

IRM

INSTITUTE OF RESEARCH IN MANAGEMENT
INSTITUT DE RECHERCHE EN MANAGEMENT

ECONOMIC FEASIBILITY OF REPUTATION MECHANISMS FOR WEB SERVICES

WORKING PAPER 0605

Anthea Monod

And

Jean-Claude Usunier

ECONOMIC FEASIBILITY OF REPUTATION MECHANISMS FOR WEB SERVICES

ANTHEA MONOD & JEAN-CLAUDE USUNIER

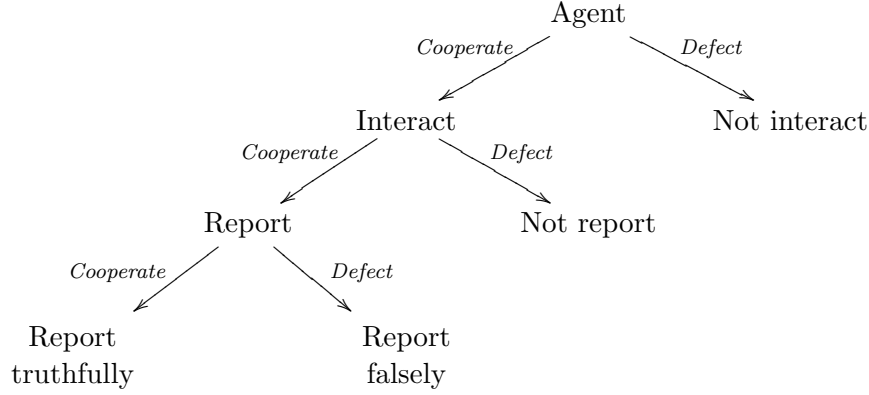
ABSTRACT. This paper presents a preliminary analysis of the economic feasibility of reputation mechanisms for web services [1]. We assess the economic and practical feasibility of the reputation mechanism by way of cost/benefit analysis from the perspective of the user, working from the hypothesis that rational users will cooperate, and thus ultimately contribute to the sustainability of the reputation mechanism when the benefits outweigh their associated costs. In addition, we perform an analysis of the examples of web services proposed in the above paper, and discuss factors in the startability and sustainability of reputation mechanisms for these examples.

1. INTRODUCTION

In order to assess the economic and practical feasibility of a reputation mechanism for web services [1], we perform cost/benefit analysis from the perspective of the user. Working from the basis of cost/benefit analysis, we then discuss factors in startability, sustainability and the economic feasibility of a reputation mechanism for examples of web services and propose settings for testing and simulation.

2. USER COOPERATION

The users of the web services and associated reputation mechanisms are assumed to be rational decision-makers. Summarizing the interaction protocol outlined in [1] from the practical perspective of the user, once she has discovered the existence of the web service, she is faced with making decisions of whether to engage in the web service (cooperate) or not (defect). Given that she decides to interact (cooperate), she can then decide whether to report (cooperate) or not (defect), and once again, given that she decides to report (cooperate), she decides whether to report truthfully (cooperate) or untruthfully (defect). This can be illustrated by way of a decision tree:



Given that the users behave rationally, we would like to hypothesize that users will be induced to cooperate at each node of the tree whenever the benefits of cooperating outweigh the associated costs. That is,

$$\text{Costs} \leq \text{Benefits} \implies \text{User Cooperation}$$

3. ANALYSIS OF COSTS AND BENEFITS

There are monetary and opportunity costs and benefits associated with each node of the decision tree. With the aim of encouraging user cooperation, we discuss these costs and benefits in this section, and explore the conditions that are necessary for user cooperation.

3.1. Interacting vs. Refraining. In deciding to interact with the web service, the user faces the monetary cost of paying the fee to use the service; she also faces the opportunity cost of foregoing the use of another web service for her needs. Thus, in order to induce the user to cooperate at this node, the fee for the use of our service must present the greatest quality/price ratio for the user when compared to the same service offered by competitors.

Depending on the service sought by the user, the benefits in terms of time opportunities and costs of utilizing the service to perform the task will generally exceed the costs of performing the task herself. Please see Section 4 below for a more detailed discussion of cost/benefit analysis and economic feasibility in regard to the particular scenarios discussed in [1].

3.2. Reporting vs. Refraining. In deciding to report on her experience with the service or not, the user faces the opportunity cost of time and effort spent on reporting; she does not face any explicit monetary costs in reporting. From the individual user's perspective, the cost of reporting should objectively be minimal in order to encourage reporting without first reflecting upon the benefits of doing so; that is, we would like the process to appear sufficiently simple and effort-minimizing to

the user so that she will consider reporting even before considering her gains from reporting. In order to minimize the time cost for the user, the step in reporting as designed by the mechanism engineers should be as user-friendly as possible, without, however, sacrificing the usefulness of the reputation mechanism to the general public.

For the user who is fully aware of and places great importance on the opportunity cost of her time and effort, and recognizes her considerable gain from utilizing the service rather than refraining (again, depending upon the service sought by the user to be discussed further), minimal required effort alone may suffice in order to induce the user to report. However, to secure the establishment of the proposed reputation mechanism, it would be necessary to consider users who require a more quantifiable gain in order to induce reporting. Addressing the needs of such users will not exclude the cooperation of other users who do not require such quantifiable gains from reporting; on the contrary, it will ensure their cooperation. These quantifiable benefits are discussed in [1] and emerge as monetary compensation in terms of rebates; that is, we offer a refund or rebate (partial refund) of the service fee in exchange for a report from the user on her experience. For certain scenarios discussed below, this particular strategy is pivotal in securing the establishment of the proposed reputation mechanism and use of the web service, since firstly, users are assumed to be one-shot players in the mentioned paper and therefore there is a strong inclination to free-ride on the reputation mechanism without contributing to its maintenance, and secondly, “free” web services for the example scenarios discussed in the same paper already exist.

One may point out the fact that the existing web services are not free at all, and that the service fees are embedded in the price of the end product that the user acquires. Indeed, this strategy successfully creates the impression that the service is free, however it sacrifices the element that is the fundamental difference between those existing web services and the proposed web service: the reputation mechanism. If a web service chooses to present itself as a provider of a free service, it cannot offer the monetary compensation in terms of service fee rebates to obtain feedback; it will have to rely upon the users who are fully aware of and place great importance on time and opportunity costs in order to obtain any feedback that pertains to the general experience of utilizing the web service. Otherwise, the feedback that such web services can expect will, in general, be limited to disappointing experiences and complaints, since we can assume that such experiences provide sufficient incentive for the user to report: revenge-seeking users who have been wronged are more abundant than altruistic users. In effect, “free” web services are more susceptible to skewed feedback and reputations.

However, it should be mentioned that our proposed reputation mechanism and web service is not expected to function as smoothly in economic practice as these “free” sites. Even under the service fee refund offer, there are still costs arising from interacting with our service instead of one of the other “free” sites. Firstly, there is the cost of the discounted time-value of money where the service fee is concerned; depending upon the time interval between the time the user reports her feedback and her service fee refund is processed, her money may be worth less. The monetary value of the opportunity cost of using our service is at least $(1 + r)f$, where r is the risk-free interest rate, and f is the fee of using the service. This particular cost issue may be considered negligible if the refund occurs electronically (and therefore promptly), or we process refunds via account credits (although this contradicts the assumption in [1] that users are considered to be one-shot players). Secondly, it is generally more inconvenient to engage in “free” services that require payment with an expected full refund at a future date than “free” services that have no explicit service fee.

3.3. Truthful Reporting vs. False Reporting. Once the user has made the decision to report on her experience with the web service, she is faced with the decision to either report truthfully, or to report falsely. Since we assume here that the appropriate considerations have been addressed by the mechanism designers so that the costs in effort of reporting are minimal, we can thus assume that once the user has decided to cooperate and report, the cost of reporting truthfully and the cost of reporting falsely equal; we can assume the cost of reporting truthfully or falsely is equal to the cost of reporting at all. Therefore, in inducing the user to cooperate at this stage, the discussion reduces to the analysis of the benefits that the user reaps from reporting truthfully or reporting falsely.

In discussing the economic feasibility of the reputation mechanism for a web service, the only assumption we have made, and indeed the only assumption that we can make at all, is that users are rational. From this assumption, we have hypothesized a decision rule that users will follow when faced with a decision to make. However, the stipulation that the users of the service are rational does not guarantee an injective correspondence between the decision to report and the motivation for doing so on any level; that is, when rational users decide to report, there cannot be a single objective that is common to all such users. A user may report because she values the opportunity costs of time and effort from using the service to perform her task as opposed to performing the task herself and decide to report truthfully as a reciprocal motive for the time-saving benefits that the service offers. She may also report truthfully (or untruthfully, since we have argued above that the costs are the same) in order to receive the service fee refund that serves as

compensation for her effort. Alternatively, she could also choose to report, but report untruthfully in order to fulfill some ulterior motive; for example, if she is satisfied with the service and finds the service useful, in order to avoid an excess in demand of the service to the limit where the service can no longer be delivered efficiently, she may even be induced to reporting falsely; at best, she may refrain from reporting altogether.

Hence, given the lax assumptions we have made in this analysis, and given that rational users will subjectively value the payoffs of acting on different motives on different levels, it would be impossible to measure the benefits on any quantifiable level that a given user would gain from reporting truthfully rather than reporting falsely without introducing further restrictions that would essentially render the analysis impractical and obsolete. As a result, it is impossible to analyze the conditions under which the hypothesized decision rule would hold true via cost/benefit analysis.

The main topic of analysis and discussion from this node would be the robustness and resistance of the reputation mechanism to user defection given the existence of different motives that are not formalizable. In [2], Nowak and Sigmund discuss the destabilization of a system when errors in perception prevail and argue that a system can be resistant to such destabilization attacks if all players (users) face the same error in perception. Would similar results hold for our reputation mechanism? Under what conditions would our reputation mechanism be resistant to defective reporting from the users, and would it be technologically feasible to incorporate such conditions in our mechanism?

In regard to the economic feasibility of the existence of a reputation mechanism, the user's decision of whether to report any feedback at all is arguably the most important node. Inducing user cooperation at the first stage to engage in the services to begin with is an issue concerning the marketing of the service; the reputation mechanism is a feature that can be exploited in order to attract more users. Inducing users to cooperate and report truthfully is a complex issue that may or may not affect the sustainability of the reputation mechanism, given that our mechanism is constructed in a decentralized manner and therefore there are no means of verifying whether truthful reporting occurs or not. However, before discussing the sustainability of a reputation mechanism, it is fundamental to assess the practical feasibility of its emergence, which reduces to user acceptance or rejection; the reputation mechanism would effectively be rejected if users decide not to report at all.

3.4. Deterministic vs. Stochastic Elements. Cost/benefit analysis must analyze both deterministic and stochastic aspects of costs and

benefits. Costs as well as benefits cannot be considered entirely certain since they are subject to the action/reaction of other players in the macro-reputation game (i.e. when a large number of players play repeated prisoner's dilemma game with strategic decisions, such as those depicted in the decision tree).

The costs of reporting are deterministic in regard to the time and effort spent in reporting. The costs of reporting also comprise a stochastic component, which in turn can be decomposed into direct and indirect reciprocity elements: the likelihood of retaliation by a player who responds to a negative rating with a negative rating exemplifies negative direct reciprocity, while the possibility of spoiling one's own reputation if one is considered by observers to be overly negative and libel-oriented exemplifies negative indirect reciprocity. Stochastic costs are typically associated with negative ratings (denoted by \mathcal{R}^-).

The benefits of reporting are deterministic in regard to the economic advantages of using reports to make one's decisions; that is, the economic advantages of the existence of the reputation mechanism as a public good. This deterministic component exists only if the information obtained from the reputation system can be viewed as truthful and reliable. The benefits of reporting also consist of a stochastic component, which similarly can be decomposed into direct and indirect reciprocity effects: the benefits arising from a potential increase in one's reputation when another player courteously responds to a positive rating with a positive rating demonstrates positive direct reciprocity, while the probability of being rated positively by a player who observes existing positive ratings (even in the absence of direct interaction with the player) demonstrates positive indirect reciprocity. Stochastic benefits are typically associated with positive ratings (denoted by \mathcal{R}^+).

Intuitively, the benefits of positive ratings will unfortunately be outweighed by the costs of negative ratings, as the deterministic component of the costs of reporting is easier to assess for participants of an online reputation system than the corresponding deterministic component of the benefits of reporting. Also, it can be assumed that punishment associated with negative ratings is far more likely to occur than reciprocal good turn following positive ratings.

With certain web services (particularly the second-hand bookstore discussed in [1]), if we assume that costs of reporting negatively strongly outweigh the benefits of reporting positively (that is, if $\mathcal{R}^- \gg \mathcal{R}^+$), then it is rational for participants of an online reputation system not to report negatively. This will result in massive underrepresentation of negative experiences, and much more abundant reports on positive experiences (which is apparently the case with eBay: negative experiences are reported in only 0.5% of the cases on average, according to [6]). Acknowledging the existence of a mediator (that is, the centralized organization that offers a web service used by both buyers and

sellers), we may also assign a reputation to the mediator via a proxy defined by the average proportion of positive ratings over the whole base of reports. It is then clear that such online reputation systems do not target the increase of trust between individual participants but the promotion of confidence in the mediator itself and in the electronic mediation it offers. The imbalance between costs of \mathcal{R}^- and benefits of \mathcal{R}^+ is exploited to inflate the reputation of the mediating web service.

This provides motivation to test the outcome when both negative and positive direct reciprocal feedback is either forbidden or signaled by making a clear difference between rate-first and rate-second strategies discussed by Dellarocas, Fan and Wood in [6]. In doing so, there is potential for massive reduction in the costs of \mathcal{R}^- , as well as significant diminishment of the benefits of \mathcal{R}^+ .

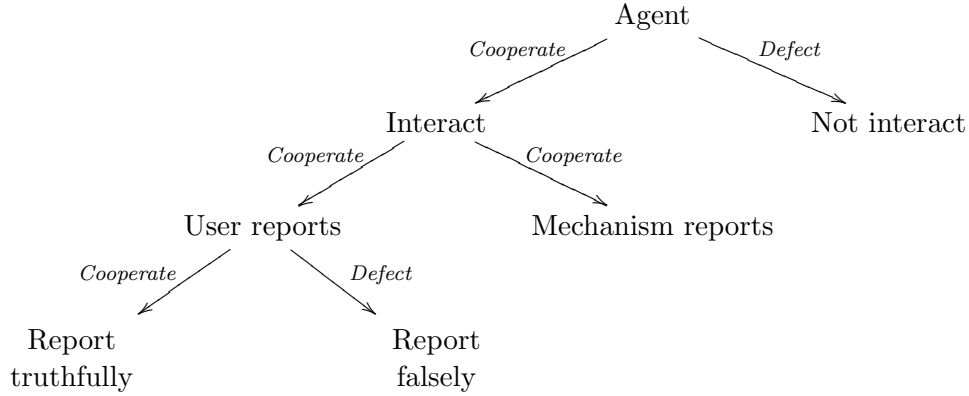
4. APPLICATION

In this section, we apply our cost/benefit discussion to the examples illustrated in [1].

4.1. Computation Service. This service fulfills the requirements of quantitative costs and benefits for concrete cost/benefit analysis. As for deciding to cooperate and interact with the service, assuming that the service functions efficiently and accurately, using the service for her computation needs would generally be more convenient and less costly than conducting the research and computation herself.

As for reporting on her experience with the service, for the particular case of the computation service, the authors of [1] introduce the implementation of an automated feedback feature of the reputation mechanism. In the absence of a report from the user, the system will automatically leave a report from the user which can be determined in two ways: either the transaction has occurred and in the absence of response from the user, the mechanism can assume that the experience was successful and leave positive feedback on behalf of the user; or the mechanism can take an average response to leave on behalf of the user (either an average of this particular user's past responses under a scenario where the user interacts multiple times, or an average of all users' responses).

This feature in effect eliminates one node of the user's decision tree, namely that which is essential for economic feasibility of the mechanism. The decision tree then becomes:



Since this feature leaves no choice to the user of whether to report or not, it effectively drives the user's costs of reporting to zero in the case where the user leaves the reporting to the mechanism. If the costs in this case were nonzero, the system and therefore the reputation mechanism would not be sustainable. In the case where the user decides to leave feedback herself, she would face opportunity costs in doing so, but these would necessarily be outweighed by the benefits she reaps, otherwise she would choose to costlessly allow the automated feedback submission feature to report on her behalf.

In this case, the feedback she leaves can be truthful or untruthful; the assumption that users are rational utility-maximizers does not rule out the possibility that some users may experience higher utility from the act of lying and "cheating the system" as opposed to passively reporting (via automated feedback submission) and will therefore choose to report falsely. This reaction may be quite frequent in the scenario where the user is aware of the automated feedback submission feature of the reputation mechanism; it is not uncommon for users to resent being deemed "predictable" by others, or dislike the idea of not being able to control their opinions and the general lack of options, and to react negatively in response, which may result in extreme feedback, regardless of what users may actually feel about the service. However, under the scenario where users are unaware of the automated feedback submission feature, this reaction will generally occur less frequently and the mechanism would therefore be more resistant to defective behavior.

Under the scenario where the user is unaware of the automated feedback submission feature, considerations of incentivizing reports (in particular, the anonymity of raters) become obsolete. The mechanism is startable due to the absence of costs, and credible reports are generated by the first reports, given the deterministic nature of the feedback under automated feedback submission. It is also sustainable; it is not threatened by cheating, collusion, lack of motivation to report on the part of the user, or legal action on the part of the service provider (since

one possible method of forcing user cooperation via reporting may be a legal stipulation on the part of the service provider, although such a stipulation would negatively affect the startability of the reputation mechanism. If a rational service provider (i.e. one that does not seek to undermine his own service) seeks to impose such a stipulation, the service must necessarily possess the features of a necessity (necessary good), which would therefore imply that the service provider acts as a monopolist so that it is indeed within her power to enforce such legal stipulations).

Taking into account all such considerations, this example would present the most ideal setting for testing economic feasibility due to relatively simple and well-defined cost/benefit analysis.

4.2. Price Comparison Service. Cost/benefit analysis would also be applicable to this type of service. The costs and benefits associated with this service are quantitative, and assuming that the service functions efficiently and accurately, the benefits of interacting and using the price comparison service outweigh the associated costs.

The user's decision to report at the next node depends upon on the completeness and correctness of the services delivered, and so it would be impractical to force user cooperation by means of an automated feedback submission feature of the reputation mechanism as in the case of the computation service. An example of a price comparison service would be an online travel reservation service, such as Orbitz or Travelocity. In the case of an online travel reservation service, completeness of service would relate to the completeness of information provided by the service, that is, the best prices are offered; correctness of service would correspond to the actual price charged being correctly reflected by the service.

In particular, reporting on the correctness of service may require a time-lag, ruling out the possibility for an automated feedback submission feature for this service. For example, using an online travel reservation service to reserve a rental car at an international destination would require a time allowance for reporting, as prices are not often correctly reflected online before arrival at the rental desk due to tax and insurance regulations imposed by different countries that may not be accurately calculable in advance by the web service. In the case that such costs are calculable, this service would extend to a combination of a price comparison service as well as a computation service, but while computation services are compatible with automated feedback submission, as discussed above, this is not the case for a combination of a computation service and a price comparison service. Using an online travel reservation service for flights would be easier to compute correctly than car reservations, since flights are a more restricted service than car rentals in that there is a unique way to travel from a given

airport to another by plane, but there are several ways to drive from one destination to another (shortest route, most scenic route, etc.). As such, there are more variables involved in different products offered by a single travel reservation service, and so automated feedback submission is inappropriate, and in order to analyze economic feasibility of our reputation mechanism, we must promote user cooperation to report as discussed previously.

With this particular web service, there are two possibilities where the reputation mechanism may be applied. First, the reputation mechanism may be applied to the variety of suppliers whose prices are being compared by the service. Using the example of the online travel reservation service, in this case, the mechanism would be applied to the different airlines or car rental companies. Applying the reputation mechanism to the suppliers may create a compounding effect of user perception of supplier reputation as users may have a preconception of each supplier's reputation which may or may not be confirmed by the reputation mechanism. In the case where the reputation of a supplier is unclear (for example, with a new supplier), the mechanism may be useful.

If all users have the same perception of a supplier, however, this would lead to either limited startability or limited sustainability of the mechanism: the reputation mechanism may be rejected altogether since it would serve no purpose, even in the case of error in the users' perception; or in the case it is not rejected, it would serve solely to confirm what the users already know, therefore providing redundant information, eventually leading to disintegration of the mechanism.

Secondly, the reputation mechanism may be applied to all price comparison services of a particular industry. Again, taking the example of online travel reservation services, the reputation mechanism would be applied to generate a reputation for each website (for example, Orbitz, Travelocity, CheapTickets, etc.) that offers such a service. In theory, the prices for each supplier being compared should be the same across all sites (for example, the price for a particular Swiss Air flight should read the same on all online travel reservation sites), and so any one price comparison service should be as effective as any other in retrieving the available prices. The reputation mechanism applied here would serve to provide all users' perception of each site's completeness and correctness of information.

It should be noted, however, that under the assumption that all information is freely available and retrievable to each provider of price comparison services of a particular industry, then there would be no issue of completeness or correctness of information provided by the price comparison service, but rather the completeness and correctness of information provided by each individual supplier being compared by the price comparison service, which reduces to the previous case discussed

above. Under this assumption, either the industry of price comparison services for some other given industry would be one of perfect competition, with perfect information and homogeneous products and prices (and therefore zero profits for all firms in the industry); or one of a monopoly where there exists one sole price comparison service in the industry. Such an assumption would render the reputation mechanism applied to all price comparison services of a particular industry irrelevant.

4.3. Second-Hand Bookstore. Cost/benefit analysis would also be applicable to this type of service. The costs and benefits are assessable, however less concretely than with the previous examples, as prices for used books can be negotiated between buyers and sellers as opposed to prices that are set by suppliers. The online second-hand bookstore provides a medium of exchange for users looking to buy used books and users looking to sell used books, and thus the benefits to both types of users of this service, buyers and sellers, outweigh the time and opportunity costs of searching for partners in order to conduct exchanges with, as well as the monetary cost of service fees. The details of the service with regard to the quality of service and service level agreements discussed in [1] are key in practical acceptance.

By the nature of the service, automated feedback submission is not applicable to the reputation mechanism, and so we must induce user cooperation to report. However, with this service, the decision to report or not, as well as the decision to report truthfully or not, is more subtle than making the same decisions with the other services. Here, although the users do engage in a service and pay a service fee for the bookstore to provide the medium of exchange, the direct interaction that must occur for users to attain their goals takes place not between a client and service provider, but between two clients: a buyer and a seller. As such, some users may take a greater responsibility of their decision (whether they decide to report or not). This scenario reduces to the situation of a game between two active players, as opposed to the previous situations which may be thought of as a game between one active player and one passive player, where the user can be considered as an active player and the firm or service provider is a passive player (since interacting with a firm or service provider is more of an abstract setting than interacting with a single other player).

An advantage of this setting is that users may experience low susceptibility to cheating, which is controlled by an understanding of direct reciprocity between the players. Since neither player knows with absolute certainty whether this “game” that they “play” with each other when they decide to exchange a second-hand book through this service is a one-shot game, finite repeated game, or infinite repeated game,

neither player will know for sure what behavior with respect to reporting would optimize her outcome for this or future games. As such, in the case of interactions that were not negative, both players may experience a sense of joint responsibility to deter from reporting. In the case of interactions that were negative, as discussed above, players will be strongly incentivized to report.

There are also several possibilities for the application of the reputation mechanism to this type of service. First, the mechanism may be applied to generate reputations for each user of the system in their dual roles of both buyers and sellers (since all users who register with the service may either be buyers or sellers or both). Secondly, the mechanism may be applied to generate a reputation for the second-hand bookstore itself as a medium of exchange for buyers and sellers, although this application would defeat the goal of our project of creating a self-propagating and self-maintaining system that is ultimately decentralized. In this situation, as the users would not rate other users but only the service itself, the only features that users would be able to comment on are efficiency of the service to bring together buyers and sellers efficiently (which in itself is a subjective feature, since most users are exposed to only a small group of users out of the entire system's user database), as we are aiming for a decentralized system and thus users are not able to comment on the service's capacity as mediator or exchange supervisor. The reputation mechanism may also be applied to both users and the service, in theory, but may reduce to a reputation mechanism only for users.

The reputation mechanism itself may be difficult to start due to this setting of a two-player game, as discussed above. Its startability is also limited by the nature of the service, as the system is not commodity oriented; that is, the service does not mass-produce commodities that all users can consume as demanded. With the second-hand bookstore, the good in question is a single item that is negotiated over by several buyers, with a single supplier with ultimately a single winner. This two-player game setting and the players' understanding of direct reciprocity may also limit the reputation mechanism's sustainability. In order to test a reputation mechanism for this type of web service, this direct reciprocity between users would need to be ruled out via stipulations that promote one-sided reporting over direct reciprocity, although such a setting is difficult to reproduce and maintain for practical use outside a testing environment.

4.4. Dating Service. In terms of economic feasibility of reputation mechanisms, this type of web service poses the greatest complications. Costs and benefits of this type of web service are not easily quantified for analysis, the understanding of quality of service for this type of web service is highly subjective, not only on the part of the service users, but

also the service providers. Perceptions of failures in service for the user (which can be viewed as opportunity costs) would correspond to failures to find any matches through the service. Even in the case of perceptions of success in service at a given instant in time (for example, success in finding a partner through the service), there is always the possibility that the relationship will not last, and thus the possibility that this failure could be perceived as a failure in service. However, from the perspective of the service provider, such failures could be perceived as a “defective” user, and not a failure of the service. With such a conducive setting for mirror criticism, costs and benefits are difficult to define, although the existence of several such services (eHarmony, Match.com, etc.) provide evidence of practical success of the services, which is not indicative of success of reputation mechanisms for such services. There is no way to assess economic feasibility through cost/benefit analysis with this type of service.

In terms of establishment of a reputation mechanism, it is unclear where the reputation mechanism should be applied. The mechanism may be applied to rate the service, but as discussed above, the meaning of any feedback, positive or negative, is obscure; the mechanism and service are both highly exposed to disintegration and criticism from users who act upon reports given by the mechanism and experience opposite results. For the same reason of obscurity, the mechanism applied to generating reputations of users is also exposed to disintegration. The service provider would also need to provide a set of courtesy rules for the reporting of other users, which increases costs of reporting for the user, which effectively attacks the sustainability of the reputation mechanism, assuming that the reputation mechanism manages to establish itself. The very startability of the reputation mechanism is also under attack by the nature of the service, since there is the possibility of users engaging in the service for mere entertainment aims. If these prevail over intended aims of partner-seeking, this would impede the establishment of the reputation mechanism.

4.5. Proposals for Simulation. We now elaborate further on the mechanics of a reputation mechanism and discuss aspects of settings for tests and simulations. The table following on end provides a summary of the cost/benefit assessment analysis, and startability and sustainability issues of reputation mechanisms for the above web services. It also provides a summary of other issues arising in the development of reputation mechanisms for these services addressed in [1]. The flow diagram motivated by the service quality model proposed by Parasuraman, Zeithaml and Berry in [5] on end provides a description of the quality of service assessment by the user.

Computation Service: This setting may be used to test the influence of an online reputation system on prices, effort levels (i.e. quality of service), probability of sale (i.e. market share), and market structure (i.e. number of service providers and their respective market share).

Price Comparison Service: Here, the quality of service (completeness and correctness) may be parametrized. This setting may be used to test the influence of reputation information on the effort level of the service provider, as well as the emergence and sustainability properties according to different levels of completeness and correctness that reflect the effort level of the service provider.

Second-Hand Bookstore: We would like to use the setting to show that, as in eBay, the online reputation system serves the purpose of increasing the mediator’s reputation rather than that of individual users. Empirical evidence in [6] suggests that negative experiences are largely unreported since negative ratings are very infrequent compared to positive ratings. This results in overstating the reliability of the intermediation system and, consequently, artificially increases user-confidence in it. The reputation system is confidence- rather than trust-driven: it targets a general sense of confidence in the system, as opposed to a sense of risk-taking with particular agents whom one may choose to trust.

Dating Service: We use the same assumption as in [1] that all users have a clear motive, which can also be considered as ulterior, in rating the service positively in order to attract more participants into the pool of potential dates. A large user base is in the best interest of all date-seekers, especially if there is an imbalance in key characteristics of date-seekers (e.g. more males than females or older date-seekers than desired by participants); consequently, service providers are not the reputation targets. This is especially true where service providers may clearly inform users about the dating service and its mechanics, establish a code of conduct, and finally transfer to the date-seekers the full moral and legal responsibilities of inadequate behavior. Thus, we consider date-seekers as the only reputation targets for the reputation mechanism. Each date-seeker acts anonymously under a pseudonym; rates (i.e. participants that are potential matches) are described by a number of characteristics, such as sex, age, sexual preferences, marital status, health status (which can be rated on a discrete scale, such as 1 to 5), outlook/physical appearance (also rated on a discrete scale), personality (also rated on a discrete scale), etc.; weights for particular characteristics may be assigned. Date-seekers may cheat on some of their true characteristics, however this will be discovered later when a face-to-face encounter (date) takes place. Defective behavior in relation

to particular characteristics of the declared profile can be parametrized in a probabilistic manner.

Date-seekers also display a preferred profile for potential partners with the same type of information as in their own declared profile. This profile is assumed to always be true. Matching may occur when a source preferred profile corresponds to the target declared profile for both date-seekers; the system offers a date for matched profiles. However, each date-seeker is in a position to refuse the date if he or she perceives the other date-seeker's reputation as inadequate. When the date takes place, the truthfulness of the declared profiles is revealed to each date-seeker. From the rater's point of view, reputation information is subjective; that is, reputation information is always true. The logic of indirect reciprocity may be applied to support this claim: each agent may rate all other agents, but reciprocal feedback has no effect on the calculation of average reputation scores. It is, however, stored in the reputation profile of a particular date-seeker under four positions: positive first, negative first, positive second, negative second. Positive second and negative second ratings are likely to be caused by positive (courtesy rating) and negative (retaliatory rating) direct reciprocity respectively. Each individual reputation profile is accessible to all date-seekers. The logic of assessment rules and action rules as described by Nowak and Sigmund in [2] can be applied to this type of setting. The cost of reporting can be neglected since participants are rewarded (that is, they are more than compensated for their efforts) by the playful nature of the dating service. The key issue is to determine whether such a service produces adequate matching between date-seekers based on such reputation information; that is, what does reputation information add to the efficiency of the system, compared to pure random matching (i.e. the average quality of matching)?

On the basis of actual profiles there is high probability that matching will be poor since the best potential targets may have already found a mate by other, more traditional ways; consequently, an imbalance between targets and sources is very likely to occur. However, online date-seekers are likely to embellish their reported profile and to ask for more in terms of the target's desired/preferred profile, and the resulting embellished profile would attract more potential matches than what their actual profile would allow them to expect, demand, and attract on average. It is possible for the dating service to generate user-disappointment when dates occur. We hypothesize that disappointment occurs only when there is discovery of a clear lie on a key characteristic (for example, age and physical appearance by at least two to three positions on the discrete rating scale). When both date-seekers discover that their reported profiles are true, neither is disappointed because they have been offered a fair match and their self-esteem is

enhanced by being matched with an honest date-seeker. They are satisfied, even if this encounter is without further promises; that is, when deeper affinity does not emerge. A probability distribution may be assigned to the event of a final match where both date-seekers leave the pool as potential dates.

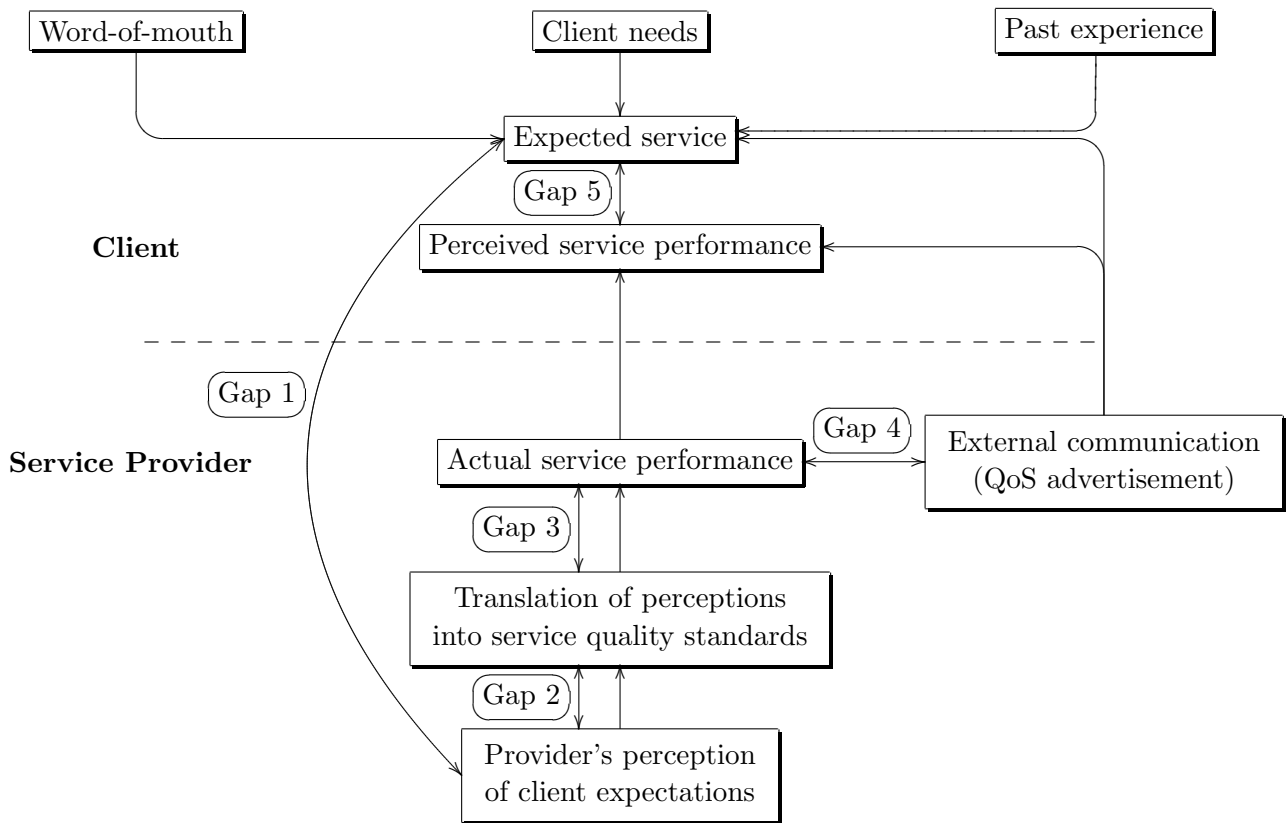
Reputation information is likely to induce a change in the behavior of a user of such an online dating service. A revealed cheater (that is, a user who embellishes his own profile, “Bad”) may retaliate by posting reciprocal negative information on the user who has revealed him, “Good”, even though Good may have reported her profile truthfully. In this manner, Bad may offset the effect of a negative first rating on his average reputation score. However, Bad’s reputation profile is accessible to all date-seekers, both those who are interested in and those who are not, a date with Bad. Depending on the rules set forth by the reputation mechanism designers, these date-seekers may also post a rating on Bad based on what they observe (an act of community enforcement): if a date-seeker like Bad has a rich profile in terms of negative second ratings compared to other date-seekers that have few or no negative second ratings and a “good” average score, it is likely that some date-seekers will rate Bad negatively first. Bad may go on retaliating, but he accumulates a very high profile in terms of negative second ratings relative to many members of the online date-seeking community. On average, the system should introduce a mechanism to incentivize honest behavior, as once cheaters are discovered, their reputation profile will be dubious and they will no longer be desirable dates. This setting may be used to test the effect of anonymity and the likelihood of deception on reputation information.

5. CONCLUSION

We have discussed necessary conditions for user cooperation to the extent of reporting under the decision rule that users will report if the associated benefits outweigh the associated costs. Are they sufficient? One could argue under the assumption that all users are rational and practical decision-makers, these conditions may indeed be sufficient, however, this assumption does not prevail unconditionally. We require further exploration into the impact of this fact, as well as the possibility of defective reporting, on the robustness and stability of the reputation mechanism.

Properties	Computation Service	Price Comparison Service	Second-Hand Bookstore	Dating Service
Forms of Reporting	Automated Feedback Sub-mission	Voluntary feedback on completeness and correctness	Qualitative reporting or multi-dimensional quantitative	Purely qualitative, subjective and ill-defined
Susceptibility to Cheating (“truthful” from the rater’s perspective, collusion between raters)	Very Low: Automated feedback makes lying unattractive; collusion is unlikely, however possible	High, due to bias of overly negative feedback and susceptibility to collusion	Low: Cheating is controlled by direct reciprocity mechanisms between buyer and seller	Not Applicable: frontier between truthful and untruthful reporting is too vague
Incentive to Report (for Users)	Not applicable	High for negative experiences; low otherwise	High for negative experiences, low otherwise; a sense of joint responsibility may deter both parties from reporting	High if anonymous (free self-expression is rewarding), low otherwise (people fear being outlawed)
Anonymity of Raters	Not necessary due to automated feedback	Anonymous rating increases susceptibility to cheating	Anonymous rating increases susceptibility to cheating	Often required before match occurs (fake identities are the rule)
Symmetry	Asymmetrical	Asymmetrical	Partial symmetry required	Asymmetrical between providers and date-seekers; symmetrical between date-seekers
External Communication	Minimal advertisement is possible	Minimal advertisement is possible	Detailed Service advertisement required	Detailed Service advertisement required

Properties	Computation Service	Price Comparison Service	Second-Hand Bookstore	Dating Service
Quality of Service Delivery	Provider only	Depends on provider and on user's ability to efficiently use the service	Depends heavily upon buyer/seller interaction; sensitive to size and symmetry of user base	Depends upon client/provider interaction; sensitive to size and symmetry of user base
Quality of Service Assessment by User	Simple gap model (gaps 4 and 5)	Simple gap model (gaps 4 and 5)	Simple gap model (gaps 4 and 5)	Complex gap model (gaps 1 through 5)
Reputation Target	Service Provider	Service Provider	Service Provider rather than buyer or seller (or each of them?)	Provider (perhaps also date-seekers?)
Startability	High: Credible reports start with the first report given the deterministic nature of feedback	Requires a sufficient number of reports to derive a probabilistic estimate (or any kind of feedback aggregation)	Relatively difficult to start, not commodity-oriented but deals with a singularized item	Very Difficult to Start: Reputation system is threatened by chaotic reporting if fun and fancy dominate over mate-seeking
Sustainability	Reputation system is not threatened by cheating, collusion, lack of client motivation to report or illegal action by provider	Reputation system is threatened by lack of client motivation to report; cheating or collusion are unlikely	Reputation system is threatened by lack of buyer and/or seller motivation to report	Provider needs to corset reporting into a set of courtesy rules, thereby reducing attractiveness
Assessment of Economic Feasibility	Easiest setting for testing economic feasibility due to relatively simple cost/benefit analysis	(Requires several assumptions to make cost/benefit analysis feasible, especially the assumption that prices announced are both complete and correct)	Best setting to test a number of mechanisms that forbid direct reciprocal feedback and promote one-sided reporting	A good setting to test what happens when anonymity is the rule and all reports are assumed to be true



REFERENCES

- [1] R. Jurca, B. Faltings, L. H. Vü, J.-C. Usunier, K. Aberer, A. Perrinjaquet *Reputation mechanisms for web services*, EPFL preprint (2005).
- [2] M. A. Nowak, K. Sigmund *Evolution of indirect reciprocity*, Nature **437** 1291–1298 (2005).
- [3] R. Jurca, B. Faltings *Reputation-based service level agreements for web services*, EPFL preprint (2005).
- [4] L. H. Vü, M. Hauswirth, K. Aberer *QoS-based service selection and ranking with trust and reputation management*, EPFL preprint (2005).
- [5] A. Parasuraman, V. Zeithaml, L. Berry *A conceptual model of service quality and its implications for future research*, Journal of Marketing **49(4)** (1985).
- [6] C. Dellarocas, M. Fan, C. Wood *Self-interest, reciprocity, and participation in online reputation systems*, MIT Sloan School of Management working paper 4500-04 (2004).
- [7] C. Dellarocas *The digitization of word-of-mouth: promise and challenges of online feedback systems*, Management Science **49(10)** 1407–1424 (2003).
- [8] C. Dellarocas, P. Resnick *Online reputation mechanisms: a roadmap for future research*, Summary report of the First Interdisciplinary Symposium on Online Reputation Mechanisms (2003).
- [9] Z. Despotović *Building trust-aware P2P systems: from trust and reputation management to decentralized e-commerce applications*, EPFL PhD thesis 3313 (2005).

UNIVERSITÉ DE LAUSANNE, ÉCOLE DES HAUTES ÉTUDES COMMERCIALES
E-mail address: Anthea.Monod@unil.ch, Jean-Claude.Usunier@unil.ch