Incentive-Compatibility, Individual-Rationality and Fairness for Quality of Service Claims

Andrew Byde
Digital Media Systems Laboratory
HP Laboratories Bristol
HPL-2006-27
February 27, 2006*

QoS, services, trust, pricing

This document discusses extensions to the pricing structures presented in [2] for the provision of IT services that ensure trust without requiring repeated interactions between service providers and users. We extend [2] by providing theorems demonstrating the existence of incentive-compatible mechanisms for a broad class of expected costs to the user, and also introduce a notion of fairness with which to constrain the space of acceptable mechanisms.

* Internal Accession Date Only
[2]http://www.hpl.hp.com/research/idl/papers/trust/index.html
Huberman, B. et al Ensuring Trust in One Time Exchanges: Solving the Qos Problem

Approved for External Publication

Incentive-Compatibility, Individual-Rationality and Fairness for Quality of Service Claims

Andrew Byde

February 3, 2006

1 Abstract

This document discusses extensions to the pricing structures presented in [2] for the provision of IT services that ensure trust without requiring repeated interactions between service providers and users. We extend [2] by providing theorems demonstrating the existence of incentive-compatible mechanisms for a broad class of expected costs to the user, and also introduce a notion of fairness with which to constrain the space of acceptable mechanisms.

2 Basic Model

In the basic model a single user and service provider come together to agree terms for delivery of a service at some time in the future, which we will refer to as period 2, to distinguish it from period 1, in which the agreement is formed. The provider has a certain probability q – referred to as his QoS – of being able to provide the service in period 2, a claim about which, q', he makes to the user to encourage the user to enter into a service delivery agreement. The form of the agreement is a payment f_1 from the user to the provider in case the service is successfully delivered, and a payment f_0 in case the service is not delivered for some reason.

The goal of this paper is to describe constraints on suitable payment functions f_i, which constitute a *mechanism*, such that

1. The mechanism is *incentive-compatible*. This means that the provider has an incentive to act truthfully – in particular to report his true QoS, $q' = q$.
2. The mechanism is *individually-rational*[1]. This means that both the user and the provider are expected to gain higher utility from participating in the mechanism than from avoiding it.

[1]Note that in [2] incentive-compatibility is referred to as *truth-telling*, while individual-rationality is referred to as *incentive-compatibility*. The usage here is standard in the Game Theory literature [1].

To investigate these properties in more detail we need utility models for each party and a cost model for the provider. In the basic model we assume that attempting to deliver service in period 2 costs the provider a constant amount c irrespective of q. The motivation for this is that the provider will need to make certain capital investments in period 1 in order to be able to provide service in period 2; whether he will in fact be able to deliver service might depend on many factors, such as over-booking of the resources, or failure of service infrastructure.

The user is assumed to have value v for receiving the service, value 0 for not receiving the service. Both parties are assumed to be risk-neutral rational utility maximizers. For a provider with true QoS q, the expected return on reporting q' is $EP(q, q') = EC(q, q') - c$, where $EC(q, q')$ is the expected cost to the user, $EC(q, q') = qf_1(q') + (1 - q)f_0(q')$. For the user, utility maximization simply means choosing to participate when the expected return on doing so assuming that the provider is telling the truth $EU(q) = vq - EC(q, q)$, is positive.

We can now examine incentive-compatibility and individual-rationality in more detail.

2.1 Incentive-Compatibility

The incentive-compatibility criterion requires that the provider truthfully reveal his quality of service. There are two possible ways for the provider to lie. Most obviously he can attempt to provide service but claim QoS q' other than his true QoS q. This will be irrational if

$$q \neq q' \Longrightarrow EC(q, q') < TC(q) = EC(q, q). \tag{1}$$

Given that this inequality holds he might still have an incentive to claim that he is attempting to provide service at quality level q while in fact making no effort to do so, thus guaranteeing failure but avoiding his costs c. This will be irrational if

$$TC(q) > c + f_0 \tag{2}$$

The second of these incentive-compatibility criteria, which was not referred to in [2] could be seen either as a truth-telling (incentive-compatibility) issue, or a participation (individual-rationality) issue. In

3 demand

when we discuss a more complex model, we discuss the equivalent of (2) in the section on individual-rationality.

For $q \in (0, 1)$ the incentive-compatibility condition (1) for differentiable functions (f_0, f_1) implies at first order that

$$0 = \big(\partial_{q'} EC\big)(q, q) = q\, f_1'(q) + (1 - q) f_0'(q) \tag{3}$$

Taking the derivative with respect to q of the right-hand side of (3) gives

$$\begin{aligned} 0 &= q\, f_1'' + f_1' + (1 - q) f_0'' - f_0' \\ &= \big(\partial^2_{q'} EC\big)(q, q) + f_1' - f_0' \end{aligned}$$

It follows that if $f_1 - f_0$ is monotonically increasing, then the second derivative with respect to q' of $EC(q, q')$ will be strictly negative when evaluated at $q = q'$, which is sufficient for the mechanism to be strictly incentive-compatible. The monotonicity of $f_1 - f_0$ is a realistic business constraint: it is the difference in payment between delivery and non-delivery, which one would expect to be higher for services that are claimed to be more reliable.

Taking the derivative of the user's expected cost,

$$\begin{aligned} TC'(q) &= q\,f_1' + f_1 + (1-q)f_0' - f_0 \\ &= f_1 - f_0. \end{aligned} \tag{4}$$

It is clear from the previous discussion that TC is an increasing convex function. Indeed it is easy to see that an incentive-compatible mechanism is completely determined by the expected cost function. By rearranging the definition of TC and (4) we get

$$\begin{aligned} f_0 &= TC - q\,TC', \\ f_1 &= TC + (1-q)TC', \end{aligned} \tag{5}$$

Theorem 3.1 *Any incentive-compatible mechanism* (f_1, f_0) *for which* $f_1 - f_0$ *is increasing has convex increasing expected cost function* $TC : [0,1] \to \mathbb{R}$*; conversely any convex increasing differentiable function* $TC : [0,1] \to \mathbb{R}$ *is the expected cost function of an incentive-compatible mechanism as defined by (5), such that* $f_1 - f_0$ *is increasing.*

Proof. Easy. □

3.1 Individual-Rationality

Given that a mechanism is incentive-compatible, it will be individually-rational given that

$$\text{user}: \quad TC(q) < vq \tag{6}$$

$$\text{provider}: \quad TC(q) > c \tag{7}$$

To be precise, these inequalities define *strong* incentive-compatibility and individual-rationality; *weak* versions of these properties have weak inequalities $\leq$ and $\geq$ in place of strict inequalities $<$ and $>$.

These equations obviously show that no incentive-compatible mechanism can be individually-rational for $q < c/v$. From a business point of view this makes sense since services that have a fixed cost to attempt to provide but which might nevertheless fail with high probability are clearly not good business.

Obvious choices given knowledge of either c or v would seem to be $TC(q) = c$ or $TC(q) = qv$ respectively. These are optimal mechanisms for the user and provider respectively, since they allocate all of the surplus to the party in question. They are only weakly individually-rational (and $TC(q) = c$ is only weakly incentive-compatible), but clearly for small $\epsilon > 0$ the mechanisms

$TC(q) = c + \epsilon$ and $TC(q) = q(v - \epsilon)$ are strictly individually-rational compatible and have asymptotically identical properties. An unbounded example is $TC(q) = k(1 - q)^{-n}$, $k, n > 0$.

4 Variable Effort

In this section we consider an augmented model in which the provider's costs are variable depending on an effort level that also affects QoS. In a sense this was already the case in section 2: we have to introduce individual-rationality constraint (2) to ensure that when the provider claims he is attempting to provide service he will actually do so, rather than simply keeping the default payment f_0. Now we are imagining that the space of actions is richer; for example if costs are associated with buying equipment that will be used to deliver a service in period 2, there might be different quality levels, higher quality equipment being more expensive but more reliable.

We consider a model in which the user selects a QoS q', and the incentive-compatibility criteria is that given the payment schedule being considered, either the provider decides not to participate at all, or adjusts his true QoS q to be q'.

4.1 Continuous QoS choice

If the provider has arbitrary choice of q and cost function $c(q)$, then the first-order incentive-compatibility condition is

$$0 = \big(\partial_q EC\big)(q, q) = f_1 - f_0 - c',$$

so that the difference between delivery and non-delivery is exactly the marginal cost at the QoS selected by the user. An incentive-compatible mechanism therefore communicates a great deal of information about the provider's business, and is probably not realistic in practice, especially since the user will not in general have knowledge of a provider's costs (and thus will not be able to verify independently that a mechanism is incentive-compatible). Where providers retain control over the likely quality of service it is unlikely that one-time exchanges can be guaranteed to be incentive-compatible.

5 Uncertain Demand

In this section we consider a generalization of the basic protocol in which the user's demand at period 2 is uncertain. In particular we suppose that the user has probability p of requiring the service in period 2, a claim for which, p' he reveals to the provider in period 1 after the provider has revealed his QoS q'.

5.1 Revelation Protocol

In period 2 we can hypothesise two different protocols for true supply and demand revelation:

1. The provider reveals his ability to deliver before the user reveals his need for the service. This model is appropriate if the sources of delivery risk are tied to over-booking: in this case the provider wants to decide first which of his users will be permitted to use the service, in order to avoid the chance that the users might coordinate their demand so as to force the provider into paying excessive fines.

2. The user reveals his demand before the provider reveals his ability to deliver. This model is appropriate if the sources of delivery risk are inherent to the service itself – such as random equipment failure – and therefore *cannot* be decided prior to delivery.

Notice that in the first case, if the provider is unable to deliver, we would expect a payment independent of the user's demand, which need not be revealed; in the second case, if the user does not want the service, we would expect a payment independent of the provider's ability to deliver, which need not be revealed. These constraints will be referred to as Independence from Irrelevant Alternatives for the provider and user: IIAp and IIAu.

5.2 Notation

The expected cost to the user EC now depends in general on all four variables (p, p', q, q'). A *mechanism* will be a quadruple of functions dependent on the claimed demand and QoS levels:

$$EC(p, p', q, q') = \alpha(p', q') + p\beta(p', q') + q\gamma(p', q') + pq\delta(p', q') \tag{8}$$

Equivalently any function $EC : [0, 1]^4 \mapsto \mathbb{R}$ that is linear in p and q is also a mechanism, and has a decomposition of the form (8), which we will refer to as its *standard-form* representation.

Considering all mutually exclusive events gives another way to decompose a mechanism $EC(p, p', q, q')$, similar to section 2:

$$\begin{aligned} EC(p, p', q, q') \quad = \quad & f_{00}(p', q') \times (1 - p)(1 - q) \\ & + f_{10}(p', q') \times p(1 - q) \\ & + f_{01}(p', q') \times (1 - p)q \\ & + f_{11}(p', q') \times pq \end{aligned} \tag{9}$$

A mechanism written in this form will be described as being in *event*-form. Each term represents a payment or penalty depending on whether the service is deliverable or not, and whether the user wants it or not: f_{11} is the payment the user makes to the provider for the service if all goes well; f_{10} is the penalty the provider pays the user for failing; f_{01} is the penalty the user pays for demanding a service un-necessarily, and f_{00} is a mutually acceptable payment from one to the other in case they both fail.

Converting between standard and event forms is simple algebra:

$$\begin{array}{rcl} \alpha &=& f_{00} \\ \beta &=& f_{10} - f_{00} \\ \gamma &=& f_{01} - f_{00} \\ \delta &=& f_{11} - f_{10} - f_{01} + f_{00} \end{array} \tag{10}$$

$$\begin{array}{rcl} f_{00} &=& \alpha \\ f_{10} &=& \alpha + \beta \\ f_{01} &=& \alpha + \gamma \\ f_{11} &=& \alpha + \beta + \gamma + \delta \end{array} \tag{11}$$

In standard- and event-forms the IIA constraints clearly translate into:

1. IIAp. The provider reveals first, so if the provider cannot supply the service, the user's payments are the same irrespective of his need: $f_{00} = f_{10}$, i.e. $\beta = 0$.
2. IIAu. The user reveals first, so if a user does not need the service, his payments are the same irrespective of the provider's ability to deliver: $f_{00} = f_{01}$, i.e. $\gamma = 0$.

The user arguably only obtains value if he wants the service and it can be delivered:

$$EU(p, p', q, q') = pqv - EC(p, p', q, q')$$

And as before we assume constant costs of service provision attempts, regardless of whether in period 2 the service can in fact be delivered.

$$EP(p, p', q, q') = EC(p, p', q, q') - c$$

The value of EC assuming that both the service provider and user are truth-telling is referred to as TC: $TC(p, q) = EC(p, p, q, q)$.

5.3 Incentive-Compatibility

The first-order incentive-compatibility condition for the user, who claims p' after the provider has claimed q', and must assume that the provider is telling the truth, is

$$0 = \partial_{p'} EC(p, p', q, q)\Big|_{p'=p} = \alpha_p + p\beta_p + q\gamma_p + pq\delta_p. \tag{12}$$

For the provider, assuming that the user will go on to truthfully report p, we require

$$0 = \partial_{q'} EC(p, p, q, q')\Big|_{q'=q} = \alpha_q + p\beta_q + q\gamma_q + pq\delta_q. \tag{13}$$

Note that whereas (12) could consistently hold for some but not all pairs (p, q), in which case the mechanism would only be incentive-compatible for those (p, q), in order for the mechanism to be incentive-compatible for the provider at QoS q, (13) must hold for every p, since the provider supplies q' first, and thus does not yet know p. In general we will work with mechanisms that are incentive-compatible for all values of (p, q), so this is not an issue.

As before we examine partial derivatives of (12) and (13) to derive second-order inequalities. Taking the partial derivative of (12) with respect to p we get

$$\begin{aligned} 0 &= \alpha_{pp} + p\beta_{pp} + q\gamma_{pp} + pq\delta_{pp} \\ &\quad + \beta_p + q\delta_p \end{aligned}$$

This implies that

$$\begin{aligned} \partial^2_{p'} EC(p, p', q, q')\Big|_{p'=p} &= \alpha_{pp} + p\beta_{pp} + q\gamma_{pp} + pq\delta_{pp} \\ &= -(\beta_p + q\delta_p) \end{aligned}$$

A similar equation holds for (13). The sufficient second-order conditions for incentive-compatibility, $\partial^2_{p'} EC < 0$ and $\partial^2_{q'} EC > 0$ are therefore

$$0 > \beta_p + q\delta_p \tag{14}$$

$$0 < \gamma_q + p\delta_q \tag{15}$$

The inequality directions are opposite because the user wants to *minimize* expected costs, whereas the provider wants to *maximize* them.

Taking the first partial derivatives of the expected cost function assuming first-order incentive-compatibility (12) and (13) gives analogous equations

$$\partial_p TC(p, q) = \beta + q\delta \tag{16}$$

$$\partial_q TC(p, q) = \gamma + p\delta \tag{17}$$

These equations in conjunction with the definition of $TC(p, q)$ are three equations in four unknowns, which allows us to solve for any three of the unknowns in terms of the fourth and TC. If we also assume that one of the IIA constraints is enforced, then since these constraints are equivalent to either $\beta = 0$ (provider first), or $\gamma = 0$ (user first), we will have a full set of equations that are uniquely solvable for α, β, γ and δ in terms of the truthful expected cost function TC.

For IIAu ($\gamma = 0$) we get

$$\begin{aligned} \alpha &= TC - p\partial_p TC, \\ \beta &= \partial_p TC - \frac{q}{p}\partial_q TC, \\ \delta &= \frac{1}{p}\partial_q TC, \end{aligned} \tag{18}$$

while for IIAp ($\beta = 0$) the solution is an analogous expression with p and q swapped:

$$\begin{aligned} \alpha &= TC - q\partial_q TC, \\ \gamma &= \partial_q TC - \frac{p}{q}\partial_p TC, \\ \delta &= \frac{1}{q}\partial_p TC, \end{aligned} \tag{19}$$

Theorem 5.1 *If $TC : [0,1]^2 \to \mathbb{R}$ is a twice differentiable function with $TC_{pp} > 0$ and $TC_{qq} < 0$, then $EC(p, p', q, q')$ defined by (8), where α, β and δ are defined in terms of $TC(p, q)$ by (19), and $\gamma = 0$ is an incentive-compatible mechanism with expected costs to the user of $TC(p, q)$ that satisfies IIAu.*

Likewise under the same conditions on TC, if α, γ and δ are defined by (18) and $\beta = 0$, then EC is an incentive-compatible mechanism with expected costs TC satisfying IIAp.

Proof. It is sufficient for $EC(p, p', q, q')$ to be incentive-compatible that it satisfies the first- and second-order conditions (12) and (14) for the user and (13) and (15) for the provider. We shall demonstrate the user case for IIAu: the other three cases are nearly identical. We can calculate directly that

$$\begin{aligned} \alpha_p + p\beta_p + q\gamma_p + pq\delta_p &= TC_p - pTC_{pp} - TC_p \\ &\quad + p\Big(TC_{pp} - \tfrac{q}{p}TC_{pq} + \tfrac{q}{p^2}TC_q\Big) \\ &\quad +0 \\ &\quad + pq\Big(\tfrac{1}{p}TC_{pq} - \tfrac{1}{p^2}TC_q\Big) \\ &= 0, \end{aligned}$$

thus proving (12). A similar calculation proves (13). To check that the second order condition is satisfied, notice that from (16) (which only requires (12) and (13)), $TC_{pp} = \beta_p + q\delta_p$, so that $TC_{pp} > 0$ implies $\beta_p + q\delta_p > 0$, which is (14). □

5.4 Fairness

Although it is not necessary to constrain the space of mechanisms any further in order to ensure that all parties report truthfully, it is still possible to construct mechanisms that are unlikely to be "realistic" in the sense that they would not be acceptable in a real market.

The following four constraints define a *fair* mechanism.

1. $f_{11} \geq f_{10}$, which says that the provider should receive at least as much money for delivering a demanded service as he gets for not delivering it.
2. $f_{11} \geq f_{01}$, which says that the user should pay at least as much for a service he actually uses as for one he does not use.
3. $f_{01} \geq f_{00}$, which says that if the service is not demanded, then the provider should not be penalized for being able to deliver it.
4. $f_{00} \geq f_{10}$, which says that if the provider cannot deliver the service, then the penalty he pays should be at least as big when the service is actually in demand as if it isn't.

It can be shown using either (18) or (19) that the four fairness constraints in event-form are equivalent under IIAu or IIAp to either

$$0 \leq TC_q \quad \textit{and} \quad (q-1)TC_q \leq pTC_p \leq qTC_q, \tag{20}$$

or

$$0 \leq pTC_p \leq qTC_q, \tag{21}$$

respectively.

5.5 Individual-Rationality

As before there are two classes of individual-rationality constraints. The first set is that the expected costs of each party from following the mechanism should be positive. For an incentive-compatible mechanism specified as above via TC and δ, this is simply

$$\begin{aligned} EU(p,q) &= pqv - TC(p,q) \geq 0, \\ EP(p,q) &= TC(p,q) - c \geq 0. \end{aligned}$$

We must also consider the possibility that the provider will claim to be participating in the mechanism while in fact making no effort towards service provision. For the provider to have no incentive to do this we require

$$TC(p,q) - c \geq pf_{10} + (1-p)f_{00},$$

which becomes

$$TC(p,q) - c \geq f_{00}$$

for fair mechanisms (for which $f_{10} \geq f_{00}$). These together give the individual-rationality conditions

$$c + \max(0, f_{00}) \leq TC(p,q) \leq pqv. \tag{22}$$

These inequalities define weak individual-rationality; strong forms are obtained by making the inequalities strict. As before the boundary functions $TC(p,q) = c$ and $TC(p,q) = pqv$ are weakly incentive-compatible, and weakly individually-rational in the maximal feasible region $pq \geq c/v$. Strongly individually rational functions can be constructed by perturbing either of these trivial solutions. For example:

$$TC(p,q) = c + k_p(1-p^2) + k_q q^2$$

is strongly incentive-compatible, and strongly individually-rational for $pq > c/v + \epsilon$ whenever $0 < k_p, k_q < \epsilon/2v$.

For more general costs $c(p,q)$ the second of these functions $TC(p,q) = pqv$ is in general incentive-compatible.

5.6 Variable effort

In [2] provider costs are proportional to p, but likelihood of service delivery is not. On the one hand the former indicates that the service provider is scaling his effort, as captured by the cost of preparing to deliver service, in proportion to the user's reported likelihood of using the service. On the other hand the likelihood of delivering the service q is independent of this effort. This inconsistency indicates that a more sophisticated model of the relationship between the provider's cost of attempting to provide service and his likelihood of being able to do so is waranted. However, as seen in section 4, any effort to do so is likely doomed.

5.7 Summary

Combining the results of the previous sections we see that any function $TC : [0,1] \to \mathbb{R}$ that is both convex in p and concave in q gives rise to an incentive-compatible mechanism via either (18) or (19) depending on whether the user reveals his demand before the provider reveals his ability to deliver or vice-versa. In the former case, the mechanism thus constructed is fair if (20) holds,

$$0 \leq TC_q \quad and \quad (q-1)TC_q \leq pTC_p \leq qTC_q,$$

and furthermore is individually-rational if

$$c + \max(0, TC - pTC_p) \leq TC \leq pqv.$$

In the latter case, the mechanism is fair if (21) holds,

$$0 \leq pTC_p \leq qTC_q,$$

and furthermore is individually-rational if

$$c + \max(0, TC - qTC_q) \leq TC \leq pqv,$$

References

[1] D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, Oct. 1991.

[2] B. Huberman, F. Wu, and L. Zhang. Ensuring trust in one time exchanges: Solving the QoS problem, 2006. http://www.hpl.hp.com/research/idl/papers/trust/index.html.