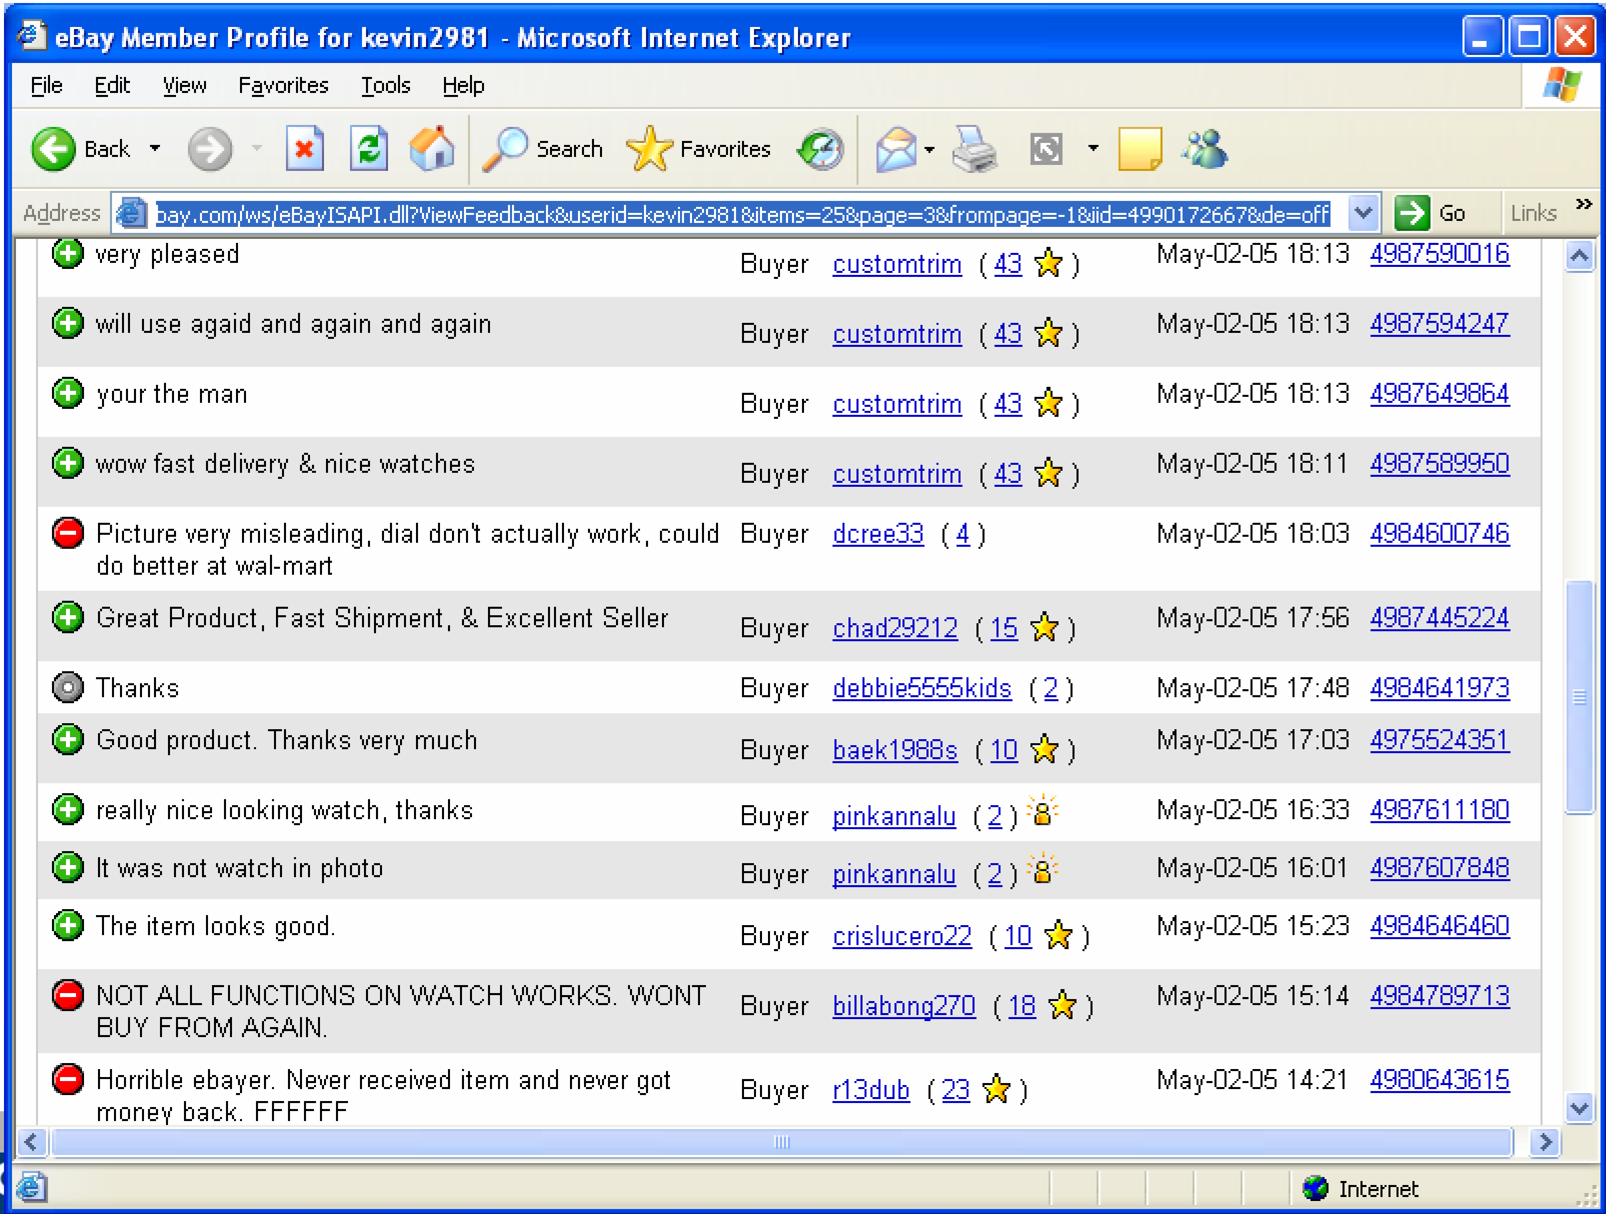

very pleased	Buyer	customtrim (43)	May-02-05 18:13	4987590016
will use agaid and again and again	Buyer	customtrim (43)	May-02-05 18:13	4987594247
your the man	Buyer	customtrim (43)	May-02-05 18:13	4987649864
wow fast delivery & nice watches	Buyer	customtrim (43)	May-02-05 18:11	4987589950
Picture very misleading, dial don't actually work, could do better at wal-mart	Buyer	dcree33 (4)	May-02-05 18:03	4984600746
Great Product, Fast Shipment, & Excellent Seller	Buyer	chad29212 (15)	May-02-05 17:56	4987445224
Thanks	Buyer	debbie5555kids (2)	May-02-05 17:48	4984641973
Good product. Thanks very much	Buyer	baek1988s (10)	May-02-05 17:03	4975524351
really nice looking watch, thanks	Buyer	pinkannalu (2)	May-02-05 16:33	4987611180
It was not watch in photo	Buyer	pinkannalu (2)	May-02-05 16:01	4987607848
The item looks good.	Buyer	crislucero22 (10)	May-02-05 15:23	4984646460
NOT ALL FUNCTIONS ON WATCH WORKS. WONT BUY FROM AGAIN.	Buyer	billabong270 (18)	May-02-05 15:14	4984789713
Horrible ebayer. Never received item and never got money back. FFFFFF	Buyer	r13dub (23)	May-02-05 14:21	4980643615

Reputation extortion on eBay

- Serious sellers
  - want satisfied customers
  - don’t want negative feedback
- Dissatisfied buyers can contact seller before giving negative feedback
- Threat of negative feedback can work better in customer’s favour than actual negative feedback
- Proves that reputation systems work

AllExperts

- Free advice from volunteer experts
- Ratings given on scale [1,10] for
  - Knowledgeable, Clarity of response, Timeliness and Politeness
- Score = average of ratings
- Most experts have scores ≈ 10
- Business model:
  - Low profile advertisement
  - Prestige to volunteer experts

Example AllExperts profile

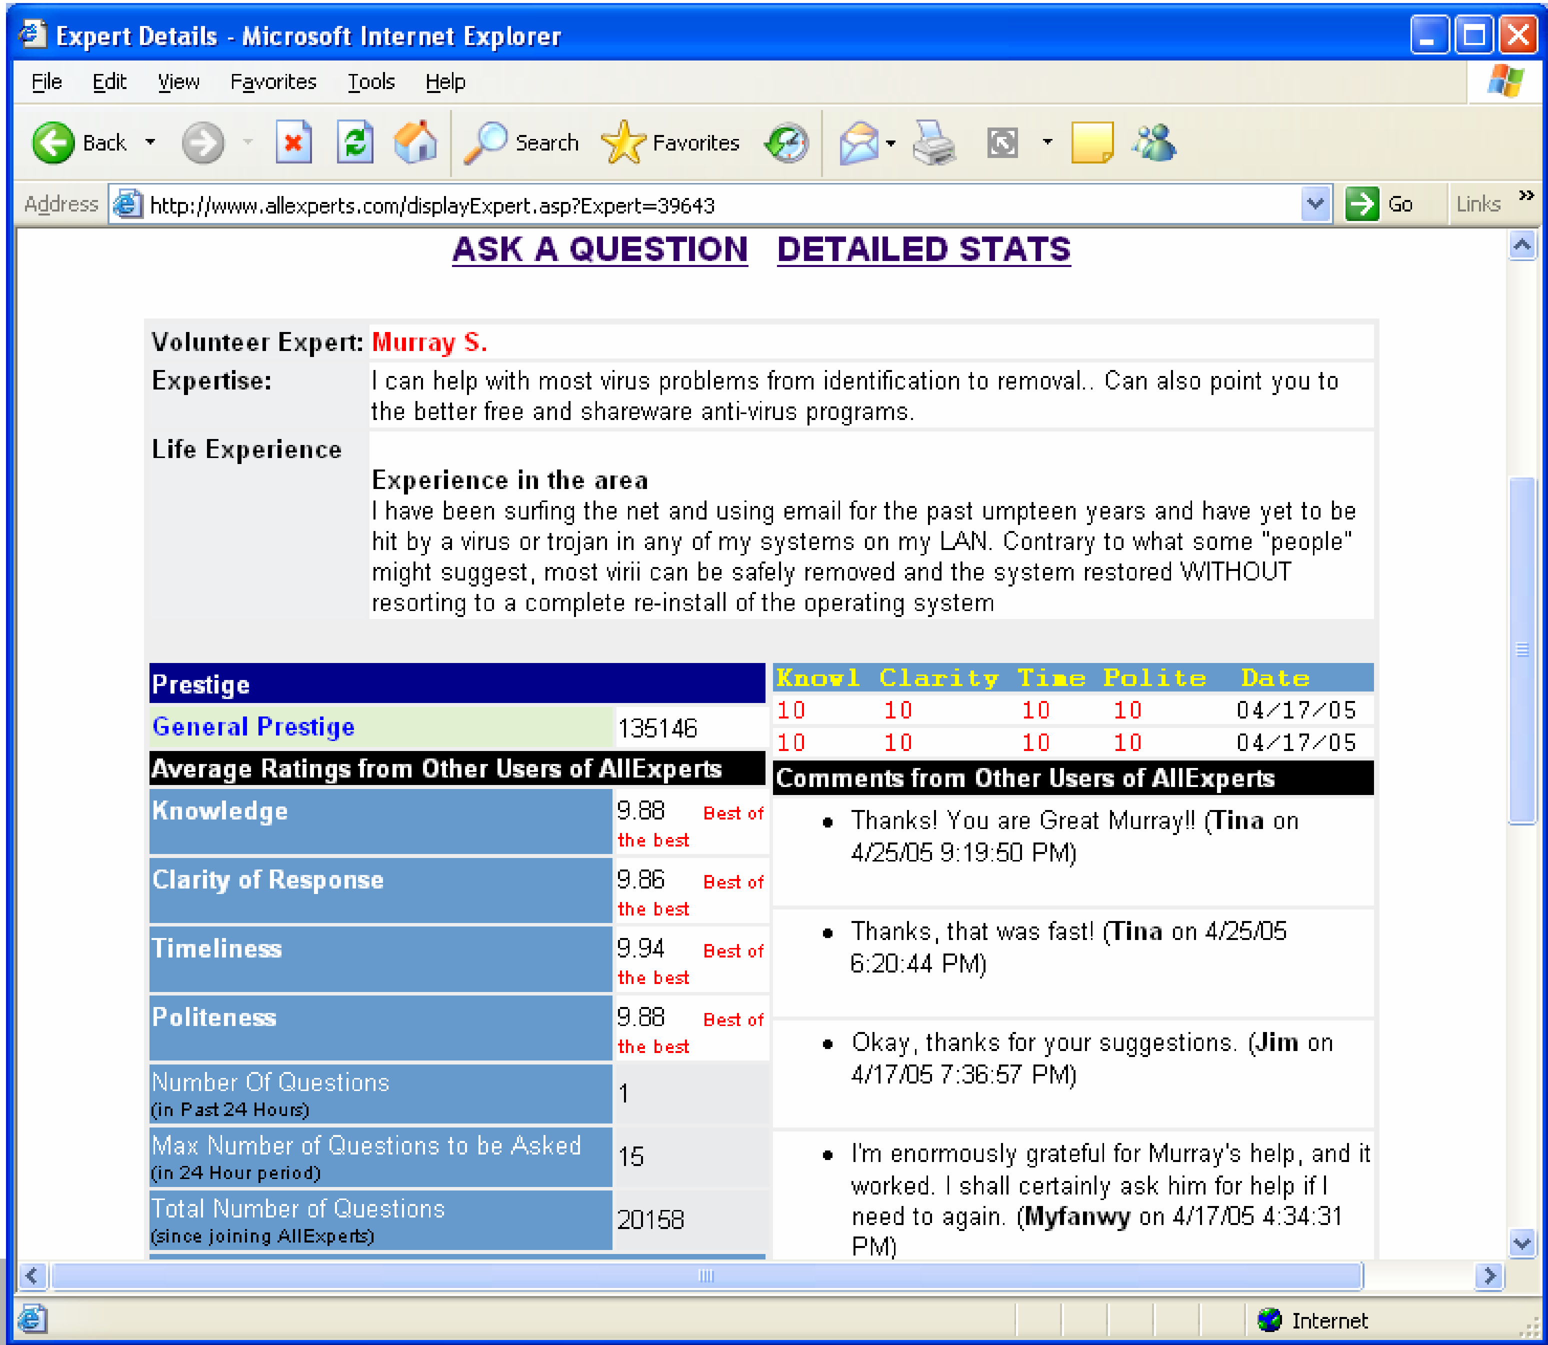

Expert Details - Microsoft Internet Explorer

File Edit View Favorites Tools Help

Back Search Favorites

Address http://www.allexperts.com/displayExpert.asp?Expert=39643 Go Links

ASK A QUESTION DETAILED STATS

Volunteer Expert: Murray S.

Expertise: I can help with most virus problems from identification to removal.. Can also point you to the better free and shareware anti-virus programs.

Life Experience

Experience in the area
I have been surfing the net and using email for the past umpteen years and have yet to be hit by a virus or trojan in any of my systems on my LAN. Contrary to what some "people" might suggest, most virii can be safely removed and the system restored WITHOUT resorting to a complete re-install of the operating system

Prestige		
General Prestige	135146	
Average Ratings from Other Users of AllExperts		
Knowledge	9.88	Best of the best
Clarity of Response	9.86	Best of the best
Timeliness	9.94	Best of the best
Politeness	9.88	Best of the best
Number Of Questions (in Past 24 Hours)	1	
Max Number of Questions to be Asked (in 24 Hour period)	15	
Total Number of Questions (since joining AllExperts)	20158	

Knowl	Clarity	Time	Polite	Date
10	10	10	10	04/17/05
10	10	10	10	04/17/05

Comments from Other Users of AllExperts

- Thanks! You are Great Murray!! (**Tina** on 4/25/05 9:19:50 PM)
- Thanks, that was fast! (**Tina** on 4/25/05 6:20:44 PM)
- Okay, thanks for your suggestions. (**Jim** on 4/17/05 7:36:57 PM)
- I'm enormously grateful for Murray's help, and it worked. I shall certainly ask him for help if I need to again. (**Myfanwy** on 4/17/05 4:34:31 PM)

Internet

Advogato open-source community

- Community of programmers
- Hierarchic flow model reputation system
- Flow capacities assigned as a function of distance from root seed
- Computation based on Ford-Fulkerson algorithm for flow through graphs
- Recommendations as
  - Apprentice, journeyer, or master
- Purpose: give prestige to members

Advogato trust structure

Advogato's trust metric - Microsoft Internet Explorer

File Edit View Favorites Tools Help

Back Search Favorites

Address http://www.advogato.org/trust-metric.html Go Links

Assignment of capacities

The next step is to assign a capacity to each node in the graph. This is done by breath-first searching the graph from the seed, computing a shortest distance from the seed to each node.

Then, capacities are assigned simply as a function of this distance. Nodes closer to the root have high capacity, which diminishes with distance. Currently:

cap(0) = 800
cap(1) = 200
cap(2) = 200
cap(3) = 50
cap(4) = 12
cap(5) = 4
cap(6) = 2
cap(*i*) = 1 for *i* > 6

An example of the capacity assignment is shown below, although with much smaller capacities than used on Advogato:

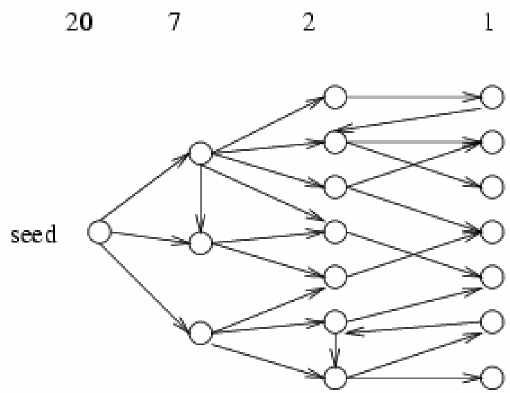

Done Internet

Example Advogato profile

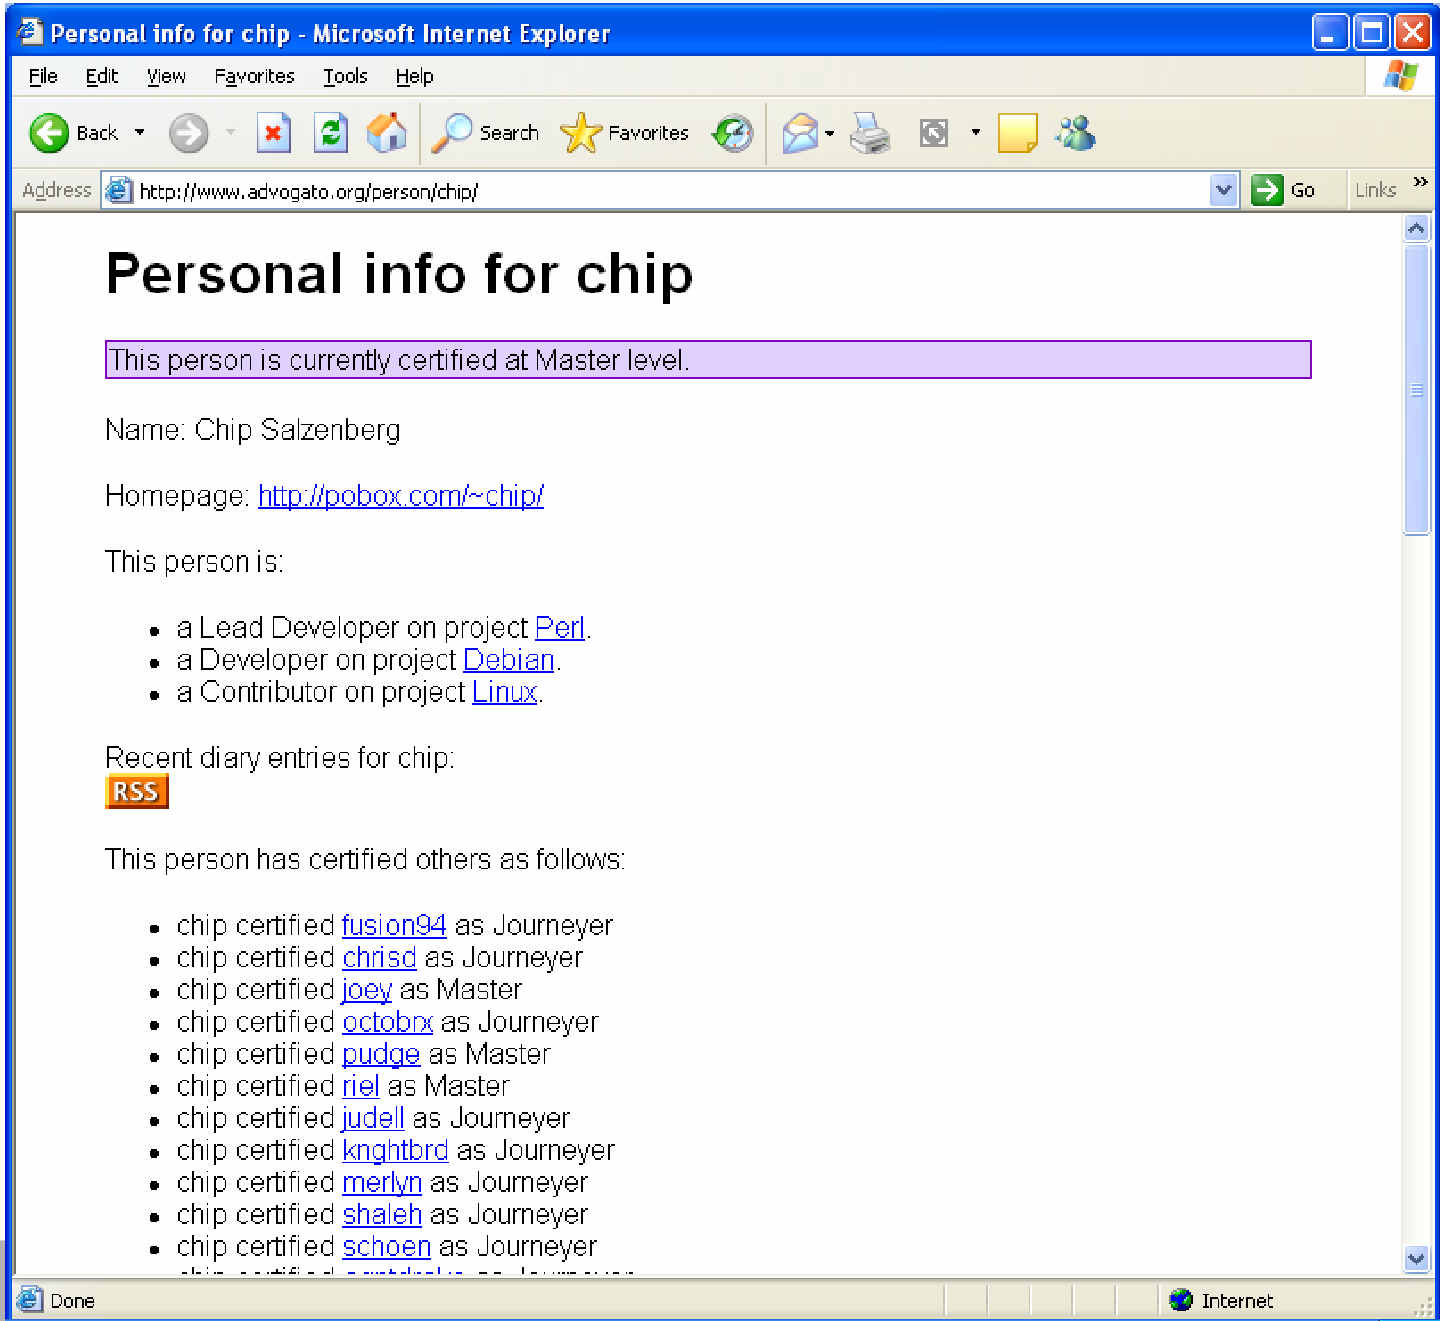

Epinions product review site

- Reviews consumer products
- Product ratings
  - in range 1 – 5 stars
  - Score = average of product ratings
- Review ratings
  - Not helpful, somewhat helpful, helpful, very helpful
  - Review score = average of review ratings
- Reviewer status
  - Member, advisor, top reviewer, category lead
- Income share program
  - Gives cash to reviewers with high number of very helpful reviews

Example Epinions product profile

Example Epinions product reviews

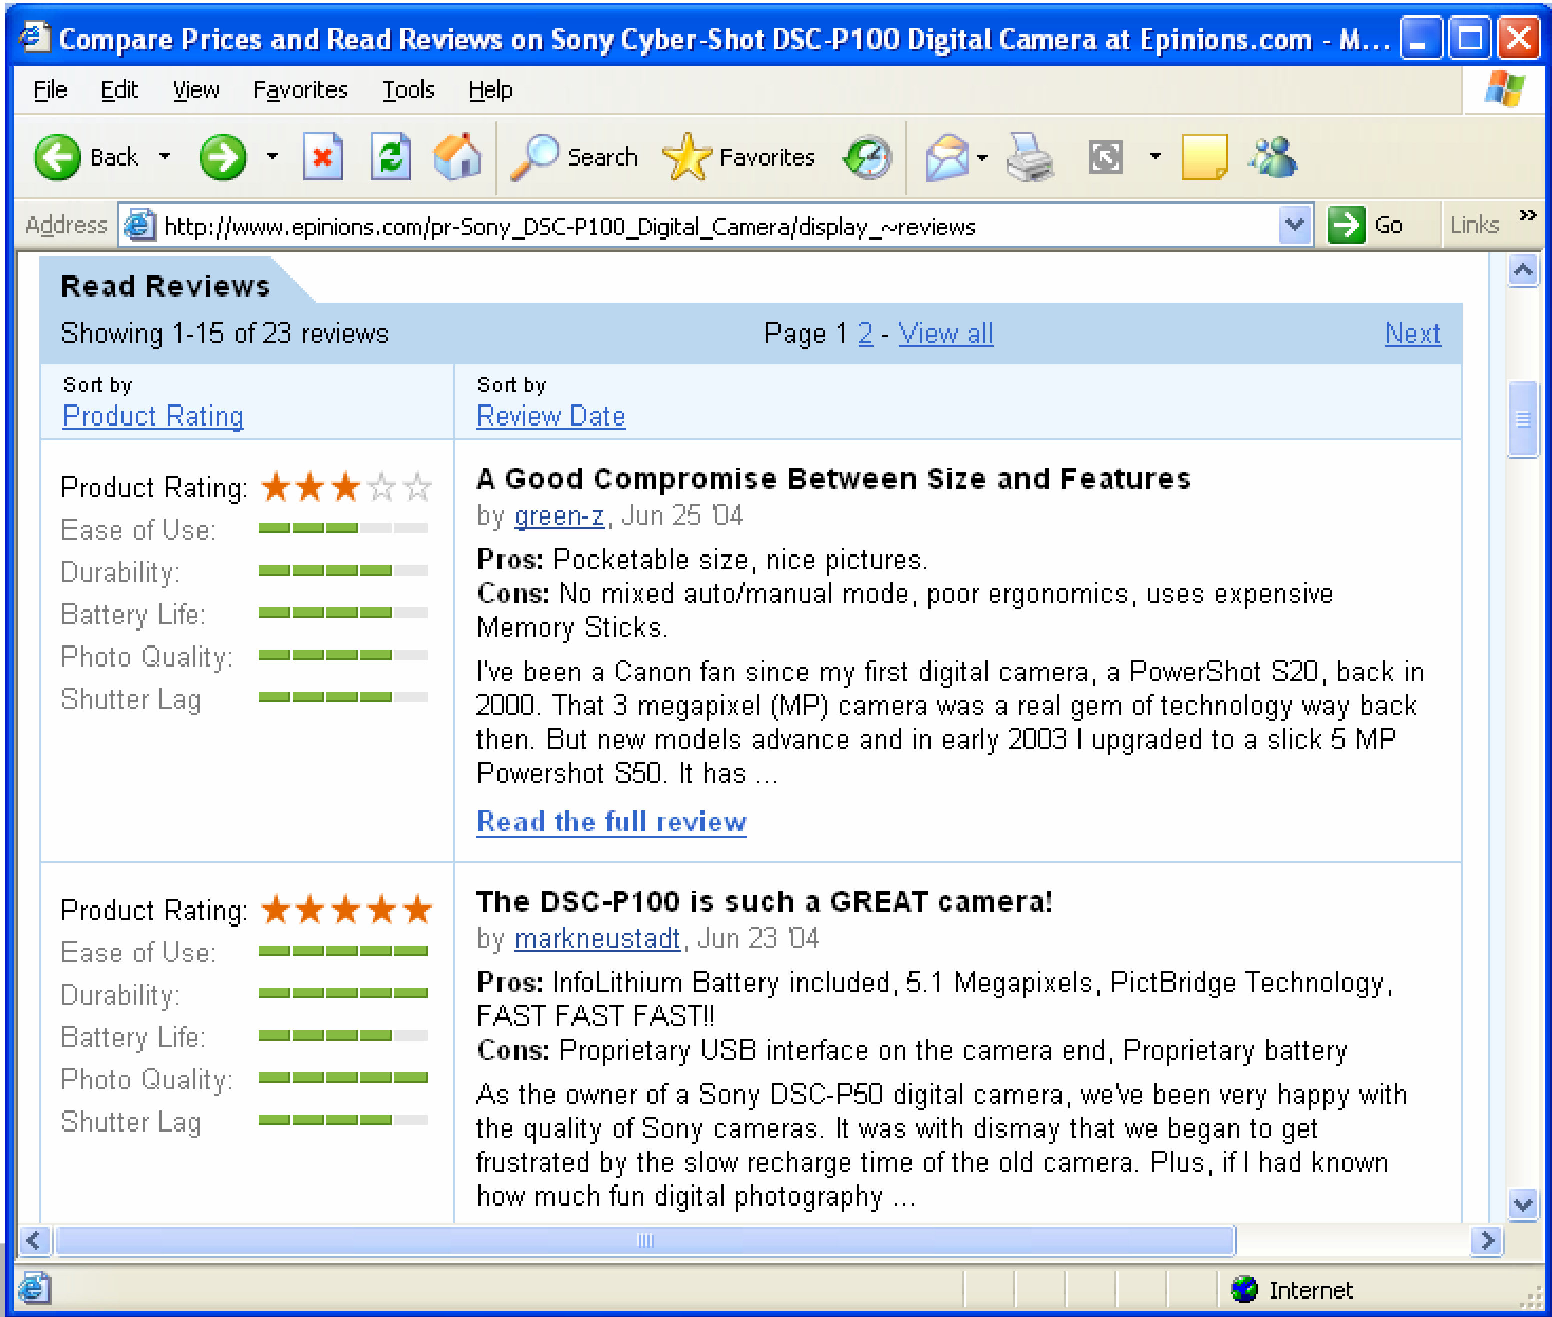

Example Epinions member

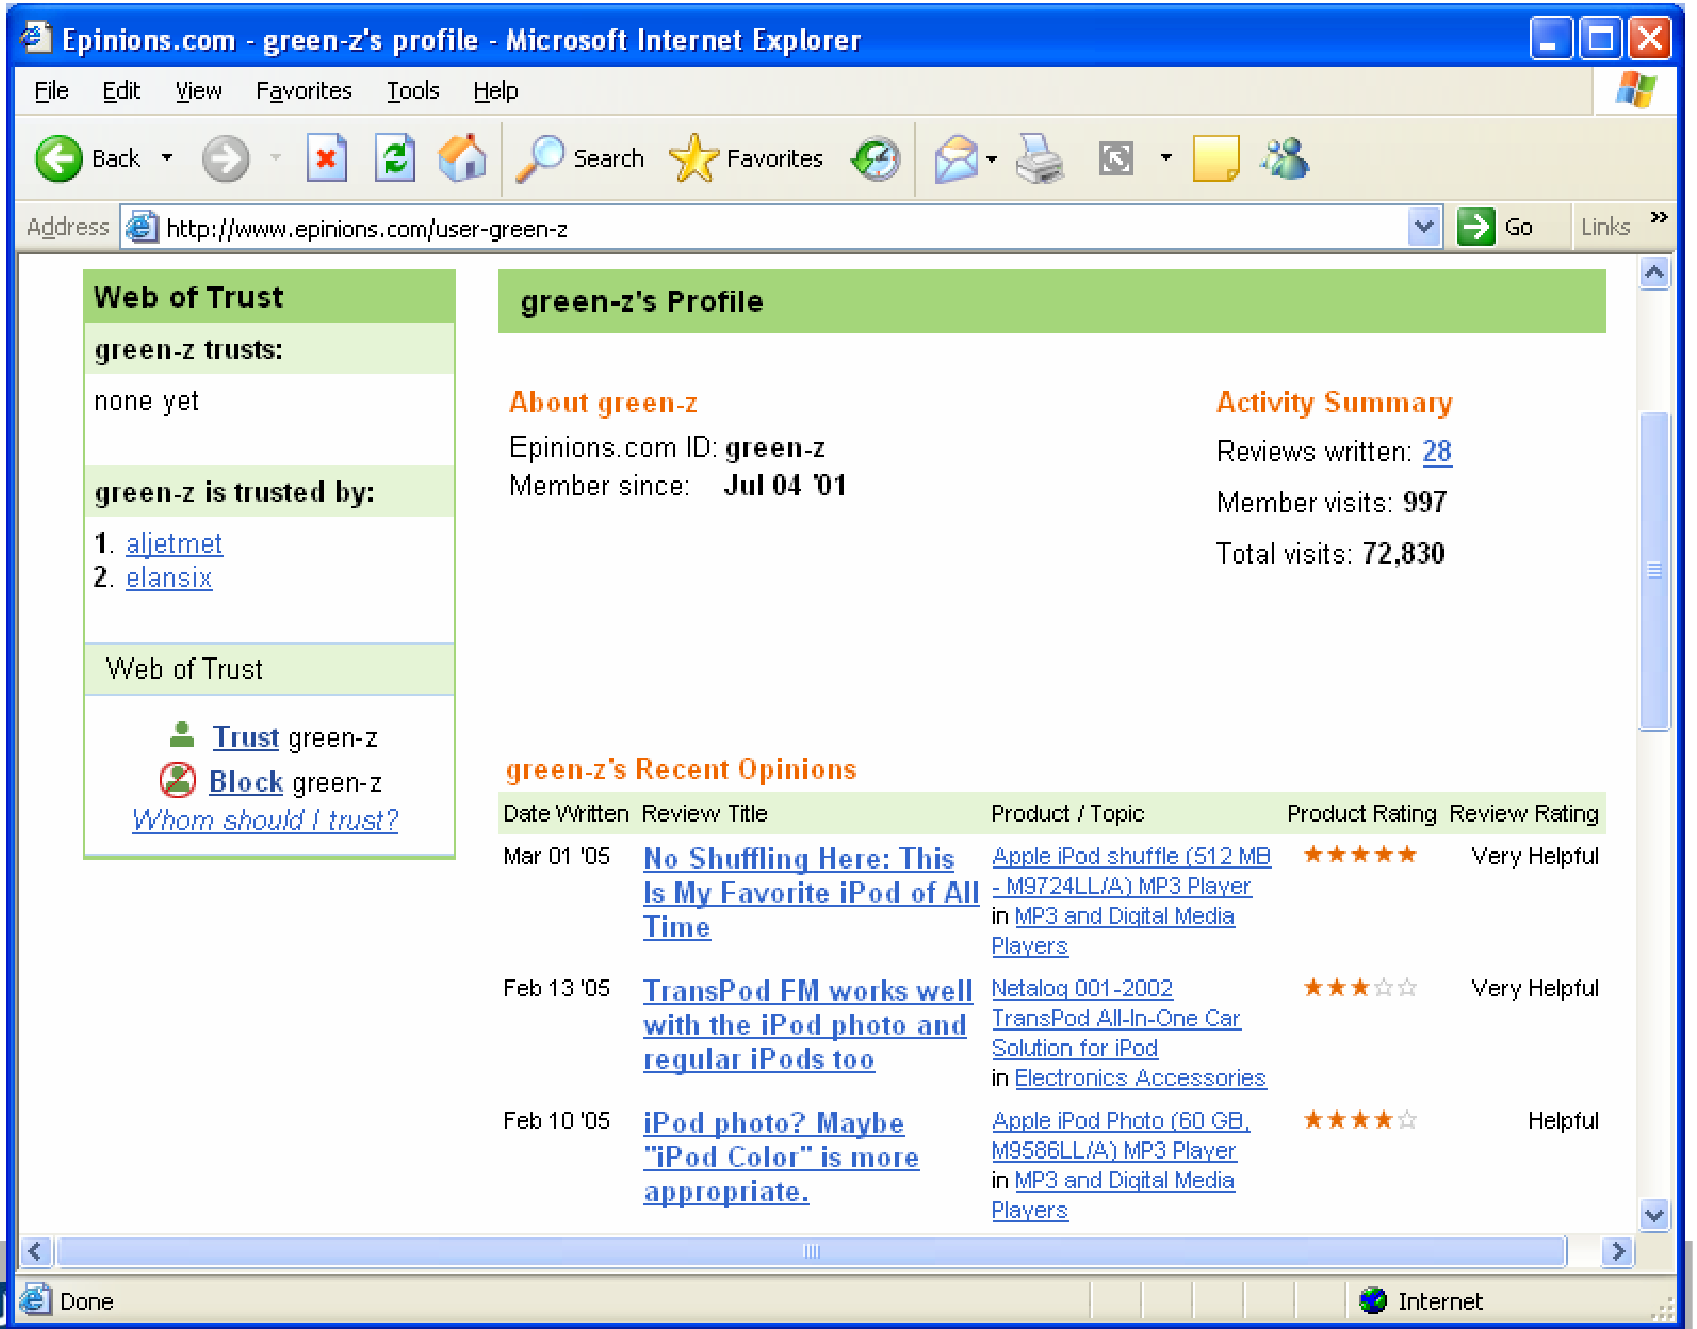

Amazon

- Online book store, with reviews by members
  - Book review in prose
  - Book ratings: 1 – 5 stars
  - Book score = average of book ratings
- Review ratings
  - Helpful or not helpful
  - Reviewer score = Σ helpful ratings
- Reviewer status
  - #1, top 10, top 50, top 100, top 500, top 1000
  - To be the #1 reviewer, you must read more books than any living person could do.

Example Amazon reviews

Amazon.com: Books: Applied Cryptography: Protocols, Algorithms, and Source Code in C, Second Ed - Microsoft Inte...

File Edit View Favorites Tools Help

Back Search Favorites

Address http://www.amazon.com/exec/obidos/ASIN/0471117099/qid=1115343019/sr=2-1/ref=pd_bbs_b_2_1/102-5509256-0267349 Go Links

Customer Reviews

Average Customer Review:

Write an online review **and share your thoughts with other customers.**

The Bible of the Crypto-world, April 19, 2005

Reviewer: Gabriel E. Borlean (San Jose, CA United States) - See all my reviews

REAL NAME™

For anyone working or studying the crypto-world (IT Security, Cryptographer, Cryptologists, PKI/PGP/SecureID Technologists, etc.) this is a must read.

For the fascinating world of Cryptography (the practical application of Cryptology) this is an essential textbook that any graduate course of digital security and cryptology should include. It is inteded as a reference book as well as a practical book to have handy for the working professional.

It's no wonder the NSA had asked the author not to publish this book. It leaves the Crypto-world naked to the average viewer/reader.

Was this review helpful to you? yes no (Report this)

3 of 4 people found the following review helpful:

Excellent Text - Not Too Mathematical, August 26, 2004

Reviewer: S GARDNER (Fair Lawn, NJ) - See all my reviews

REAL NAME™

This book provides excellent coverage of the mayor cryptography algorithms. It is a must have (for academic study or implementation), if only for the completeness of coverage and the comprehensive references. The C code provided is adequate, but reasonable programming skills are assumed. The book is not overly mathematical (which many of the other cryptography texts are - this is great for me, as I am more intereseted in practice). There is also good information on assessing / comparing the merits of different algorithms. Great for the practictioner or student taking a first course (it's still not bedtime reading). Mathematicians will look for more ......

Internet

Example Amazon member profile

Slashdot

- *“News for nerds”* message board
- Article postings, at Shlasdot’s discretion
- Comments to articles posted by members
- Comment moderation by members
  - Positive: insightful, interesting, informative funny, underrated
  - Negative: offtopic, flamebait, troll, redundant, overrated
  - Comment score ≈ Σ positive(Karma) - Σ negative(Karma),
  - Moderation by members with high Karma carries more weight
- Comment viewing filtered by score
- Member Karma
  - Terrible, bad, neutral, positive, good, excellent
  - Based on moderation of comments.
- Metamoderation, to combat unfair moderation
  - Rate the moderations: fair, unfair, neutral
  - Affects Karma of member who gave the moderation
- Arbitrary moderation by Shlashdot staff
- Purpose: Directing massive collaborative moderation effort

Hierarchic reputation architecture

Shlashdot type

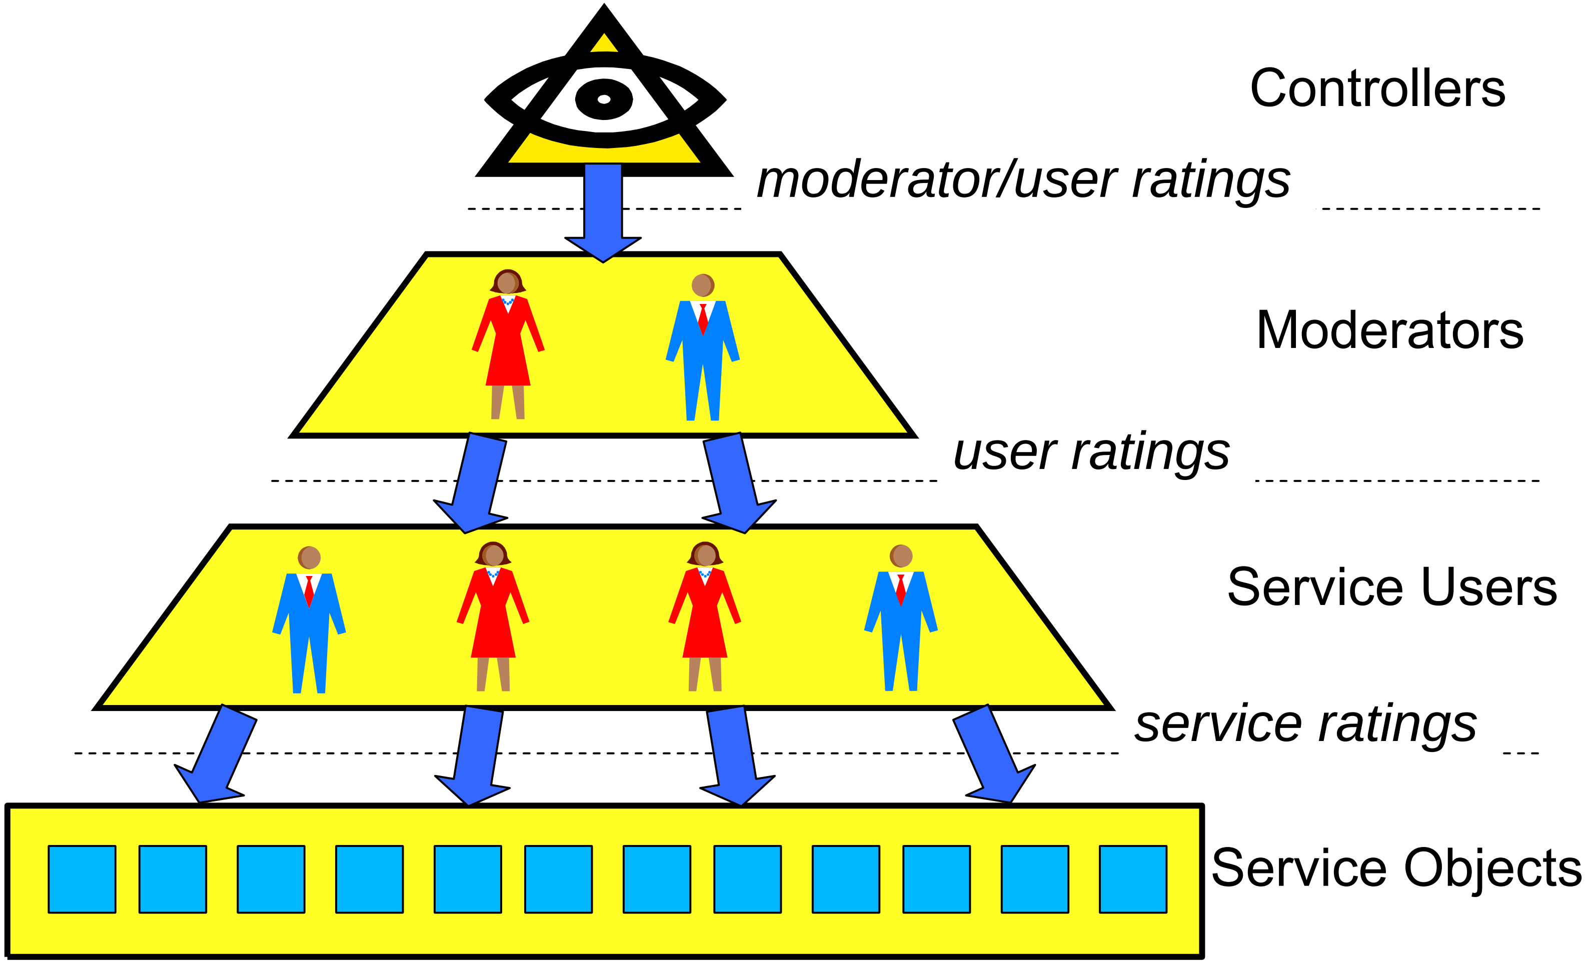

Example Slashdot posting

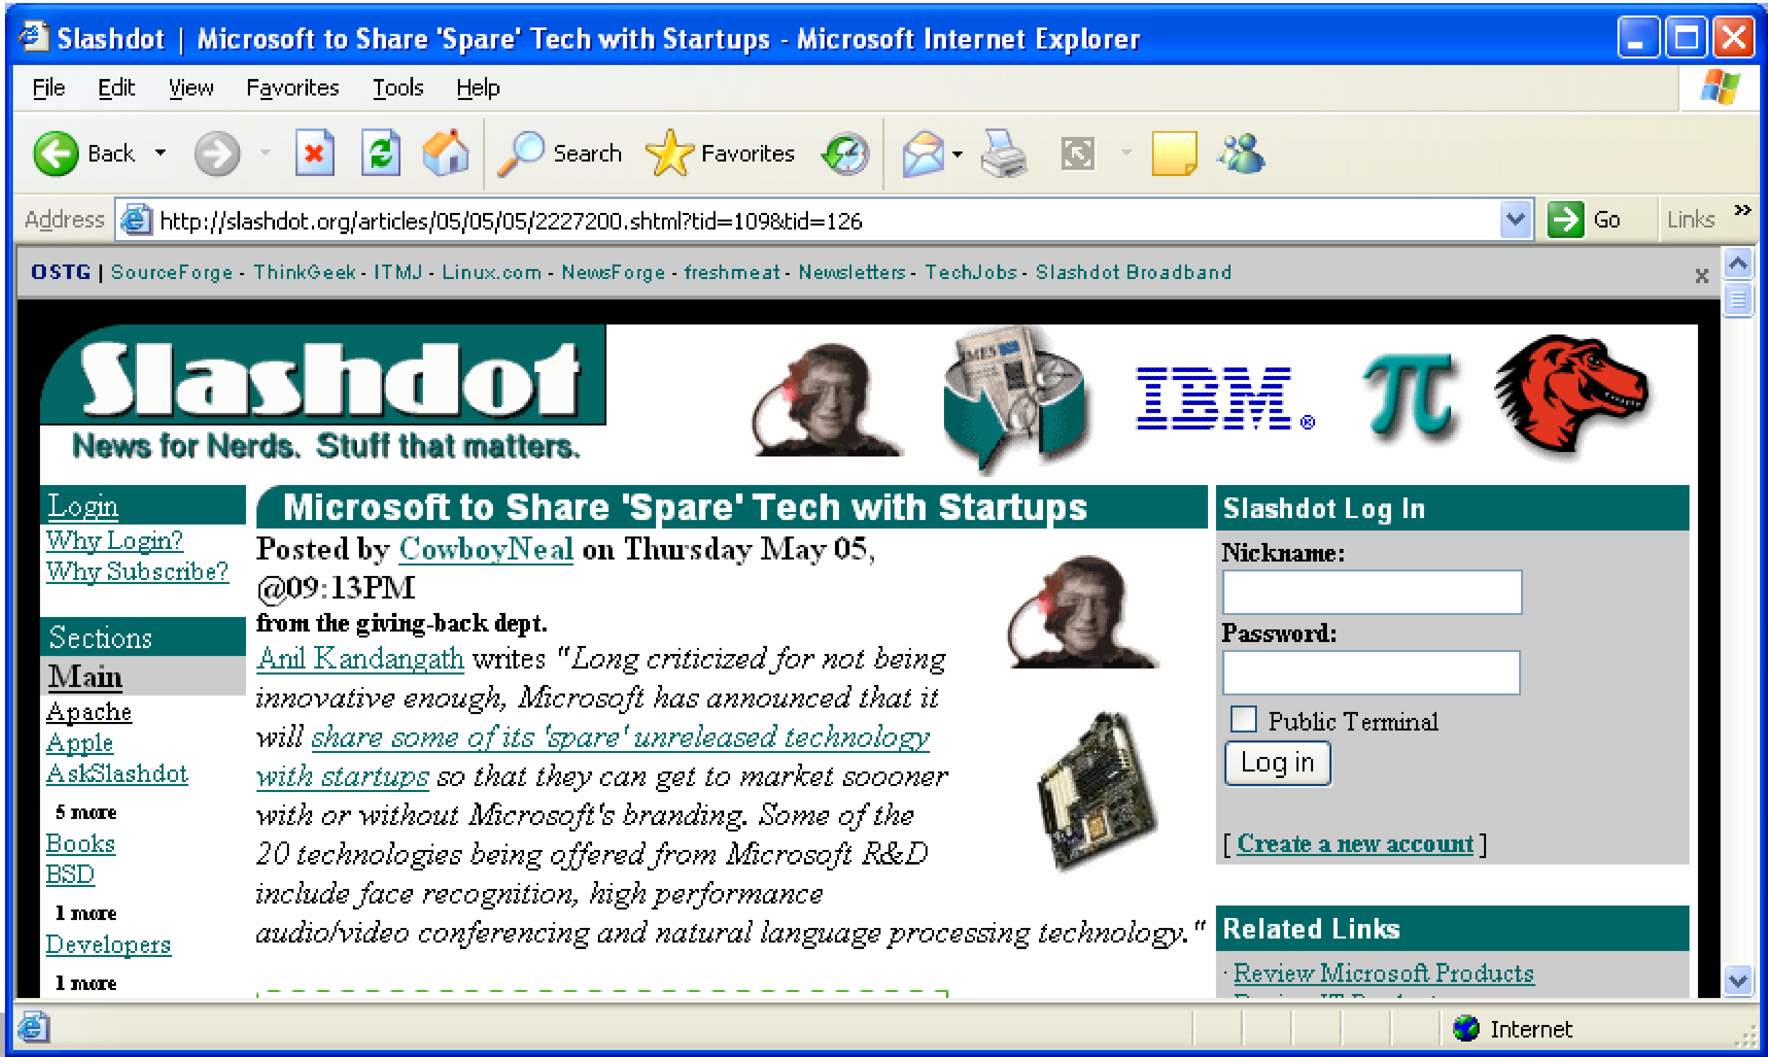

Example Slashdot comments

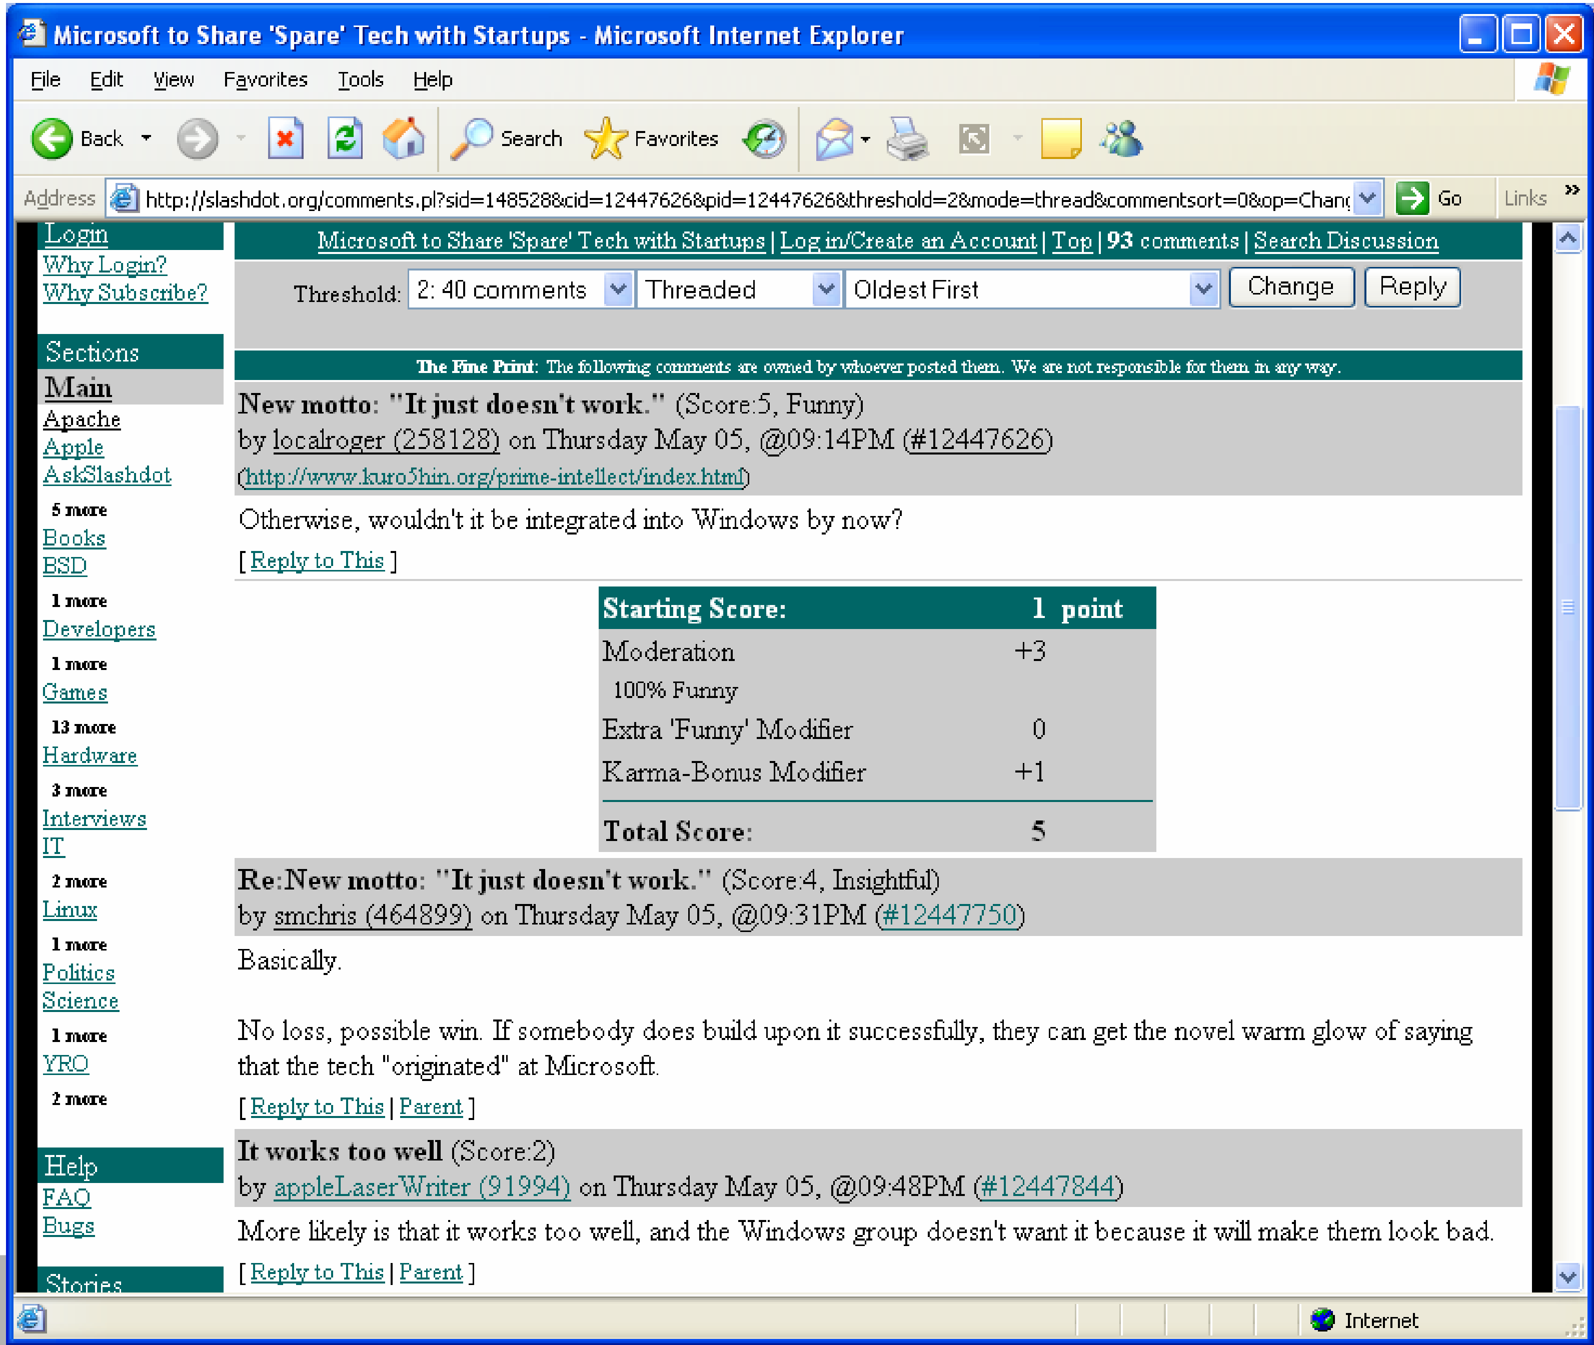

Example Slashdot member

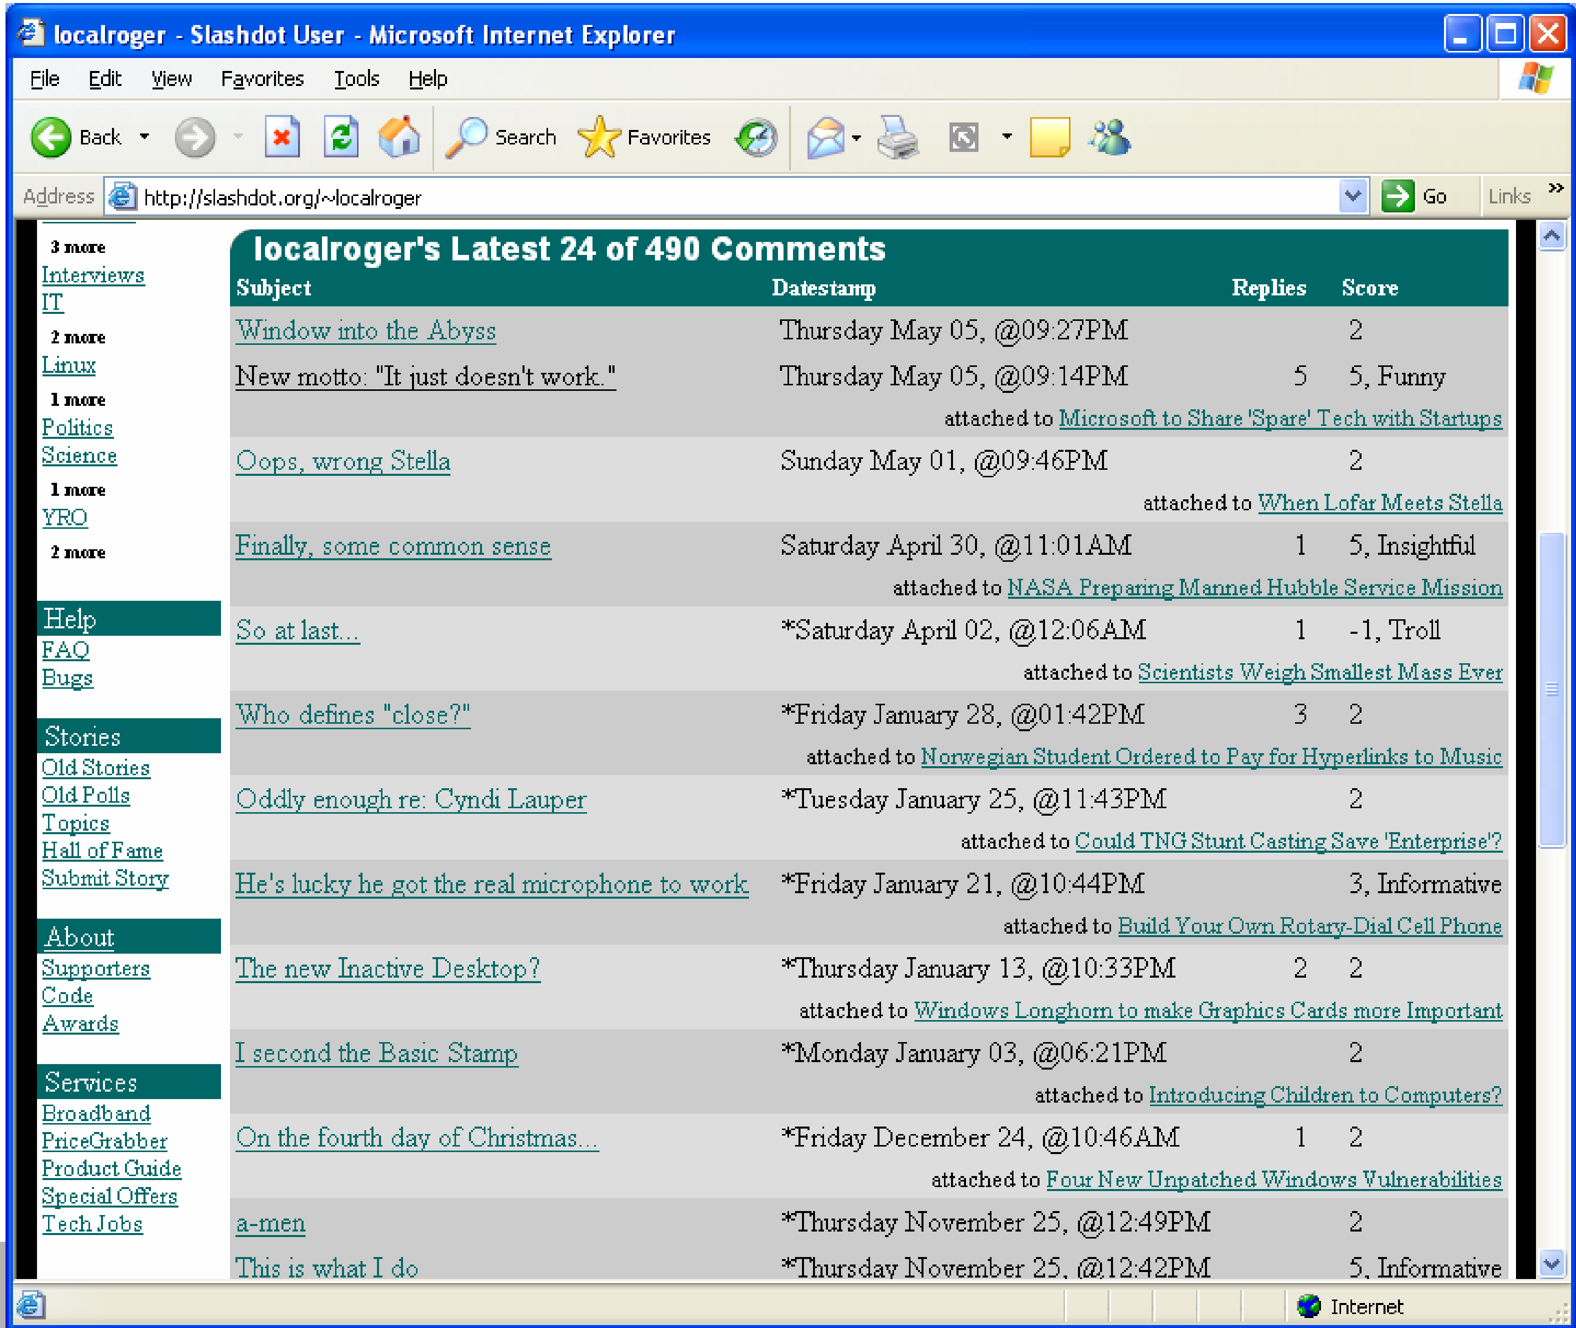

Problems and proposed solutions

Reputation System Challenges

- Ad hoc computation
- Collusion
- Unfair ratings
- Change of identity
- No incentive to provide ratings
- Hard to elicit negative feedback
- Discrimination
- Is past performance = future performance ?

Reputation systems as attack instruments

- Dependence on reputation systems makes it necessary to assess their reliability
- Strategic manipulation of reputation systems can harm the entities through reputation destruction
- Robustness of reputation systems requires hard security
  - Authentication
  - Anonymity
  - Rating tokens
  - etc.

What about subjective taste?

- Recommender systems based on collaborative filtering
  - Assumes different taste
  - Identifies like-minded with same taste
  - Recommender systems
- Reputation System
  - Assumes consistent quality judgement
  - Sanctions poor quality
  - “Collaborative Sanctioning System”

Collaborative filtering

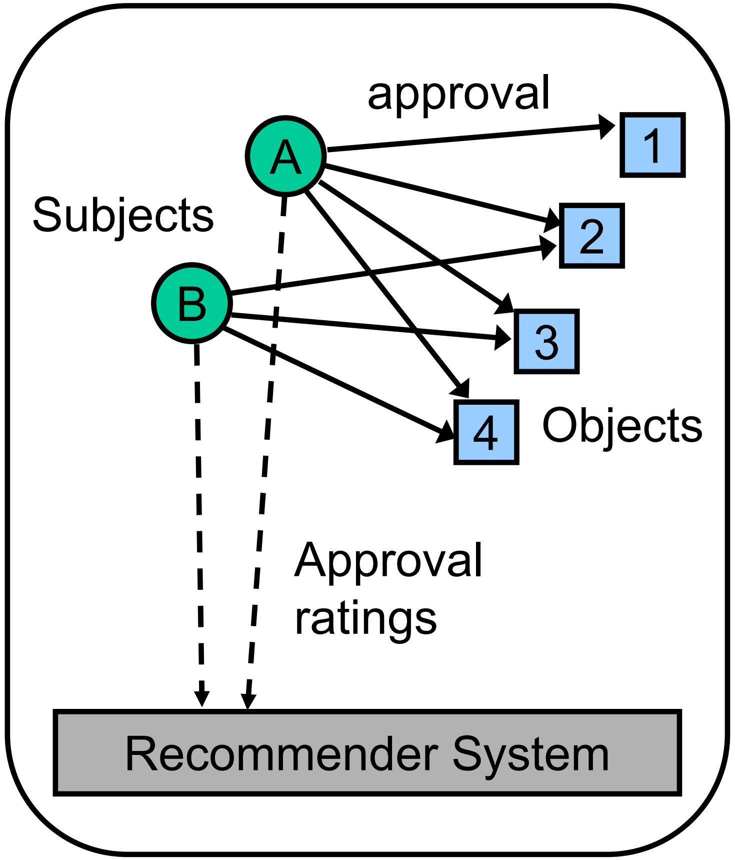

a) Past

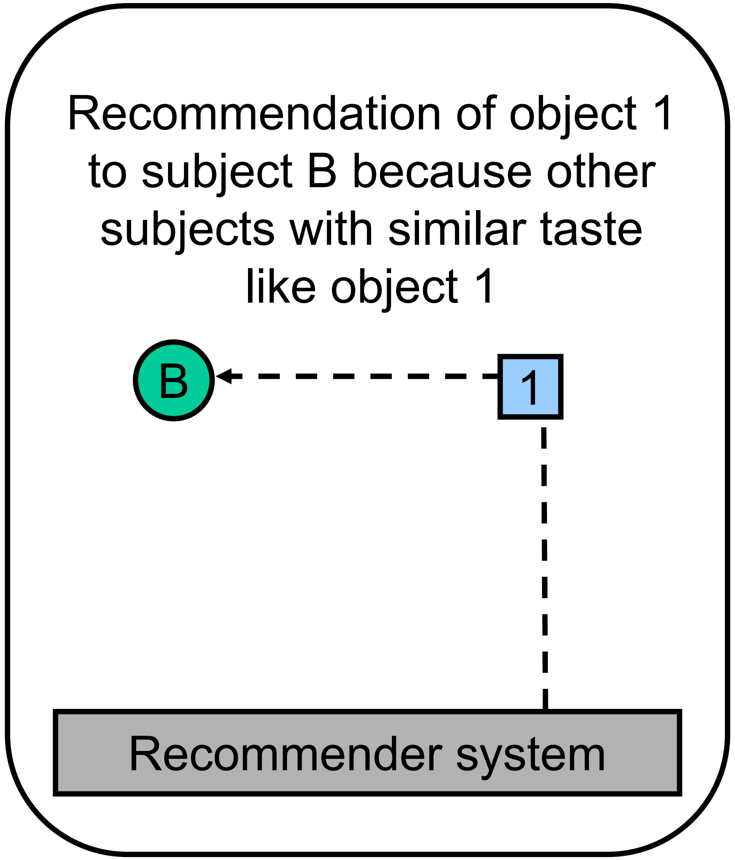

b) Present

Recommender systems in practice

Date: 4 Sep 2006 03:33:48 -0700
From: Amazon.com <store-news@amazon.com>
To: a.josang@qut.edu.au
Subject: Amazon.com recommends The CISSP Prep Guide: Gold Edition and more

Audun Josang, Amazon.com has new recommendations for you based on items you purchased or told us you own.

We recommend: The CISSP Prep Guide: Gold Edition
by Ronald L. Krutz
http://www.amazon.com/o/ASIN/047126802X/ref=pe_ar_x1047126802X

List Price: $80.00
Price: $47.88
You Save: $32.12 (40%)

Recommended because you purchased or rated:
* Official (ISC)2 Guide to the CISSP Exam

Combining recommender and trust systems

- Reputation systems can be used to determine trust relationships
  - Little purpose of getting trust recommendations about Microsoft products from Linux freaks
- Trust recommendation based on collaborative filtering

Collaborative filtering and transitive trust

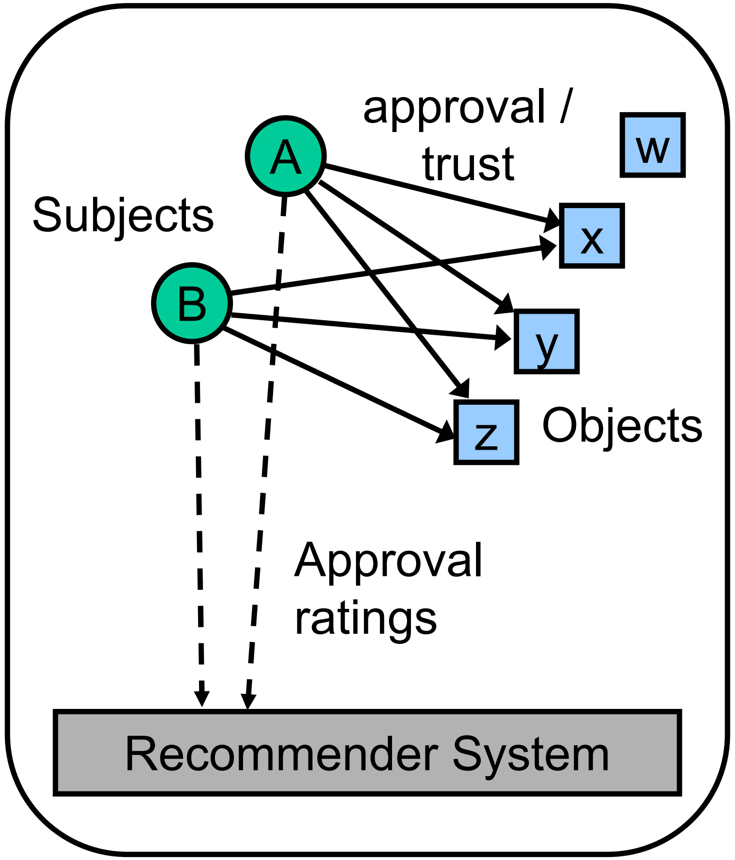

a) Past

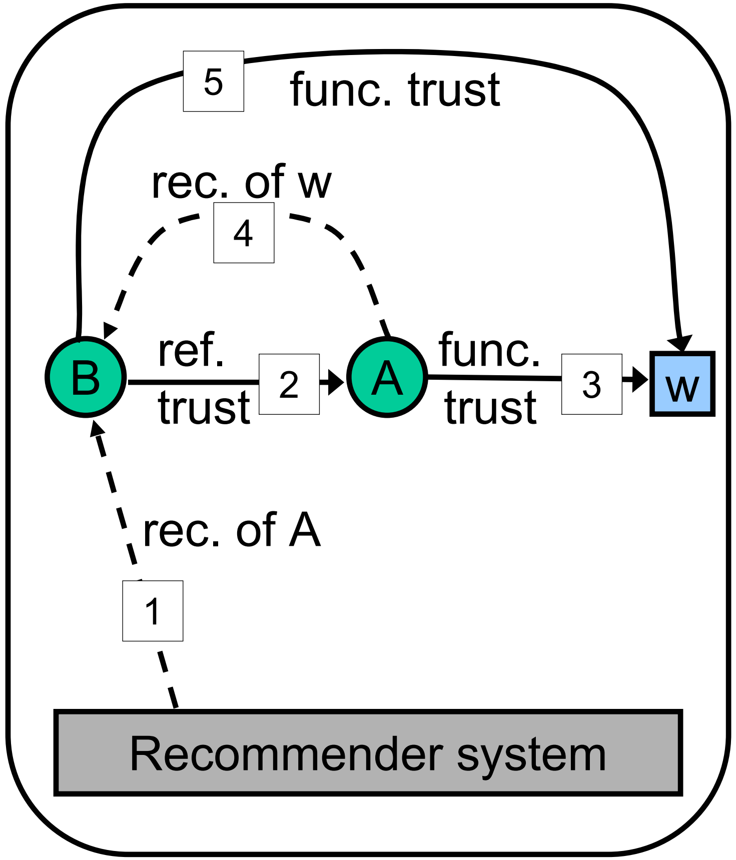

b) Present

Yhprum's Law
(systems that shouldn't work sometimes do)

- People provide ratings despite having no rational incentive to do so.
  - Helps others, but not self
  - Can create competition over scarce resource
- Negative ratings are hard to elicit.
- Relatively easy to mount attacks against existing reputation systems.

- A reputation system works when people can relate to it
- Supports community building

Countermeasures against attacks

- Sound computation engines
- Authentication/security
  - Prevents change of identity
- Statistical filtering, and discounting
  - To prevent unfair ratings, discrimination and collusion
- Anonymity
  - To prevent fear of retaliation
- Benefits / special offers
  - To provide incentive

Concluding remarks 1

- Very primitive commercial systems
  - It is important that users can relate to the systems
  - Community building is an important factor, in addition to enhancing market quality
- Many different proposed theoretic systems
  - Little coherence among researchers
  - Pioneering period
  - No one system is optimal in for all applications

Concluding remarks 2

- Challenging to make systems robust against attacks
  - Limits the potential of reputation systems
  - Requires hard security
- Benefits of trust management
  - Complements traditional security mechanisms
  - Provides incentive for good behaviour
  - Sanctions bad behaviour
  - Increases the quality of online markets and communities

References

- Survey of Trust and Reputation Systems for Online Service Provision
  - http://www.fit.qut.edu.au/~josang/papers/JIB2006-DSS.pdf
- Multinomial Bayesian Reputation Systems (draft)
  - http://www.fit.qut.edu.au/~josang/papers/Jos2007-multirepu.pdf
- Subjective Logic demonstrations
  - http://sky.fit.qut.edu.au/~josang/sl/