

Reputation-based Pricing of P2P Services

Radu Jurca
Ecole Polytechnique Fédérale de Lausanne
(EPFL)
Artificial Intelligence Laboratory
CH-1015 Lausanne, Switzerland
radu.jurca@epfl.ch

Boi Faltings
Ecole Polytechnique Fédérale de Lausanne
(EPFL)
Artificial Intelligence Laboratory
CH-1015 Lausanne, Switzerland
boi.faltings@epfl.ch

ABSTRACT

In the future peer-to-peer service oriented computing systems, maintaining a cooperative equilibrium is a non-trivial task. In the absence of Trusted Third Parties (TTP's) or verification authorities, rational service providers minimize their costs by providing ever degrading service quality levels. Anticipating this, rational clients are willing to pay only the minimum amounts (often zero) which leads to the collapse of the market.

In this paper, we show how a simple reputation mechanism can be used to overcome this moral hazard problem. The mechanism does not act by social exclusion (i.e. exclude providers that cheat) but rather by allowing flexible service level agreements in which quality can be traded for the price. We show that such a mechanism can drive service providers of different types to exert the social efficient effort levels.

Categories and Subject Descriptors

I.2.11 [Distributed Artificial Intelligence]: Multiagent Systems

General Terms

Design, Economics, Security

Keywords

Reputation, Service Level Agreement, Service Oriented Computing

1. INTRODUCTION

Service oriented computing systems represent an attractive paradigm for the business world of tomorrow. User requests (e.g. ranging from a trip reservation to determining the optimal control of a specific process) are no longer atomically treated by monolithic organizations, but rather decomposed into smaller components which are separately

addressed by different service providers [10]. While the advantages of such a scenario are uncontested (simplicity, easy of management and customization, fault tolerance and scalability), a whole new set of problems need to be addressed.

Trusting service providers is a central issue for the practical success of such markets. The peer-to-peer interactions between clients and providers do not benefit from the presence of traditional trust-establishing tools: trusted third parties or enforcing authorities are not available to detect and punish violations of the service level agreement. Free from most forms of legal retaliation, service providers can maximize their profits by investing the least amount of effort in provisioning the service. Rational clients anticipate this moral hazard and will be willing to pay for the services only the minimum possible amount. This vicious circle leads to a "market of lemons" [2] and eventually to the total collapse of the market.

Reputation mechanisms have been widely accepted as social trust-establishing systems. The reputation of a provider (i.e. information about past behavior) influences the present and future revenues of an agent. Present cheating translates into bad future reputation, and therefore, lower future rewards. When carefully designed, reputation mechanisms make cheating economically uninteresting and support an efficient market. [8]

Most of the existing reputation mechanisms (eBay, Slashdot, Amazon) use the average of past feedback reports to assess the reputation of one agent. Clear semantics is not attributed to reputation information and clients use ad-hoc rules (i.e. trust if reputation is above a certain threshold) to take trusting decisions. Such systems tend to separate agents in two categories (trustworthy or untrustworthy) and do not leave much room in between.

Black and white mechanisms cannot work in a service-oriented environment. The perfect service is impossible to provide (or prohibitively expensive), and various providers might prefer different production quality levels. Such settings require a more flexible mechanism in which quality of service can be traded, for example, with the price: i.e. clients consciously accept a lower service quality level for a lower price.

In this paper we show how a very simple mechanism (based on averaging feedback) can also work in a service-oriented environment. Repeated failures do not automatically exclude a provider from the market (as they would on eBay), but rather influence the price the provider can charge for a future service. Service providers can thus rationally plan

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SIGCOMM'05 Workshops, August 22–26, 2005, Philadelphia, PA, USA.
Copyright 2005 ACM 1-59593-026-4/05/0008 ...\$5.00.

how to invest available resources in order to maximize their (and the social) revenue. Such a reputation mechanism does not act by social exclusion but rather through incentive-compatible service level agreements.

Before we go into the details of the mechanism, let us consider a practical example. Consider a service which provides computation power: clients submit processing jobs and wait for the answer. A client is satisfied given that the correct¹ answer is provided before a certain fixed timeout. We assume that an incorrect answer (or a late answer) does not have any value for the clients.

The probability that request are satisfied directly depends on the capacity of the provider, and on the number of service requests in a given “day”.² Given that environmental perturbations are small, the provider can fine-tune in every day the expected success rate of the provided service, by deciding what computing resources to allocate, or how many client requests to accept.

We start from the assumption that the perfect service is impossible to provide (i.e. has infinite cost). Unless the provider has unlimited resources, there always is a small chance of service failure due to a conflict of resources. On the other hand, clients can be happy with less than perfect service if the tradeoff between price and quality is right. When the provider has an idea about the price curve of the clients (i.e. how much clients are willing to pay for various rates of successful service) it can find the optimal success rate that should be guaranteed: i.e. the success rate q^* which maximizes the total reward: $price(q) - cost(q)$. q^* is the socially efficient “quality” level of the provider.

Different providers have different characteristics (also denoted as “types”) and therefore, different cost functions, and different efficient quality levels. It is usually the case that the set of efficient quality levels of all types spans a continuous interval between q_{min} and q_{max} , the generally accepted limits in which the success rate can vary.

The role of the reputation mechanism is to drive every provider (regardless of its type) to provision the optimal success rate, despite the temptation to cheat. We show how a simple mechanism based on averaging past feedback and reputation-based service level agreements can push the market towards an efficient equilibrium point.

Section 2 formally presents the setting in which we situate our work. A detailed description of the reputation mechanism, as well as proofs about its properties are presented in Section 3. Section 4 comments on the validity of the assumptions that we make, and discusses practical issues regarding the implementation of our mechanism in real settings. We conclude by presenting related and future work.

2. THE MODEL

We consider a P2P system where rational service providers (sellers) repeatedly offer the same service to the interested clients (buyers), in exchange for payment. We assume that service requests are uniformly distributed over a (possibly infinite) number of “days”, with the understanding that our “day” can have any pre-established length of physical time.

¹We assume that the client can easily verify if the answer is correct or not. e.g. the solution of a constraint satisfaction problem can be easily verified for correctness.

²We consider the “day” to be the atomic decisional time unit of the service provider. Our “day” can in fact represent a second, an hour or a week of physical time.

Clients observe only two quality levels: “satisfactory” quality for a successful invocation or “unsatisfactory” quality for a failure. A satisfactory service has value v to the clients, while unsatisfactory service is worthless.

The observation of the clients depends on the effort exerted by the provider, and on external factors characterizing the environment (e.g. network or hardware failures). External factors are assumed to be “small enough”, and time-independent. Effort is expensive, however, it positively impacts the probability that clients experience a satisfactory service.

Providers have different characteristics (cost functions, capabilities, etc) which define their *type*. We denote the set of possible types as Θ ; members of this set are denoted as θ . Excepting environmental factors, the type of the provider, and the effort invested in providing the service, completely determine the probability that the clients will be satisfied.

Let $c : [q_{min}, q_{max}] \times \Theta \rightarrow \mathbb{R}$ describe the cost function of service providers, such that $c(q, \theta)$ is the cost of the effort that needs to be exerted by a type θ seller in order to provide a satisfactory service with probability q . q can be viewed as a measure of the “production quality”; thus, the experience of the clients is an imperfect discrete observation of the continuous production quality q . As it is the case for most computational services, higher probability of success is increasingly more expensive to guarantee: i.e. $c_1(q, \theta) > 0, c_{11}(q, \theta) > 0$, for all $\theta \in \Theta$.³

The expected utility of the clients can also be expressed as a function of the production quality level, q . While different agents may have different utilities for the same value of q , we assume that the market has a common price function $p : [q_{min}, q_{max}] \rightarrow \mathbb{R}$ such that all clients agree to pay at most $p(q)$ for a service that will prove satisfactory with probability q . In other words, all clients have utilities above $p(q)$ but for some (private) reasons they do not agree to pay more than the market price $p(q)$. The price function can be learned by analyzing previous markets, and is assumed to be known by all agents (providers as well as clients). The function is convex, linear or concave when the clients are risk-averse, risk-neutral, respectively risk-seeking.

We assume that the price function and the cost function of all provider types can be piecewise linearized on the set of intervals $\{[q_i, q_{i+1}]\}$, $i = \{0, \dots, N - 1\}$, such that: $q_0 = q_{min}$, $q_N = q_{max}$ and $q_i < q_{i+1}$ for all i . The cost functions and the price function can thus be approximated by:

$$c(q, \theta) \simeq C_0(\theta) + \sum_{i=1}^{t-1} C_i(\theta)(q_{i+1} - q_i) + C_t(\theta)(q - q_t);$$

$$p(q) \simeq P_0 + \sum_{i=1}^{t-1} P_i \cdot (q_{i+1} - q_i) + P_t \cdot (q - q_t);$$

when $q_t < q < q_{t+1}$. The marginal costs $C_i(\theta)$ become increasingly bigger for all seller types (the cost function is convex), and we assume that each provider can choose each day the effort it wants to exert the next day. However, all requests coming during one day are satisfied with the same “quality” (i.e. probability of success) q . The marginal prices

³ $c_1(\cdot, \cdot), c_2(\cdot, \cdot)$ denote the first order partial derivatives of $c(\cdot, \cdot)$ with respect to the first, and respectively the second parameter; $c_{11}(\cdot, \cdot), c_{12}(\cdot, \cdot), c_{22}(\cdot, \cdot)$ denote second order partial derivatives.

P_i are increasing (risk-averse clients), equal (risk-neutral clients) or decreasing (risk-seeking clients).

Providers advertise a service level agreement (SLA) for each day. The SLA contains (among others) the promised quality level, and the price of the service. If clients could perfectly observe the type and the effort level of the providers, the market would function at the efficient level: All provider types would provide (and advertise in the SLA) the *efficient production quality level*, $q^*(\theta) = \operatorname{argmax}(p(q) - c(q, \theta))$, and all clients would pay the price $p(q^*(\theta))$. When using the piecewise linear approximations of the price and cost functions, $q^*(\theta)$ corresponds to the quality level $q_{i_\theta^*}$ such that $C_{i_\theta^*}(\theta) < P_{i_\theta^*}$ and $C_{i_\theta^*+1}(\theta) > P_{i_\theta^*+1}$. We assume that this point is unique.

Figure 1 shows a risk-averse and a risk-seeking price function, a cost function of a certain type θ , and the corresponding efficient production quality level of the type θ .

However, since perfect information is not available, the expected utility of a service invocation can be estimated only from the previous behavior of the provider. For that purpose, a reputation mechanism collects feedback from the clients about the quality of service of different providers. Feedback is binary (1 for satisfactory service, 0 otherwise) and aggregated on a daily basis.

Let R_t be the reputation of a service provider at the beginning of day t , and let r_t be the set of feedback reports submitted about the same provider during the day t . The reputation mechanism computes the new reputation of the provider for day $t+1$ as $R_{t+1} = f(R_t, t, r_t)$ where f is some stationary (i.e. time-independent) function.

By taking into account the SLA and the reputation of the provider, the clients decide whether or not the purchase the service. Each service provider type will choose for every day the SLA and the production quality level which maximizes its overall payoff. If $d(sla_t, R_t)$ is the price paid by a buyer in day t to a provider with reputation R_t advertising the service level agreement sla_t , the overall payoff of the provider of type θ is:

$$V(R_0, \theta) = \max_{(q_t, sla_t)} \sum_{t=0}^{\infty} \delta^t (d(sla_t, R_t) - c(q_t, \theta));$$

where δ is the daily discount factor of the provider.

The remaining question is (1) how to chose the reputation updating function f , and (2) what decision recommendation to give to the clients (i.e. the function $d(sla_t, R_t)$) such that:

- providers have the incentive to exert the effort level that maximizes social revenue;
- clients do not have the incentive to deviate from the recommendations of the reputation mechanism.

3. REPUTATION MECHANISM

The reputation mechanism we propose is very simple:

1. the reputation of a service provider in day t is the probability of a successful service invocation in day $t-1$,
2. buyers should not buy if the SLA in day t specifies a price greater than $p(R_t)$, i.e. the price corresponding to a quality level equal to the reputation in day t .




Figure 1: A risk-averse and risk-seeking price function, and the cost function for a service provider type.

Formally,

$$R_{t+1} = f(R_t, t, r_t) = r_t = \frac{\text{nr of positive reports in day } t}{\text{total nr of reports from day } t};$$

$$d(sla_t, R_t) = \begin{cases} \text{price}(sla_t) & \text{if } \text{price}(sla_t) \leq p(R_t); \\ 0 & \text{otherwise;} \end{cases}$$

Given that clients report the truth, and that “enough” clients submit feedback in day t , by the law of large numbers, the value of r_t is equal to q_t , the production quality chosen by the provider for day t . However, mistakes in reporting, as well as environmental factors influencing the experience of clients introduce a noise in the value of r_t . To account for this influence, we assume that the value of r_t is normally distributed around q_t with the variance σ . σ can be approximated by the system designer from previous experience. As long as the noise variance is small enough, it doesn’t influence the properties of the mechanism.

Proposition 1 establishes that the market is socially efficient when daily discount factors are greater than a given threshold:

PROPOSITION 1. *If $\delta > \frac{2C_{i_\theta^*}(\theta)}{P_{i_\theta^*} + P_{i_\theta^*+1}}$, (where δ is the daily discount factor, and i_θ^* is index corresponding to the efficient quality level of the provider of type θ) the above reputation mechanism has a Nash equilibrium in which the clients always purchase the service, the provider always exerts the socially efficient effort level and the service level agreement indicates the price $p(R_t)$ in every day t .*

PROOF. See appendix. \square

4. DISCUSSION AND FUTURE WORK

The mechanism presented above describes a way of generating SLAs based on reputation such that the market functions efficiently (i.e. providers exert the socially efficient effort level, and clients buy the service). Possible applications of such a mechanism are numerous, therefore this section discusses practical issues which might occur in the implementation of such a mechanism.

The mechanism has very little memory requirements. The reputation of a provider “survives” only one day, and therefore does not need to be protected in a usually volatile p2p

system. The mechanism is also resistant to failures: a certain fraction of the votes can be lost without greatly impacting the performance of the mechanism⁴.

The implementation of the reputation mechanism does not have to be centralized (as might be deduced from our presentation). Semi-centralized (where replicated, specialized agents implement parallel reputation mechanisms) or fully decentralized (where feedback is broadcasted and privately processed by each peer) systems are perfectly viable. The essence of the mechanism is that it allows rational agents (clients and providers) to act only based on local information (the feedback from the previous day). This model is perfectly compatible with a P2P setting.

Reporting mistakes and interferences from the environment introduce a normally distributed noise which affects the reputation of providers. In the proof of Proposition 1 we had to link the magnitude of the noise with the length of the shortest linearized interval. This assumption introduces a tradeoff between the requirements on the noise, and the precision of the cost and price function approximations. The more precise is the linear approximation, the smaller will be the length of the shortest linear interval, and therefore the smaller the noise is required to be.

One important assumption that we make when proving Proposition 1 is that clients report the truth. In real applications feedback can be distorted by the strategic interests of the reporters. Providing feedback may be expensive, and conflicts of interest could make the report unreliable. Fortunately, incentive-compatible schemes (i.e. they make it rational for clients to report the truth) exist when the report of one agent can be statistically correlated with reports coming from other agents [12, 11]. These schemes are not usually applicable when the provider behaves strategically. However, the equilibrium strategy of our reputation mechanism allows the use of such schemes (i.e. providers always exert the same effort level and therefore reports of different agents are correlated).

Another problem is the resistance to collusion. Returning agents could collude and submit negative feedback in order to pay smaller prices in the following day. While the solution to this problem is far from clear, we argue that such coalitions are not stable when services are scarce. There are two reasons behind this assertion. First, colluding agents know that the service is actually worth more than advertised. In order to secure the access to the service, they will afford to pay more than the price dictated by the reputation mechanism. These actions are visible to the other clients, and will soon raise the price of the service to the correct one. Second, colluding clients risk taking providers out of business. As services are already scarce, bankrupting providers is likely to increase the prices, not decrease them. This question needs however a deeper investigation in our future work.

We leave it up to the mechanism designer to estimate the physical length of one “day”. Longer periods of time might violate the assumption that the provider cannot change its behavior during one day. Shorter periods on the other hand, might not allow to gather enough binary reports in order to have a reliable estimation of the production quality level. In the same time, for shorter days, the impendency assumption on failures can be violated. At a macroscopic level, failures

⁴assuming that failures are sufficiently small and uncorrelated with the value of the report

are correlated: the failure of the present invocation most likely attracts the failure of the next invocation. Only when averaged over longer periods of time failures can be regarded as random.

The equilibrium strategy enounced by Proposition 1 is probably not unique. Other Nash equilibria might exist which do not have the desired properties. As future work, we’ll search and try to eliminate these “undesired” equilibrium points.

One final remark raises the question of human rationality. The desired equilibrium behavior can emerge only if clients and providers try to maximize their revenues. This is not always true with human users. However, in the future online market driven by agents and web services, the human intervention will be minimized, and rational strategies will predominate. This is the kind of environment for which we target our work.

5. RELATED WORK

Examples of computational trust mechanisms based on reputation are numerous, ranging from mechanisms based on direct interactions (in [4] agents learn to trust each other by keeping track of past interactions) to complex social networks [14] in which agents ask and give recommendations to their peers. Centralized implementations as well as peer-to-peer based systems [1, 9, 6] have been investigated.

Our work is closest to [8], [5], [7] and [3].

[5] considers exchanges of goods for money and proves that a market in which agents are trusted to the degree they deserve to be trusted is equally efficient as a market with complete trustworthiness. By scaling the amount of the traded product, the authors prove that it is possible to make it rational for sellers to truthfully declare their trustworthiness. Truthful declaration of one’s trustworthiness eliminates the need of reputation mechanisms and significantly reduces the cost of trust management. The difference from our work is that [5] considers only sellers of one given type and that it requires perfect knowledge about the seller’s cost function.

For e-Bay-like auctions, the Goodwill Hunting (GwH) reputation mechanism [7] provides a way in which the sellers can be made indifferent between lying or truthfully declaring the quality of the good offered for sale. The price of an item sold in day t depends on all items sold before and all previous feedback submitted. While our mechanism and the GwH mechanism provide similar guarantees, the settings in which the mechanisms should be used are different.

[8] addresses the same problem in a setting in which the provider can exert only two effort levels. The author makes a complete analysis of such a scenario and describes a family of reputation mechanisms which can be used to make the market efficient. However, the limitation to only two possible effort levels makes this work inapplicable in most service-oriented computing contexts.

Finally, [3] presents optimal pricing policies for service providers using recommender systems. While reputation mechanisms allows clients to differentiate between service providers, recommender systems allow clients to identify the best service based on their preferences (and feedback of previous clients with similar preferences). While such policies can be successfully used in contexts in which different services are provided by the same (or possible more) seller(s), our mechanism is applicable in a context in which the same service is offered by multiple providers.

6. CONCLUSION

This paper describes how a simple reputation mechanism allows a p2p service-oriented market to function efficiently. Reputation information is computed using an average-based aggregation rule, and drives service providers of different types to exert the socially efficient effort levels. The basic principle behind the mechanism is not social exclusion (separate and exclude “untrustworthy” agents) but rather incentive-compatible service level agreements which trade quality for price.

We have also discussed practical issues arising from the implementation of our mechanism in real applications. Acknowledging the limitations of our model, we believe that such mechanisms can bring significant improvements in decentralized service-oriented computing systems.

7. REFERENCES

- [1] K. Aberer and Z. Despotovic. Managing Trust in a Peer-2-Peer Information System. In *Proceedings of the Ninth International Conference on Information and Knowledge Management (CIKM)*, 2001.
- [2] G. A. Akerlof. The market for ‘lemons’: Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, 84(3):488–500, 1970.
- [3] D. Bergemann and D. Ozmen. Optimal Pricing Policy with Recommender Systems. In *Proceedings of P2PEcon 2004*, Harvard, Cambridge, 2004.
- [4] A. Birk. Learning to Trust. In R. Falcone, M. Singh, and Y.-H. Tan, editors, *Trust in Cyber-societies*, volume LNAI 2246, pages 133–144. Springer-Verlag, Berlin Heidelberg, 2001.
- [5] S. Braynov and T. Sandholm. Incentive Compatible Mechanism for Trust Revelation. In *Proceedings of the AAMAS*, Bologna, Italy, 2002.
- [6] S. Buchegger and J. L. Boudec. A Robust Reputation System for P2P and Mobile Ad-hoc Networks. In *Proceedings of P2PEcon 2004*, Harvard, Cambridge, 2004.
- [7] C. Dellarocas. Goodwill Hunting: An Economically Efficient Online Feedback. In J. Padget and et al., editors, *Agent-Mediated Electronic Commerce IV. Designing Mechanisms and Systems*, volume LNCS 2531, pages 238–252. Springer Verlag, 2002.
- [8] C. Dellarocas. Sanctioning Reputation Mechanisms in Online Trading Environments with Moral Hazard. MIT Sloan Working Paper #4297-03, 2004.
- [9] Z. Despotovic and K. Aberer. A Probabilistic Approach to Predict Peers’ Performance in P2P Networks. In *Eighth International Workshop on Cooperative Information Agents, CIA 2004*, Erfurt, Germany, 2004.
- [10] M. N. Huhns and M. P. Singh. Service-Oriented Computing: Key Concepts and Principles. *IEEE Internet Computing*, 9(1):75–81, 2005.
- [11] R. Jurca and B. Faltings. An Incentive-Compatible Reputation Mechanism. In *Proceedings of the IEEE Conference on E-Commerce*, Newport Beach, CA, USA, 2003.
- [12] N. Miller, P. Resnick, and R. Zeckhauser. Eliciting Informative Feedback: The Peer-Prediction Method. Forthcoming in *Management Science*, 2005.

- [13] N. Sokey and R. Lucas. *Recursive Methods in Economic Dynamics*. Harvard University Press, 1989.
- [14] B. Yu and M. Singh. Detecting Deception in Reputation Management. In *Proceedings of the AAMAS*, Melbourne, Australia, 2003.

APPENDIX

A. PROOF OF PROPOSITION 1

Assume that clients follow the recommendation of the reputation mechanism (i.e. buy only if the price specified by the SLA is smaller or equal to $p(R_t)$). A rational provider maximizes its revenue by specifying in the SLA a price exactly equal to $p(R_t)$. The continuation payoff (from, and including day t) of the service provider thus is:

$$\begin{aligned} V(R_t, \theta) &= \max_{q_t} (p(R_t) - c(q_t, \theta)) + \delta E_{R_{t+1}} [V(R_{t+1}, \theta)]; \\ &= \max_{q_t} (p(R_t) - c(q_t, \theta)) + \delta E_{r_t} [V(f(R_t, t, r_t), \theta)]; \\ &= \max_{q_t} (p(R_t) - c(q_t, \theta)) + \delta E_{r_t} [V(r_t, \theta)]; \end{aligned}$$

where $E_x[\cdot]$ denotes the expected value with respect to possible values of x .

Given that r_t is normally distributed around q_t , we have $r_t = q_t + z$ where z is the noise, normally distributed around 0 with variance σ . Therefore:

$$\begin{aligned} V(R_t, \theta) &= \max_{q_t} (p(R_t) - c(q_t, \theta)) \\ &\quad + \delta \int_{\mathcal{Z}} V(q_t + z, \theta) \cdot \pi(z) dz; \end{aligned} \quad (1)$$

where $\pi : \mathcal{Z} \rightarrow \mathbb{R}$ is the probability distribution function of the noise z (i.e. the normal prob. distribution function).

Let q^* the optimal control satisfying equation (1). We then have:

$$\begin{aligned} (p(R_t) - c(q^*, \theta)) + \delta \int_{\mathcal{Z}} V(q^* + z, \theta) \cdot \pi(z) dz &\leq \\ (p(R_t) - c(q^* + \delta q, \theta)) + \delta \int_{\mathcal{Z}} V(q^* + \delta q + z, \theta) \cdot \pi(z) dz; \end{aligned} \quad (2)$$

for all possible deviations δq . For piecewise linear price and cost functions, equation (2) can be rewritten as:

$$-c_1(q^*, \theta) \cdot \delta q + \delta \int_{\mathcal{Z}} V_1(q^* + z, \theta) \delta q \cdot \pi(z) dz \leq 0; \quad (3)$$

$V_1(R_t, \theta) = \frac{\partial (p(R_t) - c(q_t, \theta))}{\partial R_t} = p_1(R_t)$ (see [13], page 85), for all types, θ . By ignoring the noise outside the 3σ interval and by assuming that the noise is small enough (i.e. $3\sigma < (q_{i+1} - q_i)$ for all i), we have:

$$\int_{\mathcal{Z}} V_1(q_i + z, \theta) \pi(z) dz = (P_i + P_{i+1})/2 \quad (4)$$

for all points q_i situated at one end of the linear intervals.

q_{i^*} is the optimal control of (3), if and only if:

$$C_{i^*}(\theta) - \delta \cdot \frac{P_{i^*} + P_{i^*+1}}{2} \leq 0 \quad \text{if } \delta q < 0; \quad (5)$$

$$-C_{i^*+1}(\theta) + \delta \cdot \frac{P_{i^*} + P_{i^*+1}}{2} \leq 0 \quad \text{if } \delta q > 0; \quad (6)$$

For risk-averse and risk-neutral clients, $P_{i_\theta}^* \leq P_{i_\theta+1}^*$. Since $P_{i_\theta}^* > C_{i_\theta}^*(\theta)$ and $P_{i_\theta+1}^* < C_{i_\theta+1}^*(\theta)$ (these conditions hold in the point $q_{i_\theta}^*$), and by considering the hypothesis, both (5) and (6) are straightforward to prove.

For risk-seeking clients, $P_{i_\theta}^* > P_{i_\theta+1}^*$. In this case, a supplementary assumption is needed to verify (6), i.e.:

$$\delta < \frac{2C_{i_\theta+1}^*(\theta)}{P_{i_\theta}^* + P_{i_\theta+1}^*} \quad (7)$$

This condition is likely to be inactive for most scenarios. The condition:

$$\frac{2C_{i_\theta+1}^*(\theta)}{P_{i_\theta}^* + P_{i_\theta+1}^*} < 1 \quad (8)$$

is satisfied only if the price change from $P_{i_\theta}^*$ to $P_{i_\theta+1}^*$ is too abrupt. In such cases, the provider has the incentive to provide a quality that is slightly lower than the optimal one, since the positive noise in the reporting brings only a smaller marginal profit.

On the other hand, given that the provider always produces at the same quality level (i.e. the efficient one), and that the SLA specifies a price equal to $p(R_t)$, all clients have the incentive to buy (i.e. accept the recommendation of the reputation mechanism).

As a consequence, the reputation mechanism described in Section 3 has a Nash equilibrium in which providers exert the socially efficient effort, the SLA indicates a price equal to $p(R_t)$ and clients follow the recommendation of the mechanism. Q.E.D.