The Dynamics of Reputation Systems

[Extended Abstract] *

Amir Ban
Institute for the Study of Rationality
Hebrew University, Givat Ram
Jerusalem, Israel
amirban@netvision.net.il

Nati Linial
School of Computer Science and Engineering
Hebrew University, Givat Ram
Jerusalem, Israel
nati@cs.huji.ac.il

ABSTRACT

Online reputation systems collect, maintain and disseminate reputations as a summary numerical score of past interactions of an establishment with its users. As reputation systems, including web search engines, gain in popularity and become a common method for people to select sought services, a dynamical system unfolds: *Experts'* reputation attracts the potential *customers*. The experts' *expertise* affects the probability of satisfying the customers. This rate of success in turn influences the experts' reputation. We consider here several models where each expert has innate, constant, but unknown level of expertise and a publicly known, dynamically varying, reputation.

The specific model depends on (i) The way that experts' reputation affects customers' preferences, (ii) How experts' reputation is modified as a result of their success/failure in satisfying the customers' requests.

We investigate several such models and elucidate some of the key characteristics of reputation in such a market of experts and customers.

Categories and Subject Descriptors

L.6.1 [**Society/Community**]: Virtual Community

General Terms

Theory

Keywords

Reputation, expert, social learning, search engine, reputation system

1. INTRODUCTION

*A full version of this paper is available at `www.cs.huji.ac.il/~nati/PAPERS/amir_ban_tark_11.pdf`

TARK 2011, July 12-14, 2011, Groningen, The Netherlands.

1.1 How to Find a Good Expert

How do you find a good restaurant to celebrate that special occasion? How to find a lawyer, mechanic or medical specialist when you need one? Where do you find the citation you need in your paper for a subject that is peripheral to your own work? How to find a good movie, a book, a travel destination or a cool YouTube video? Traditional methods of answering these questions are more and more being augmented, or even replaced, by online reputation systems.

According to Friedman et. al.[6], online reputation systems collect, maintain and disseminate *reputations*, aggregated records of past interactions of each participant in a community. While reputation systems have been used to track the trustworthiness of a participant, the property being tracked by a reputation system is more generally a measure of the quality or ability of a participant in a certain class of interactions. Online systems for grading restaurants, travel destinations, movies as well as academic papers and researchers are examples.

We use "expertise" as a catch-all term for the attribute on which a reputation system is designed to report. *Expertise* may be a measure of the quality of services (of a certain kind) rendered, ability to perform a task well, ability to correctly answer questions or predict events of a certain kind, or trustworthiness in interactions. Participants in a reputation system may be divided into *experts*, who provide a service and have an aggregated reputation record, and *users*, who use the expert's services and the system's disseminated reputations. In so-called *peer-to-peer* reputation systems each participant serves in both roles.

We define **expertise** to be the probability that an expert's service will meet a user's requirements, with the probability meant to cover the great variability in this user-expert interaction: variations between the requirements of members of the community, variations of requirements on different occasions, and variations in the skill of the expert according to his[1] own circumstances. We make the assumption that this expertise probability is an objective and fixed attribute of an expert, which is however hidden and can only be estimated by trials. Hence the value of the reputation system in providing a community of users information that is helpful in estimating expertise.

We treat expertise as an innate ability of an expert, not a strategic choice. We assume that an expert cannot perform better than his fixed expertise and has no cost in doing his best (and therefore has no motive in performing worse).

[1]Throughout this paper we refer to experts as masculine and to members of the public as feminine.

1.2 Reputation

The reputation provided by a reputation system may consist of the aggregated record of reported past interactions. More commonly it is a numerical score derived from that record, enabling experts rated by the system to be ranked. The reputations disseminated by a reputation system help users estimate the objective expertise (defined as probability for success) of each expert. Often, it is enough if the reputations assist users in *ranking* the estimated expertise levels, since a user's strategy typically consists of selecting the best available expert.

A reputation system may therefore be called effective if, given the way past interactions are reported and aggregated, a higher reputation score more often than not reflects a higher expertise. Consequently the effectiveness of a reputation system depends not only on its design, but on its sources of information as well.

When all past interactions, successful or not, are recorded by a reputation system, designing it for effectiveness is simple: The success rate of each expert (the ratio of successes to trials) is an unbiased estimator of expertise and therefore may be used as the reputation score. Users may make the natural and simple choice of the highest-score expert, or may engage in longer-range and ultimately superior k-armed bandit strategies[7].

However, full reporting, or even representative reporting of interactions is not realistic in most real-world situations. Often, the motivations of users who provide reports to a reputation system lead to one-sided reporting, of successes only, or, less commonly, of failures only. Some systems, such as search engines (where the "endorsements" consist of links), or academic importance (where the "endorsements" consist of citations), simply have no scope for negative mention. Statistics made of sales figures, such as for books or movies, also reflect only positive choices, while a customer's decision not to buy is a silent, unrecorded act. Nor is the problem solely a question of the reputation system's design: Users seem to be inclined to report positive interactions, while generally glossing over negative ones. For example, a system of "thumbs up//down" introduced to rate YouTube videos shows that "thumbs up" occurrences consistently outnumber "thumbs down" occurrences so heavily that it is doubtful that their relative proportion measures viewer approval/disapproval.

1.3 Expertise vs. Reputation

We aim to investigate the relation between expertise and reputation in a reputation system: Does reputation, the public's perception of expertise, reflect real expertise in the long run? Can we be assured that, between several experts of varying skill, eventually the expert of highest expertise will have the highest reputation? Or, alternatively, might reputation be self-perpetuating, with a high reputation merely reflecting a favorable head start?

Both possibilities are plausible. On the one hand, an expert's reputation is reinforced by a customer's positive experience, which becomes more likely the higher his expertise. On the other hand, the number of his customers depends on his current reputation. Both expertise and current reputation therefore contribute positively to future reputation and either may conceivably dominate in the long run.

The expertise-reputation contest is well-known to the corporate world: Entry into an established market is generally difficult and costly. Having an excellent product may not suffice, and the newcomer may need to spend a lot to close his reputation gap: For example, in the mid-1990's Netscape ruled the web browser world, until Microsoft's Internet Explorer managed (with considerate effort) to sideline it. On the other hand, the success of Mozilla's Firefox, a non-profit open source browser, teaches that entry against an entrenched leader is possible, and on merit alone.

Nobody in their right mind expects a better-tasting but no-name cola drink to supplant Coca-Cola and Pepsi-Cola. Nor is it commonly believed that the unrivaled supremacy of these mega brands rests on the unrivaled quality of their soft drinks.

The situation is well-known in the cultural world, where being "in vogue" is in large part self-sustaining: The people who flock to see a Van Gogh exhibition, a Rolling Stones concert or a performance of Verdi's Aida seem to be driven at least in part by the respective artists being acknowledged as "all-time greats". Indeed Van Gogh's wretched career during his own lifetime indicates that there is nothing inevitable about his posthumous fame.

In arts, music and literature there is still an intangible but definite "expertise", but when considering the "reputation" of TV personalities, movie stars, supermodels and so on, it becomes less and less clear what it is that sustains them in their elevated position in the face of the hordes of wannabes who would love to take their place and seem just as qualified. This has led a cynic to quip "a celebrity is someone who is famous for being well-known".

1.4 Search Engines as Reputation Systems

Web search engines, and in particularly the ubiquitous Google, play the part of universal "managers" of reputation. The Google search engine is easily the world's most popular reputation system. Google famously employs the Page-Rank algorithm [3] to rank the importance of pages. Briefly, the importance of a page is the sum of the importance of each page that hyperlinks to it, plus a constant self-importance.

In the full version of this paper we demonstrate a close relationship between page-rank values and the value of reputation as defined in our model.

Google and the Page-Rank algorithm exemplify well the interaction of "expertise" and "reputation": In response to a search query, Google ranks pages in order of their "reputation" (in fact, their page rank), which is indicative of their "expertise" (in fact, their likelihood to be what was searched for). A page found in a search is likely to acquire new links (for example, when looking for a travel destination, which will later be mentioned in the traveler's blog). It is tacitly assumed (by Google and its users) that this procedure ultimately causes the objectively best pages to be ranked high. Whether or not this is indeed the likely outcome is the subject of our investigation.

1.5 Main Results

Analyzing how the reputation ranking of experts evolves given their expertise and initial ranking, we show that being ahead in reputation confers a quantifiable advantage that may balance inferiority in expertise ("No. 2 Tries Harder"). In a system where reputation is enhanced positively, a steady-state order is always reached, and is unique given initial conditions. On the other hand, when negative feedbacks are included, reputation orders become chaotic. We also prove that reputation is indeed a positive signal for expertise in (almost) all conceivable reputation systems.

1.6 Reputation in Economics and Game Theory

The subject of reputation is extensively discussed in the literature of game theory. It was introduced by Selten in the "chain-store paradox" [12] to mean the belief of players in games that another player takes actions that fall within a certain class, e.g. "aggressive". Kreps and Wilson [9] showed how reputation may affect behavior in Bayesian games where there is uncertainty about players' payoff structure. Reputation is usually used in order to capture strategic choices that a player selects at will (such as honesty or aggresiveness), rather than intrinsic attributes, such as quality of service or expertise in a field, which a player cannot choose at will.

The concept of brand as a carrier of a firm's reputation was put forward by Kreps [8], in the context of moral hazard. Cabral [4] discusses firm reputation as a posterior belief of its customers of the firm's quality level given the firm's history of performance, in the context of whether a firm with a strong brand is well-advised to use the same brand for a new product.

The term *reputation system* was originally used for online systems whose participants grade each other's trustworthiness, or competence, of which the one used by eBay is the archetype. eBay's reputation system has been extensively studied, e.g. by Dellarocas[5].

Tadelis[13] as well as Mailath and Samuelson[10] consider reputation as a tradable asset: A firm's reputation is a noisy signal of its competence or effort observable by customers. Firms may trade in reputations, and such trades are only partially observable by customers.

Information cascades [2] study situations in which it is optimal for an individual, having observed the actions of those ahead of her, to follow the behavior of the preceding individual without regard to her own information. Information cascades have been advanced as an explanation of the localized conformity of behavior and the fragility of mass behavior. Information cascades have some features in common with reputation systems, but typically their models lack user feedback which is key to the formation of reputation.

1.7 Organization of this paper

Section 2 presents our model.

In Section 3 we analyze in detail the behavior of the reputation-expertise model and prove most of our main results.

In Section 4 we discuss the rationality of the behavior outlined in our model, and prove that it is rational, under very broad assumptions.

In Section 5 we summarize and discuss future work.

In the full version of this paper the interested reader will find a demonstration of the relation between our model's notion of reputation and the measures used by web search engines to rank search results, proofs of the theorems and lemmas, as well as supplementary details on the behavior of the model.

2. THE MODEL

We use a model wherein N users repeatedly seek the services of n experts, each expert having a publicly known reputation. The reputations evolve dynamically and represent the aggregate feedback of the users based on their satisfaction with the services they received. The probability of providing satisfactory service to any user request is modeled as the expert's expertise, a fixed but hidden quantity characterizing each expert.

When seeking service, users engage experts according to some selection order, stopping when they are satisfied with the service (or when they have exhausted all experts). The selection order may be simply the experts ordered by descending reputation (the "Reputation Ordered Scheme"), or may also take into account a user's previous experience with experts ("Loyalty Schemes").

2.1 Rounds

Time is discrete: At each integral time a **round** takes place, in each of which each user seeks a service that may be provided by any of the experts. The service provided by an expert may either succeed or fail. Each expert has his own **reputation** and **expertise**. The reputation of expert i is a real-valued function of the discrete time, $r_i = r_i(t)$ representing accumulated user feedback (at round t) on expert i's success rate. The expertise of expert i is his actual success probability $\epsilon_i \in [0, 1]$ in satisfying users requests. At any given round t, the current reputation values $r_1(t), \ldots, r_n(t)$ are common knowledge to all users. On the other hand, the expertise values $\epsilon_1, \ldots, \epsilon_n$ are unknown to the experts and the users.

2.2 Success and Failure Probabilities

The outcome of expert i's service is a random Bernoulli event with probability ϵ_i for success and probability $1 - \epsilon_i$ for failure. This outcome is independent of any other user-expert interaction in any round. However, a repeat service request to an expert in the same round produces the same result as the first request. Therefore there is no point in seeking the service of the same expert more than once in a round.

2.3 Selection Order

The order in which experts are asked in each round is called the selection order. Users follow their selection order until success, or until no more experts are available. The selection order may be based on the reputations publicly known at the start of the round, or on the user's previous experience, or both. A user has no direct information of other users' experience. The selection order may be non-deterministic.

A selection scheme that depends only on the order of experts' reputations (and not, e.g., on actual reputation values) is called **order-based**.

A user that queries an expert in some particular round, is called a **customer** of the expert in that round.

2.4 Reputation Update Rule

At the end of each round, each expert's reputation is updated according to his customers' experience in that round. For every successful service (i.e., for every satisfied customer) in the current round, the expert's reputation is incremented by β, and for every failed service (i.e., a dissatisfied customer) in the current round, it is decremented by $1 - \beta$. The total of all reputation updates for an expert in a round is called the expert's **feedback**. The parameter $0 \leq \beta \leq 1$ is called the **reward/penalty factor**. Note that for $\beta = 1$,

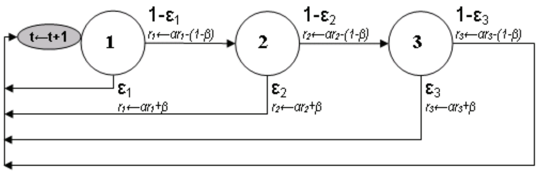

Figure 1: Flow, timing and feedback in the reputation-order scheme

only successes are rewarded, while for $\beta = 0$, only failures are penalized.

At the beginning of each round, the previous round's reputation is multiplied by a persistence (discount) factor $0 \leq \alpha \leq 1$. No discounting corresponds to $\alpha = 1$.

We denote the number of customers of expert i in round t by $c_i(t)$, and his feedback in round t by $w_i(t)$. Let $S_i(t) \subset [N], U_i(t) \subset [N]$ be i's set of satisfied and unsatisfied customers, respectively. Then $|S_i(t) \cup U_i(t)| = c_i(t)$, and the full reputation update rule is:

$$w_i(t) = |S_i(t)|\beta - |U_i(t)|(1-\beta) \tag{2.1}$$

$$r_i(t+1) = \alpha r_i(t) + w_i(t) \tag{2.2}$$

It is easy to express the expectation of expert i's feedback in round t. If he has $c_i(t)$ customers at that round, then

$$\mathbb{E}[w_i(t)] = \beta\epsilon_i \, \mathbb{E}[c_i(t)] - (1-\beta)(1-\epsilon_i)\, \mathbb{E}[c_i(t)] \Rightarrow \tag{2.3}$$
$$\mathbb{E}[w_i(t)] = (\epsilon_i + \beta - 1)\, \mathbb{E}[c_i(t)] \tag{2.4}$$

It follows that the expected feedback of an expert is positive if and only if his expertise ϵ is $\geq 1-\beta$, regardless of the selection order or any other detail.

2.5 Selection Order as a Markov Chain

Figure 1 demonstrates the flow of a user in a round in the reputation-ordered scheme with 3 experts, shown as a Markov chain with feedback side-effects.

The 3 experts are ranked by reputation so that $r_1 > r_2 > r_3$. The diagram represents the flow so long as this order is stable (an order change, if it happens, is noted only at the start of a round, at the $t \leftarrow t+1$ oval). Each of the three states (circles) represents a possible query of one of the experts, with two emanating edges representing failure (emanating right), and success (emanating downwards), with their respective feedback updates.

3. BEHAVIOR OF THE MODEL

We explore the behavior of the model under various selection schemes. The main question that will interest us is whether the rank order of expert reputations reaches a steady state in the long run, and if so, to what extent the reputation ranking in the long run reflects the experts' objective levels of expertise $\epsilon_1, \ldots, \epsilon_n$, vs. the initial order of reputations $r_1(0) > \ldots > r_n(0)$.

Since the model is stochastic in nature, the notion of *steady state* should be clarified. To this end we define the notion of *quasi-stability*:

3.1 Quasi-Stability

DEFINITION 1. *A **setting** is a combination of a reputation ranking $r_1(t) > \ldots > r_n(t)$, of experts with expertise levels $\epsilon_1, \ldots, \epsilon_n$ and a selection scheme $\mathcal{S}$ (with reward-penalty factor β).*

*The i'th **pair** within a setting consists of a **leader**, the expert at index i, and the pair's **follower** - the expert with index $i+1$.*

*The i'th pair within a setting is said to be **quasi-stable** (at round t) if the expected reputation of the pair's leader given the setting (as defined in (2.1)) is greater or equal to the expected reputation of the pair's follower, at all rounds starting at t, i.e.:*

$$\mathbb{E}[r_i(T)] \geq \mathbb{E}[r_{i+1}(T)], \qquad \forall T \geq t \tag{3.1}$$

*where the probability space for the expectation is taken over all possible user-expert interactions in rounds after t, **assuming that the expert order does not change**.*

*If a pair is not quasi-stable, it is called **unstable**.*

For order-based schemes, feedback expectations stay constant so long as the reputation order does not change. Therefore in view of (2.2) the definition of quasi-stability is equivalent to:

$$\mathbb{E}[w_i(t)] \geq \mathbb{E}[w_{i+1}(t)] \tag{3.2}$$

Quasi-stability encompasses the notion of stochastic stability of the order between the leader and the follower in a pair: The leader is expected to retain his position in subsequent rounds, perhaps indefinitely. Since the feedback is a random variable, this is not guaranteed: An unlikely lucky streak may close the gap and put the follower in the lead. In the context of our model, a "lucky streak" for an expert would be to have a higher rate of success than his true level of expertise. The probability for this happening vanishes as the number of users N grow without limit, as by Law of Large Numbers the (relative) variance in reputation feedback per round vanishes.

We use the term WITH HIGH PROBABILITY or w.h.p. for short to describe an attribute that, in the scope of a given time frame, is expected to be true with probability $1 - o(1)$ as $N \to \infty$. Thus an equivalent to saying that a pair is quasi-stable is the statement that it is w.h.p. stable, or that the leader's reputation is w.h.p. greater than the follower's.

In contrast, an unstable pair is guaranteed to experience an order change between leader and follower, in a time frame that does not depend on the number of users.

We use the following additional definitions on stability:

DEFINITION 2. *A setting is called **quasi-stable** if all pairs within it are quasi-stable.*

*A setting that is not quasi-stable, i.e. if any pair within it is unstable, is called **unstable**.*

*A pair within a setting is called **two-sided quasi-stable** if it is quasi-stable and if it would be quasi-stable in a setting wherein the pair's leader and follower have traded places.*

*A pair within a setting is called **two-sided unstable** or **chaotic** if it is unstable and if it would be unstable in a setting wherein the pair's leader and follower have traded places.*

Two-sided quasi-stability describes a situation wherein whoever leads retains the lead. Two-sided instability describes a situation wherein no lead is stable.

3.2 The Reputation-Ordered Scheme

In the reputation ordered scheme all users select experts in descending order of their current reputations.

The scheme is "memoryless": Users have no recall of what happened in previous rounds. Therefore it is irrelevant whether they are the same users each round: In memoryless schemes, there is no significance to users' identity, and all users employ the same selection order for experts.

Let us analyze the dynamics of this scheme. Assume the experts are indexed by their reputation order (at round t) $r_1(t) > r_2(t) > \ldots > r_n(t)$ and have respective expertise $\epsilon_1, \ldots, \epsilon_n$. All N users will be the 1st expert's customers, $c_1(t) = N$.

By expectation, $\epsilon_1 c_1(t)$ of the customers will be satisfied, and $(1-\epsilon_1)c_1(t)$ will be dissatisfied. 1st expert's dissatisfied customers will become 2nd expert's customers. In general, i'th expert's dissatisfied customers will become $(i+1)$'th expert's customers. Formally (and marking by $c_{n+1}(t)$ the number of users who failed with **all** experts):

$$\mathbb{E}[c_{i+1}(t)] = (1-\epsilon_i)\,\mathbb{E}[c_i(t)], \qquad \forall i \in [n] \tag{3.3}$$

Therefore:

$$\mathbb{E}[c_i(t)] = N\prod_{j=1}^{i-1}(1-\epsilon_j), \qquad \forall i \in [n+1] \tag{3.4}$$

Under what conditions will the reputation order be quasi-stable?

Theorem 1. *Let n experts be indexed by their reputation order (at round t) $r_1(t) > r_2(t) > \ldots > r_n(t)$ and have respective expertise $\epsilon_1, \ldots, \epsilon_n$. Let n_1 be the smallest index for which $\epsilon_{n_1} = 1$, or let $n_1 = n$ if no such index exists. Then the order is quasi-stable under the* ***reputation-ordered scheme*** *if and only if the following equivalent inequalities apply for each $i \in [n_1 - 1]$:*

1. leader's advantage: $\frac{\beta}{1-\epsilon_i}\epsilon_i \geq \epsilon_{i+1}$
2. follower's handicap: $\epsilon_i \geq \frac{1}{\beta+\epsilon_{i+1}}\epsilon_{i+1}$
3. reciprocal difference: $\frac{1}{\epsilon_i} - \frac{\beta}{\epsilon_{i+1}} \leq 1$

Proof. The expected feedback of each expert (see (2.4)) is:

$$\mathbb{E}[w_i(t)] = (\epsilon_i + \beta - 1)\,\mathbb{E}[c_i(t)], \qquad \forall i \in [n] \tag{3.5}$$

Since this scheme is order-based, it is quasi-stable if $\forall i \in [n-1]$ $\mathbb{E}[w_i(t)] \geq \mathbb{E}[w_{i+1}(t)]$. If $i \geq n_1$ the inequality holds trivially, as both expectations are zero. Otherwise the condition translates to:

$$(\epsilon_i + \beta - 1)\,\mathbb{E}[c_i(t)] \geq (\epsilon_{i+1} + \beta - 1)\,\mathbb{E}[c_{i+1}(t)] \tag{3.6}$$

By (3.3):

$$(\epsilon_i + \beta - 1)\,\mathbb{E}[c_i(t)] \geq (\epsilon_{i+1} + \beta - 1)(1-\epsilon_i)\,\mathbb{E}[c_i(t)] \tag{3.7}$$

Dividing both sides by $\mathbb{E}[c_i(t)]$ (a positive number as $i < n_1$):

$$\epsilon_i + \beta - 1 \geq (\epsilon_{i+1} + \beta - 1)(1-\epsilon_i) \tag{3.8}$$

Which, by rearrangement, leads to each of the three equivalent inequalities. □

Corollary 1. *Under the reputation-ordered scheme, the quasi-stability of a pair within a setting depends only on the expertise of the pair's leader and follower.*

Under the reputation-ordered scheme, rank certainly has its privileges: Considering the situation where experts get positive feedback only ($\beta = 1$), Corollary 2 shows that being ahead in reputation confers on an expert with expertise ϵ an advantage factor of $1/(1-\epsilon)$ in expertise over followers:

Corollary 2. *Under the conditions of Theorem 1 with a reward-only scheme ($\beta = 1$), quasi-stability requires the following equivalent inequalities for each $i \in [n_1 - 1]$:*

1. leader's advantage: $\frac{1}{1-\epsilon_i}\epsilon_i \geq \epsilon_{i+1}$
2. follower's handicap: $\epsilon_i \geq \frac{1}{1+\epsilon_{i+1}}\epsilon_{i+1}$
3. reciprocal difference: $\frac{1}{\epsilon_i} - \frac{1}{\epsilon_{i+1}} \leq 1$

It is worth noting that this advantage becomes insurmountable when $\epsilon > 1/2$, for in this case $\frac{1}{1-\epsilon}\epsilon > 1$, and therefore the pair is quasi-stable against any follower.

A different perspective on this is to consider the disadvantage of being second: Corollary 2 shows that this inflicts a handicap factor of $1/(1+\epsilon)$ on an expert with expertise ϵ. In other words, an expert can reasonably expect to overtake a leader over whom his advantage in expertise is greater than his handicap. Again, it is worth noting that $\frac{1}{1+\epsilon}\epsilon \leq 1/2$. That is, a follower can never expect to overtake a leader with expertise of more than 1/2.

A pair whose leader and follower have the same expertise is always two-sided quasi-stable, i.e. between equals, the order determined by initial conditions is preserved. Generally, by Corollary 2 (3), two-sided quasi-stability exists iff:

$$-1 \leq \frac{1}{\epsilon_i} - \frac{1}{\epsilon_{i+1}} \leq 1 \tag{3.9}$$

On the other hand, when negative feedback applies ($\beta = 0$), chaos reigns, as spelled out by Corollary 3:

Corollary 3. *Under the conditions of Theorem 1 with a penalty-only scheme ($\beta = 0$), quasi-stability is possible only if $\epsilon_1 = 1$ or $\epsilon_i = 0$ for all but the first expert. (Substituting $\beta = 0$ in Theorem 1 (3) we derive either $\epsilon_{i+1} = 0$ or $\epsilon_i \geq 1$ which is impossible as we assumed $i < n_1$).*

As negative feedback is gradually added (i.e. as β gradually decreases below 1), the above situation changes in several respects. By Theorem 1:

- The leader's advantage decreases from $1/(1-\epsilon)$ to $\beta/(1-\epsilon)$, disappearing (i.e. equaling 1) at the critical point $\beta = 1-\epsilon$.
- A leader with expertise greater than $1/(1+\beta)$ has an unassailable position.
- The follower's handicap decreases from $1+\epsilon$ to $\beta+\epsilon$, disappearing (i.e. equaling 1) at the critical point $\beta = 1-\epsilon$.
- Note (see (2.4)) that experts with expertise above the critical $1-\beta$ have positive feedback expectation, while experts with expertise below the critical value have negative feedback expectation.

- A pair with leader expertise of less than $1-\beta$ and follower expertise of at least $1-\beta$ is always unstable, while in reverse order it is always quasi-stable.
- Two-sided quasi-stability is possible only if both follower and leader have expertise greater or equal to $1-\beta$.

Corollary 4. *Under the conditions of Theorem 1, two-sided instability, i.e. chaos, is possible only if an expert pair exists such that $\epsilon_i < 1-\beta$ as well as $\epsilon_{i+1} < 1-\beta$.*

Proof. By Theorem 1 two-sided instability requires both of the following inequalities to be true:

$$\frac{1}{\epsilon_i} - \frac{\beta}{\epsilon_{i+1}} > 1 \tag{3.10}$$

$$\frac{1}{\epsilon_{i+1}} - \frac{\beta}{\epsilon_i} > 1 \tag{3.11}$$

Multiplying both sides of (3.11) by β and adding (3.10) results in:

$$\frac{1}{\epsilon_i} - \frac{\beta}{\epsilon_{i+1}} + \frac{\beta}{\epsilon_{i+1}} - \frac{\beta^2}{\epsilon_i} > 1+\beta$$

$$\Rightarrow \quad \frac{1}{\epsilon_i}(1-\beta^2) > 1+\beta$$

$$\Rightarrow \quad \frac{1}{\epsilon_i}(1-\beta) > 1$$

$$\Rightarrow \quad 1-\beta > \epsilon_i$$

Similarly, multiplying both sides of (3.10) by β and adding (3.11) leads to $1-\beta > \epsilon_{i+1}$. □

Note, though, that Corollary 4 states a necessary, but not a sufficient condition for two-sided instability. For example, $\beta = 0.5, \epsilon_1 = 0.4, \epsilon_2 = 0.1$ is quasi-stable, and becomes chaotic only for $\beta < 0.15$.

3.3 Steady State Orders

3.3.1 Definitions

Previously in this section we defined the notion of quasi-stability and instability of a setting, and formulated criteria for it in reputation-ordered settings. Assuming unstable pairs will eventually flip their leader-follower order, and assuming quasi-stable pairs to be stable[2], we ask:

1. Will the setting converge to a steady state, i.e. to a quasi-stable setting?
2. If so, given the setting, what will the steady state setting be?
3. Given an initial setting, is the steady state unique?

We will answer these questions for the reputation-ordered scheme and for other schemes, but first we need some definitions:

Definition 3. *A selection scheme is called **regular** if in settings that employ it the quasi-stability of a pair depends only on the expertise values of the leader and of the follower.*

By Corollary 1 the reputation-ordered scheme is regular.

Definition 4. *The notation $\epsilon_1 \prec_s \epsilon_2$ means that a pair with leader expertise ϵ_1 and with follower expertise ϵ_2 is unstable under the regular selection scheme $\mathcal{S}$.*

The notation $\epsilon_1 \not\prec_s \epsilon_2$ means that a pair with leader expertise ϵ_1 and with follower expertise ϵ_2 is quasi-stable under the regular selection scheme $\mathcal{S}$.

*A setting is called **chaotic** with respect to selection scheme $\mathcal{S}$ if it includes a **chaotic pair**: a pair of experts with expertise ϵ_1, ϵ_2 such that $\epsilon_1 \prec_s \epsilon_2$ and $\epsilon_2 \prec_s \epsilon_1$.*

The instability operator is transitive, i.e. $\epsilon_1 \prec_s \epsilon_2$ and $\epsilon_2 \prec_s \epsilon_3 \Rightarrow \epsilon_1 \prec_s \epsilon_3$. This follows from the definition of instability (Definition 1).

Recall that in a regular setting the quasi-stability or instability of an expert pair depends only on their respective expertise levels. Therefore, in a non-chaotic setting, the instability operator is a *partial order* on the levels of expertise.

3.3.2 Existence and Uniqueness

Starting at some initial setting in which the experts are arranged by descending reputation and numbered from 1 to n, $r_1(0) > \ldots > r_n(0)$ with respective expertise $\epsilon_1, \ldots \epsilon_n$, the initial order is described by the permutation $(1, 2, \ldots, n)$ of [n]. If this setting is not quasi-stable, then at some future round, two neighboring experts in an unstable pair will trade places. E.g. if the leader and follower at pair i trade places, the resulting permutation is $(1, \ldots, i-1, i+1, i, i+2, \ldots, n)$.

Let us mark the expert order at round t by the permutation $\pi(t) = (\pi_1(t), \pi_2(t), \ldots, \pi_n(t))$. $\pi_i(t)$ is the expert at the i'th position at round t.

The meaning of reaching a steady state is that there exists a round T with a quasi-stable setting, i.e.: $r_{\pi_1(T)}(T) > \ldots > r_{\pi_n(T)}(T)$ and $\epsilon_{\pi_1(T)} \not\prec_s \ldots \not\prec_s \epsilon_{\pi_n(T)}$.

The path from the initial setting to the steady-state setting consists of successive swapping of the leader and follower in unstable pairs, until no more such swapping is possible.

We now derive a general result in the theory of partially ordered sets (=posets). This result will be applied below to the instability operator, on the way to answering our questions regarding steady states.

Let $(P, \prec)$ be a finite poset. If $x, y \in P$ and either $x \prec y$ or $y \prec x$ holds, we say that x, y are **comparable**. Otherwise we say they are incomparable and write $x \parallel y$. Let $\pi = (x_1, \ldots, x_n)$ be an ordering of P's elements. A **swap** changes this permutation to $(x_1, \ldots, x_{i-1}, x_{i+1}, x_i, x_{i+2}, \ldots, x_n)$ for some index i. This swap is **permissible** if $x_i \prec x_{i+1}$. We say that a permutation σ of P's elements is *reachable* from π if it is possible to move from π to σ through a sequence of permissible swaps. A permutation of P's elements is called *terminal* if no swap is permissible. It is an easy observation that starting from any permutation of P, any series of permissible swaps is finite, since every two elements can be swapped at most once.

Theorem 2. *For every permutation π of the elements of a finite poset $(P, \prec)$ there is exactly one terminal permutation reachable from π.*

Proof. See the full version of this paper. □

Armed with this general result on posets, we derive a general theorem regarding the existence and uniqueness of steady-state reputation orders:

[2]Meaning that we neglect the $o(1)$ (as $N \to \infty$) probability for their experiencing an order change

THEOREM 3. *Given a setting with a regular selection scheme:*

1. *If the setting is not chaotic it will converge to a quasi-stable setting.*
2. *If the setting converges to a quasi-stable setting, it will w.h.p. converge to the same setting.*

PROOF. The theorem is an immediate consequence of Theorem 2 by noting:

1. The instability relation under a non-chaotic, regular setting defines a partial order on the levels of expertise.
2. Order changes in a setting are w.h.p. between some unstable expert pair.
3. In a steady-state order all expert pairs are quasi-stable, i.e. none are unstable.

□

As consequence of Theorem 3 there exists a simple algorithm to determine the steady-state order arising out of any given setting: Calculate the quasi-stability or instability of all pairs in a setting. Switch the order of any unstable pair. Repeat until reaching a setting with no unstable pair.

EXAMPLE 1. *In the reward-only reputation- ordered scheme, let $n = 4$, with initial setting*

$$(\epsilon_1 = \frac{1}{4}, \epsilon_2 = \frac{1}{3}, \epsilon_3 = \frac{1}{2}, \epsilon_4 = 1)$$

This initial setting is already quasi-stable. A slightly different initial setting

$$(\epsilon_1 = \frac{1}{4}, \epsilon_2 = \frac{1}{2}, \epsilon_3 = \frac{1}{3}, \epsilon_4 = 1)$$

eventually settles on the quasi-stable

$$(\epsilon_2 = \frac{1}{2}, \epsilon_4 = 1, \epsilon_1 = \frac{1}{4}, \epsilon_3 = \frac{1}{3})$$

If β is lowered to $\frac{3}{4}$, the behavior changes: Both initial settings converge to the naturally-ordered sequence:

$$(\epsilon_4 = 1, \epsilon_2 = \frac{1}{2}, \epsilon_3 = \frac{1}{3}, \epsilon_1 = \frac{1}{4})$$

3.4 Indifference to Expert Delays

Our model (Section 2) posits that all queries and replies occur within a single integral unit of time: a round.

We now wish to expand the model to allow flexible timing: Let each expert i have a delay of δ_i between query and answer (or, between query time and the time at which satisfaction or disappointment of the user manifests itself).

By example, Figure 2 is the flow diagram and Markov chain of the 3-expert reputation-ordered scheme with generalized delays. It is a generalization of the basic model's flow diagram given in Figure 1 in which $\delta_1 = 1, \delta_2 = \delta_3 = 0$.

Clearly, a similar Markov chain and flow diagram exists for every selection order scheme: For every possible and relevant history of the scheme, there is a subchain of n expert nodes, ordered by the selection order. Each expert node, for an expert with expertise ϵ, has two emanating edges in the chain: First, a failure edge with probability $1 - \epsilon$, leading to the next expert in the selection order (or, for the last expert, back to the pre-round delay). Second, a success edge with

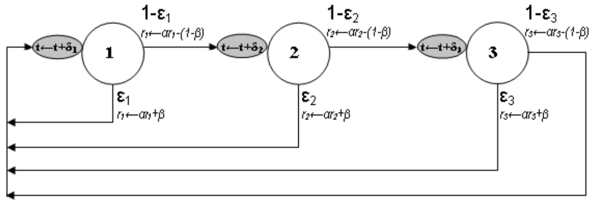

Figure 2: Flow, timing and feedback in the reputation-ordered scheme with generalized delays

probability ϵ leads to the pre-round delay of the updated user history.

In the general case, there are $M = mn$ expert nodes in the flow diagram, where m is the number of distinct user histories, and M delays, one per expert node.

So long as the reputation order of the experts does not change, each participating user may be seen as taking a random walk through the Markov chain. An order change changes the Markov chain.

We now claim that the delays are of no importance to the behavior of the model:

THEOREM 4. *Let there be n experts ordered by their reputations $r_1 > \ldots > r_n$ and M expert nodes in the selection order's Markov chain. Let $\Delta_1 = (\delta_{1,1}, \ldots, \delta_{1,M})$ be a set of expert delays for each expert node and $\Delta_2 = (\delta_{2,1}, \ldots, \delta_{2,M})$ be another such set. Then the values of reputation expert feedback under the two delay sets are proportional. Specifically, there exist constants for each delay set, $C_{\Delta_1}, C_{\Delta_2}$ such that for each expert i:*

$$w_i(t; \Delta_1)C_{\Delta_1} = w_i(t; \Delta_2)C_{\Delta_2} \quad (3.12)$$

Where the notation $w_i(t; \Delta)$ generalizes (2.1): The feedback of expert i at round t with set of expert delays Δ.

PROOF. Assume that in addition to node delays, there is a fixed edge delay $\delta > 0$. Consider the probability of finding the user in some particular edge, conditional on her being at any edge: This probability depends only on the Markov chain's graph and transition probabilities, and is independent of the node delays (or of δ).

With each edge is associated a reputation feedback. Per unit time, the expected reputation feedback from any particular edge is the edge feedback multiplied by the probability of finding the user at that particular edge during a unit of time. Since the edge feedback is constant, and the edge probabilities are in fixed proportions to each other, the theorem follows for any particular δ. In particular, it holds while $\delta \to 0$, and so holds in the limit, with no edge delays. □

COROLLARY 5. *The quasi-stability, instability and other behavioral aspects of a setting are independent of expert delays. The behavior under different sets of delays is identical with a suitable scaling of the time.*

3.5 The Loyalty Scheme

We now introduce a selection scheme in which users have memory:

Each round, each user remembers the expert that succeeded for her in the previous round (if there was such an expert), and selects him first in the current round. If this expert fails, the user will revert to the reputation-ordered

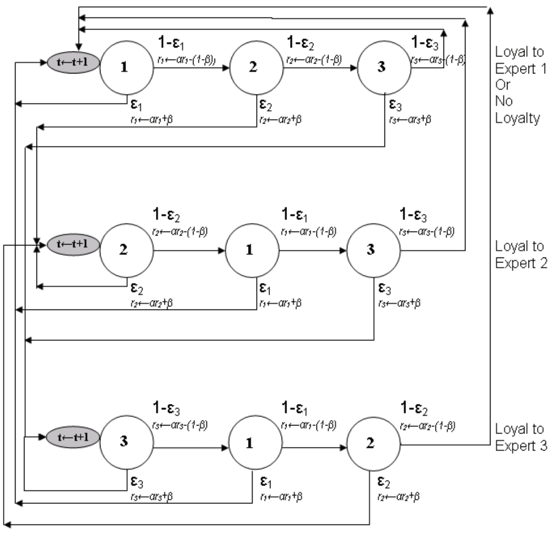

Figure 3: Flow, timing and feedback in the loyalty scheme

scheme, i.e. the second expert selected will be the highest-reputation so far unselected, etc.

We call it the **loyalty scheme**. The loyalty scheme attributes limited recall to the user: She remembers what worked out for her last time, but not more. Clearly the loyalty scheme models many real-world situations better than the memoryless reputation-ordered scheme. Our analysis will show that this limited recall significantly affects the quantitative and qualitative behavior of the model.

An example is given in Figure 3 for the 3-expert loyalty scheme, in which there are 3 expert nodes for each of the 3 relevant histories: the 3 distinct user loyalty states.

However, to avoid some complexities of the loyalty scheme, we first concentrate on a variation that is "well-behaved", which we will call the **dual loyalty scheme**. In the dual loyalty scheme, as in the loyalty scheme, the user first queries the expert who succeeded for her last time, and if this expert fails, will revert to the reputation-ordered scheme, querying all experts in order of descending reputation, but **without skipping over the first expert** to which the user was loyal. This necessarily means that the user may query this expert twice. So, for this modified scheme, we waive our standing assumption that asking the same question twice of an expert always yields the same answer.

This waiver, while admittedly artificial, has the merit of simplifying the analysis of the scheme, and, as we will see later, of making it regular.

The (unmodified) loyalty scheme is analyzed in the full version of this paper.

THEOREM 5. *Let n experts be indexed by their reputation order (at round t) $r_1(t) > r_2(t) > \ldots > r_n(t)$ and have respective expertise $\epsilon_1, \ldots, \epsilon_n$ all of which are < 1. Then the order is quasi-stable under the* ***dual loyalty scheme*** *if and only if for each $i \in [n-1]$:*

$$\frac{\epsilon_i + \beta - 1}{(1-\epsilon_i)^2} \geq \frac{\epsilon_{i+1} + \beta - 1}{1 - \epsilon_{i+1}} \quad (3.13)$$

PROOF. See the full version of this paper. □

Note that the requirement that all expertise levels be smaller than 1 is necessary, because if any experts have expertise of 1, the only possible steady-state is where **all** users are loyal to one of these experts.

COROLLARY 6. *The dual loyalty scheme is regular. This is directly observed from* (3.13).

COROLLARY 7. *Under the terms of Theorem 5 with a reward-only scheme ($\beta = 1$), quasi-stability requires the following equivalent inequalities for each $i \in [n-1]$:*

1. $\frac{1}{(1-\epsilon_i)^2}\epsilon_i \geq \frac{1}{1-\epsilon_{i+1}}\epsilon_{i+1}$
2. LEADER'S ADVANTAGE: $\frac{1}{1-\epsilon_i+\epsilon_i^2}\epsilon_i \geq \epsilon_{i+1}$
3. RECIPROCAL DIFFERENCE: $\frac{1}{\epsilon_i} - \frac{1}{\epsilon_{i+1}} \leq 1 - \epsilon_i$

Comparing these results with the corresponding results for the reputation-ordered scheme (Corollary 2), we observe that the loyalty property diminishes the value of a lead in reputation, although it does not nullify it: The leader's advantage is reduced from a factor of $1/(1-\epsilon)$ to $1/(1-\epsilon+\epsilon^2)$ and the reciprocal difference from 1 to $1-\epsilon$.

However, unlike in the reputation-ordered scheme where a leader's position with expertise of above $\frac{1}{2}$ is unassailable, loyalty does not allow for unassailable positions: A leader of expertise ϵ will be overtaken by a follower of expertise $\frac{\epsilon}{1-\epsilon+\epsilon^2}$ which is not greater than 1 and so feasible.

Two-sided quasi-stability is still feasible, but the window for it is narrower. The analogue of (3.9) is:

$$-(1-\epsilon_{i+1}) \leq \frac{1}{\epsilon_i} - \frac{1}{\epsilon_{i+1}} \leq 1 - \epsilon_i \quad (3.14)$$

In addition, we observe that the penalty-only scheme ($\beta = 0$) is, like in the reputation-ordered scheme, chaotic in any setting.

EXAMPLE 2. *Let the initial setting in a dual loyalty scheme ($n = 8, \beta = 1$) be:*

$$(\epsilon_1 = 0.1, \epsilon_2 = 0.2, \epsilon_3 = 0.3, \epsilon_4 = 0.4,$$
$$\epsilon_5 = 0.5, \epsilon_6 = 0.6, \epsilon_7 = 0.7, \epsilon_8 = 0.8)$$

This setting will eventually settle on the quasi-stable

$$(\epsilon_4 = 0.4, \epsilon_5 = 0.5, \epsilon_6 = 0.6, \epsilon_7 = 0.7,$$
$$\epsilon_8 = 0.8, \epsilon_3 = 0.3, \epsilon_2 = 0.2, \epsilon_1 = 0.1)$$

3.6 Other Observations regarding Steady-States

In addition to whether a steady-state exists, we may take interest in the value of reputation at steady-state, or rather, in its expectation. Referring to (2.2), we note that in an order-based scheme, $w_i(t)$ is constant so long as the reputation order is stable, which is the definition of a steady-state. Therefore, from (2.2), letting $w_i = \mathbb{E}[w_i(t)]$ and assuming an order-based scheme, it is observed that $\lim_{t\to\infty} \frac{r_i(t)}{t} = w_i(t)$ in the absence of discounting ($\alpha = 1$). On the other hand, with discounting ($\alpha < 1$), applying expectation to both sides of (2.2) , i'th (constant) expected feedback at steady-state, $r_i = \mathbb{E}[r_i(t)]$ has a fixed point satisfying $r_i = \alpha r_i + w_i$, and therefore:

$$r_i = \frac{w_i}{1-\alpha} \quad (3.15)$$

In the chaotic case, where no steady-state exists, and two-sided instability exists between an expert pair, we may ask what part of the time, on average, each of the experts has the lead. For order-based schemes, we provide the following answer:

Label a chaotic expert pair as 1 and 2. Let $w_i^j, i = 1, 2, j = 1, 2$ be i's reputation feedback when j has the lead. Then, define:

$$\Delta_1 \equiv w_2^1 - w_1^1$$
$$\Delta_2 \equiv w_1^2 - w_2^2$$

By the definition of two-sided instability, $\Delta_1 > 0$ and $\Delta_2 > 0$. Clearly, Δ_1 is the expected per-round change to the reputation difference $r_2 - r_1$ when expert 1 leads, while $-\Delta_2$ is the same when expert 2 leads. Let p_1 and p_2 be the probabilities that expert 1 and 2, respectively, holds the lead ($p_1 + p_2 = 1$).

The per-round expected change to the reputation difference is therefore $\Delta = p_1\Delta_1 - p_2\Delta_2$. As the reputation lead is expected to change an infinite number of times, necessarily $\Delta = 0$, therefore:

$$p_1 = \frac{\Delta_2}{\Delta_1 + \Delta_2}$$
$$p_2 = \frac{\Delta_1}{\Delta_1 + \Delta_2}$$

For example, using the reputation-ordered scheme (which is order-based), and with penalties only ($\beta = 0$):

Referring to (3.6) and (3.7), and noting that $\mathbb{E}[c_i(t)]$ depends only on experts ranked higher than i, which we may therefore mark as a constant C:

$$w_1^1 = C(\epsilon_1 - 1)$$
$$w_2^1 = C(\epsilon_2 - 1)(1 - \epsilon_1)$$
$$w_1^2 = C(\epsilon_1 - 1)(1 - \epsilon_2)$$
$$w_2^2 = C(\epsilon_2 - 1)$$

Therefore:

$$\Delta_1 = w_2^1 - w_1^1 = C\epsilon_2(1 - \epsilon_1)$$
$$\Delta_2 = w_1^2 - w_2^2 = C\epsilon_1(1 - \epsilon_2)$$

Expert 1's and 2's time-shares of the lead are in proportion $\frac{1}{\Delta_1} : \frac{1}{\Delta_2}$, and therefore in proportion $\frac{\epsilon_1}{1-\epsilon_1} : \frac{\epsilon_2}{1-\epsilon_2}$.

4. THE RATIONALITY OF USING REPUTATION AS AN INDICATOR OF EXPERTISE

As part of the model, we ascribed to each user a selection order of experts in which she queries experts in descending order of their reputation (barring private information about experts' history, as in the loyalty scheme).

The justification for employing such a selection scheme is that reputation is a positive indicator of expertise, i.e. that between any two experts, chances are 50% or better that the expert with the higher reputation has the higher expertise. But is this indeed the case? In other words, can a user, who knows that reputation is tallied through the feedback of a community of other users, each using some known or unknown selection rule, make a reasoned deduction that reputation signals expertise?

The question may appear puzzling in light of the results we have obtained, showing that often it is the less accomplished expert that is able to hold an indefinite lead in reputation over his betters, but to be aware of a possibility is not the same as knowing that it indeed occurred.

The question of rationality may be posed as follows: Suppose a particular user is aware of the expert selection methods employed by all other users. If in aggregate a community of users is more likely to reward the "losers" than the "winners" with a high reputation, the user would do well *not* to follow their collective advice, and eschew the high-reputation experts. However, if the community's behavior is sufficiently "well-behaved" to exclude such possibilities, using reputation as a signal for expertise is rationally justified.

That a user community may conceivably *not* be "well-behaved" is shown by the following example: Let all users have total recall of all their previous expert interactions, and let them each base their selection exclusively on expert success percentage, but using a *reverse* order: They give precedence to the expert that *failed* them most. Clearly such a scheme would reward the worst experts with the most customers, an advantage that may easily outweigh their lower success percentage and so provide the worst experts with the higher reputation feedback.

The anomaly in the above example is the irrational behavior of the users with their private information, i.e. their own experience: It makes no sense for them to prefer the experts that failed them most. We shall show that excluding irrational user behavior with their own experience is enough to make reputation a reliable signal of expertise:

In doing so, we want to allow the most general schemes through which users consider their experience with an expert: A user may choose to remember all previous encounters, or only the most recent m encounters, and she may attach significance to the order of experiences, e.g. A user who remembers the past two encounters with an expert, only one of which was successful, may value the experience higher if the success was on the most recent encounter, rather than in the penultimate one. However, for the valuation to be rational, a user must not value failure higher than success in any *particular* encounter, i.e. if a user remembers the most recent m encounters, and encounter $k \in [m]$ is the k'th most recent, then, changing that experience from success to failure, all other parameters held constant (i.e. all encounters except k, all experiences with other experts, and all expert reputations), may *not* advance the expert in the user's selection order. Briefly, in a rational user's selection order, an expert does not gain by failing in any particular trial.

As advancing in the selection order means a greater probability of the user becoming the expert's customer, we define this rationality requirement as *experience-monotonicity* of that probability:

DEFINITION 5. *Let a user remember her past m encounters with an expert. Let her experience be $Z \subset [m]$, such that $i \in Z$ iff the i'th most recent encounter was successful. Let $C(Z)$ be the probability for the user, with experience Z before some round to become a customer of the expert in that round. A selection order, and a user that employs it, are called* ***experience-monotone*** *if:*

- *The order depends on nothing other than the current expert reputations and on a (possibly partial) recall of experts' success-failure ratio.*

- *for each pair of expert experiences $Z1$ and $Z2$, $Z1 \subset Z2 \Rightarrow C(Z1) \leq C(Z2)$.*

All selection orders previously considered are experience monotone: Clearly the loyalty scheme is experience monotone, since a recall of a previous round's success moves an expert to first position. So are all selection orders that have no recall. Note that the monotonicity requirement is on experience alone, and does not rule out taking a non-monotone, and apparently illogical view of reputation.

The (rather weak) restriction of experience monotonicity allows us to prove the following general result:

THEOREM 6. *Observing the reputation of two experts (at some round t), the expert with the higher reputation is likely (with probability $\geq 50\%$) to have the higher expertise, provided:*

- *All users are experience-monotone.*
- *There is no prior information on expertise.*
- *There is no additional information showing that the reputation difference was previously (smaller t) larger or that it is smaller in the future (bigger t).*

PROOF. See the full version of this paper. □

Immediate corollaries are that observing all current expert reputations, the expert with the highest reputation is most likely to have the highest expertise, and that the expertise ranking order is most likely (out of all possible rankings) to be the reputation ranking order.

5. CONCLUSION AND FUTURE STUDY

We have presented a model to demonstrate and study the origins of reputation and its dynamics. Our key findings are that superior expertise may trail indefinitely behind superior reputation, and that reliance on reputation as a positive signal of expertise is correct strategic behavior for users of any reasonable reputation system. As consequence of the latter we may conclude that reputation systems may emerge by the participation and cooperation of autonomous agents.

For widest applicability of these results, we aimed to make the model as general as possible. Further study is needed in this regard: Our analysis was framed in the context of formal reputation systems, which by their nature enforce a uniform presentation of reputation on users. It would be natural to extend the analysis to the broader, informal phenomenon of reputation, formed in economic and social environments through various mechanisms, including word-of-mouth, the influence of mass media, and the influence of acknowledged authorities. This may lead to probabilistic selection rules proportional to reputation rather than rank-ordered by it.

Extensions of the model including money should be considered, with the **firm** in the role of expert. Access by users may be priced by the experts, or alternatively, may be free to the user but worth money to the expert (as in pay-per-click setups). Such a model may be augmented by the possibility of "buying" reputation (which we have eschewed in the current paper) by expenditure on advertising and other forms of marketing to raise a firm's brand/reputation. A further natural extension is the possibility of "buying" expertise, modeling expenditure on R&D as a means of increasing the quality of a firm's products and services, represented as expertise in our model.

We also propose to study differentiation between users by location or by social milieu, by treating reputation as a local rather than a global attribute of an expert. "Local" here should be understood in terms of a social or geographical network in which user feedback influences only the immediate neighbors of the user in the social or geographical graph, rather than contributing to a global, common-knowledge reputation as in the current study. Such a framework would then be suitable for studying social learning, a subject already extensively studied in the literature with various learning models, e.g. [1] [11], and the plausibility of the formation of stable barriers to influence in a social or geographical network.

6. REFERENCES

[1] BALA, V. AND S. GOYAL: Learning from Neighbors. *Review of Economic Studies*, 65, 595-621, 1998.

[2] BIKHCHANDANI, S., HIRSHLEIFER, D. AND WELCH, I: A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, **100**, 992-1026, 1992.

[3] BRIN, S. AND L. PAGE: The Anatomy of a Large-Scale Hypertextual Web Search Engine, *Computer Networks and ISDN Systems*, **33**: 107-117, 1998.

[4] CABRAL, L.: Stretching Firm and Brand Reputation, *RAND Journal of Economics*, **31**, 658-673, 2000.

[5] DELLAROCAS, C.: The Digitization of Word-of-Mouth, *Management Science*, **49** Issue 10, October 2003.

[6] FRIEDMAN, E., P. RESNICK AND R. SAMI: Manipulation-Resistant Reputation Systems in *Algorithmic Game Theory, Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani, editors Cambridge University Press*, 677-697, 2007.

[7] GITTINS, J. C.: Multi-armed bandit allocation indices. Wiley-Interscience Series in Systems and Optimization.. Chichester: John Wiley & Sons, Ltd.. ISBN 0-471-92059-2, 1989.

[8] KREPS, D.: Corporate Culture and Economic Theory. in J Alt and K Shepsle (Eds), *Perspectives on Positive Political Economy*, Cambridge: Cambridge University Press, 90-143, 1990.

[9] KREPS, D.M. AND R. WILSON: Reputation and imperfect information, *Journal of Economic Theory*, **27**, 253-279, 1982.

[10] MAILATH, G.J. AND L. SAMUELSON: Who Wants a Good Reputation?, *Review of Economic Studies*, **68**, 415-441, 2001.

[11] ROSENBERG, D., E. SOLAN, AND N. VIEILLE: Informational externalities and emergence of consensus. *Games and Economic Behavior* **69** (2009), 979-994.

[12] SELTEN, R.: The Chain-Store Paradox, *Theory and Decision*, **9** (1978), 127-159.

[13] TADELIS, S.: What's in a Name? Reputation as a Tradeable Asset, *The American Economic Review*, Vol. 89, No. 3, pp. 548-563, June 1999.