

Eliciting Honest Reputation Feedback in a Markov Setting

Jens Witkowski

Department of Computer Science

Albert-Ludwigs-Universität

Freiburg, Germany

witkowski@informatik.uni-freiburg.de

Abstract

Recently, online reputation mechanisms have been proposed that reward agents for honest feedback about products and services with fixed quality. Many real-world settings, however, are inherently dynamic. As an example, consider a web service that wishes to publish the expected download speed of a file mirrored on different server sites. In contrast to the models of Miller, Resnick and Zeckhauser and of Jurca and Faltings, the quality of the service (e. g., a server's available bandwidth) changes over time and future agents are solely interested in the *present* quality levels. We show that hidden Markov models (HMM) provide natural generalizations of these static models and design a payment scheme that elicits honest reports from the agents after they have experienced the quality of the service.

1 Introduction

Online reputation mechanisms are a prominent way to establish trust and cooperation in anonymous online interactions [e. g., Dellarocas, 2006]. Signaling reputation mechanisms applied at opinion forums, in particular, have become a popular source of information. Customers of products and services give feedback regarding the quality they experienced and the publication of this previously undisclosed information allows future customers to make better-informed decisions. From a game-theoretic point of view, however, most of these mechanisms are problematic. Two aspects of the feedback part stand out in particular. The first is under-provision. The customers are usually required to register an account and fill out forms describing their experiences. This process is time consuming and rational agents will only invest the effort of giving feedback when remunerated appropriately. The second issue is honesty. External interests (i. e. biases towards dishonest reporting) come from a variety of motivations. Imagine two companies competing for the same group of customers. Both have incentives to badmouth their competitor, praise their own products or pay the rating agents to do so. Moreover, both positive and negative externalities are wide-spread. That is, an agent's utility of a good changes if other agents consume it as well. For example, the utility of an agent using a

voice-over-IP service is higher the more agents she can call with it. Honest reputation feedback is thus crucial to incorporate into the design of the mechanism. This task is difficult because it is not clear how to decide whether a given feedback is honest. In contrast to prediction markets, for example, where a publicly observable event eventually materializes [e. g., Wolfers and Zitzewitz, 2004], the inherent quality of a good is never revealed. Thus, the designer of the mechanism has to find other ways to condition the payments.

One solution is provided by Miller, Resnick and Zeckhauser [2005] (henceforth, MRZ). They compare the quality reports of two agents about the same good with one another and apply strictly proper scoring rules [e. g., Cooke, 1991] to compute a payment scheme that makes honest reporting a Nash equilibrium. Jurca and Faltings [2006] (henceforth JF) study a largely similar setting but use automated mechanism design [Conitzer and Sandholm, 2002] to compute a budget-optimal payment scheme. Furthermore, they developed numerous extensions to the base model, such as incorporating collusion resistance [Jurca and Faltings, 2007].

The mechanisms of MRZ and JF, however, critically depend on the good's quality to stay fixed, while many real-world settings are inherently dynamic. Applying the static mechanisms to dynamic settings is problematic for two reasons: first, the payment scheme is rendered useless if the agents know about the dynamic nature of the process, as by watching the type belief updates they can derive the posted signals and use them to learn the parameters of the process. Second, the mechanism's type architecture is only capable of publishing the *average* abilities whereas the users are interested in its *present* quality.

The rest of the paper is organized as follows. In the next section we present the setting. Section 3 describes the proposed reputation mechanism that generalizes the fixed architecture of MRZ and JF to hidden Markov models. Motivated by the example of Section 4, Section 5 shows how the acquisition of signals can be used to make honest reporting the unique Nash equilibrium. Section 6 provides experimental results and Section 7 concludes with an outlook.

2 The Setting

A sequence of agents experiences the same product or service and its quality (henceforth its *type*) is drawn out of a finite set of possible types $\Theta = \{\theta_1, \dots, \theta_{|\Theta|}\}$. All agents share a

common prior belief $Pr(\theta)$ that the product is of type θ with $\sum_{\theta \in \Theta} Pr(\theta) = 1$ and $Pr(\theta) > 0$ for all $\theta \in \Theta$.

The quality observations by the agents are noisy, so that after experiencing the product, a buying agent does not know with certainty the product's actual type. Instead, she receives a signal drawn out of a set of signals $S = \{s_1, \dots, s_M\}$. Let O^i denote the signal received by agent i and let $f(s_m|\theta) = Pr(O^i = s_m|\theta)$ be the probability that agent i receives the signal $s_m \in S$ given that the product is of type $\theta \in \Theta$. These signal emissions again constitute a probability distribution: $\sum_{m=1}^M f(s_m|\theta) = 1 \quad \forall \theta \in \Theta$. We assume that different types generate different conditional signal distributions and that all $f(s_m|\theta)$ are common knowledge.

In order to incorporate dynamic settings, we introduce a common-knowledge transition matrix that models quality changes as a Markov process (MP). Together with the noisy perception of the signals, the resulting structure is a discrete-time hidden Markov model (HMM) with a finite set of states. We use an extended definition of HMMs and allow both *null transitions* and *multiple observations*. That is, in each time step t , there can be more than one or no agent receiving a signal. The transition matrix $\mathbf{P}$ is given with the problem definition and stays the same for all t (i.e. the MP is time-homogeneous). Given type θ_i , the probability to go to type θ_j in the next time step is $Pr(\theta_j^{t+1}|\theta_i^t)$. This value is stored in column i and row j so that $\mathbf{P}$ is a so-called *left* stochastic matrix (i.e. the *columns* sum up to 1). Left transition matrices allow for a simple way to determine the probability of the type vector in a certain time step. We demand that $\mathbf{P}$ is reversible and that at least two columns are different from each other.

We will allow the mechanism (henceforth, *center*) to pay agents for their feedback and we assume quasilinear utilities. For example, electronic market sites could give away rebates on future sales. However, payments are not necessarily monetary as long as agents have utility for them. Let C^i be the costs reflecting agent i 's time and effort required for the rating process and let $\Delta^i(s_j, s_h)$ be the external benefit agent i could gain by falsely announcing signal s_h instead of signal s_j (the one actually received). We assume an upper bound on C and $\Delta(s_j, s_h)$. This way the center does not require knowledge on individual agent's preferences.

3 A Reputation Mechanism for Markov Settings

The reputation mechanism consists of three parts: the payment scheme, the reference reporter choice rule and the rule for publication of updated type beliefs. Before elucidating on these parts, we describe the general procedure and present the agents' options: first, an agent buys or requests a product or service. After experiencing its quality, the center asks the agent for feedback regarding the signal she perceived. As this information is private to the agent, there are three basic alternatives: the agent can choose to report the signal actually received, she can lie (i.e. report some other signal $s \neq O^i$) or she can not report at all. If the agent has chosen to report a signal, it is compared to that of another agent $r(i)$, called the reference reporter. The point in time at which the updated

type beliefs (taking into account the announced signal) are published depends on the rule that is used to choose an agent's reference reporter. With exception of the very first agent who is rated against her successor, we will rate an agent against her predecessor. This rule has two advantages: first, the center can pay the reporting agent right away (again, with exception of the very first agent) and, second, right after the announcement it can update and publish the type beliefs considering every signal report but the very last. This quick release of the updated beliefs is generally important but especially so in dynamic contexts where information out-dates.

Let $a^i = (a_1^i, \dots, a_M^i)$ be the reporting strategy of agent i , such that she reports signal $a_j^i \in S$ if she received s_j . The honest strategy is $\bar{a} = (s_1, \dots, s_M)$, i.e. always reporting the signal received. As becomes clear from the definition of the agent strategies, we assume they are independent of both the product and the time step t . This assumption is reasonable if the reputation mechanism is located at an intermediary, such as a booking site. Here, agents can neither lie about the product nor about the time at which they consumed it since the center already knows this from the booking data. It is, however, possible that the good's consumption is postponed or brought forward if that benefits the agents in the rating process. While this has no impact on the truthfulness of the mechanism, it may result in (arguably small) inefficiencies. Depending on the application it may be required to construct an online mechanism and we leave this to future work.

3.1 Probability Computations

The central idea of comparing two signal reports is that knowing one of the received signals should tell you something about the other. This concept is called *stochastic relevance*.

Definition 1. Random variable O^i is stochastically relevant for random variable $O^{r(i)}$ iff the distribution of $O^{r(i)}$ conditional on O^i is different for different realizations of O^i .

For the fixed setting, MRZ prove that combinations of $Pr(\theta)$ and $f(\cdot|\cdot)$ that fail stochastic relevance occur with probability 0 (i.e. have Lebesgue measure 0). As $\mathbf{P}$ is required to have at least two columns that are different from each other, this readily extends to the Markov setting. In fact, the added belief perturbations give the center more power to avert these combinations by using another rating agent which then results in a different scheme. We thus assume stochastic relevance holds.

Without loss of generality, let $r(i)$ receive her signal at time t_1 while agent i receives her signal at t_2 . Let $s_k^{t_1}$ and $s_j^{t_2}$ denote the signals received by $r(i)$ and agent i , respectively. The probability that $r(i)$ received $s_k^{t_1}$ given i received $s_j^{t_2}$ is:

$$g(s_k^{t_1}|s_j^{t_2}) = Pr(O^{r(i)} = s_k^{t_1} | O^i = s_j^{t_2}). \quad (1)$$

This can be written as:

$$g(s_k^{t_1}|s_j^{t_2}) = \sum_{l=1}^{|\Theta|} Pr(s_k^{t_1}|\theta_l^{t_1}) \cdot Pr(\theta_l^{t_1}|s_j^{t_2}). \quad (2)$$

$Pr(s_k^{t_1}|\theta_l^{t_1})$ can be simplified to $f(s_k|\theta_l)$ as the probability of a signal given a certain type is independent of when it is received. For reasons of clarity, we slightly abuse the notation and add superscripts to $f(\cdot|\cdot)$ in following equations.

Applying Bayes' Theorem to Eq. 2 we obtain:

$$Pr(\theta_l^{t_1} | s_j^{t_2}) = \frac{Pr(s_j^{t_2} | \theta_l^{t_1}) \cdot Pr(\theta_l^{t_1})}{Pr(s_j^{t_2})}. \quad (3)$$

Let

$$Pr(\theta) = (Pr(\theta_1), Pr(\theta_2), \dots, Pr(\theta_{|\Theta|}))^T \quad (4)$$

be the vector of prior type probabilities. As we know both the topology and the parameters of the HMM, calculating the entire probability vector $Pr(\theta^t)$ is straightforward:

$$Pr(\theta^t) = (Pr(\theta_1^t), \dots, Pr(\theta_{|\Theta|}^t))^T = \mathbf{P}^t \times Pr(\theta) \quad (5)$$

Please note that in the context of a matrix, the superscript denotes exponentiation: $\mathbf{P}^t = \underbrace{\mathbf{P} \times \dots \times \mathbf{P}}_t$.

Using Eq. 5, we obtain the signal probability required for Eq. 3:

$$Pr(s_j^{t_2}) = \sum_{l=1}^{|\Theta|} f(s_j^{t_2} | \theta_l^{t_2}) \cdot Pr(\theta_l^{t_2}). \quad (6)$$

The probability that agent i receives signal $s_j^{t_2}$ given the type was $\theta_l^{t_1}$ (also required for Eq. 3) can be computed as follows:

$$Pr(s_j^{t_2} | \theta_l^{t_1}) = \sum_{o=1}^{|\Theta|} f(s_j^{t_2} | \theta_o^{t_2}) \cdot Pr(\theta_o^{t_2} | \theta_l^{t_1}). \quad (7)$$

For the probability of a certain type at time t_2 knowing the type at time t_1 we need to distinguish two cases:

- $t_2 \geq t_1$ Here, a minor change of Eq. 5 is sufficient:

$$Pr(\theta^{t_2} | \theta_l^{t_1}) = \mathbf{P}^{t_2-t_1} \times (0, \dots, \theta_l^{t_1} = 1, \dots, 0)^T, \quad (8)$$

i. e. the l -th column of $\mathbf{P}^{t_2-t_1}$.

- $t_2 < t_1$ From Bayes' Theorem we know

$$Pr(\theta_o^{t_2} | \theta_l^{t_1}) = \frac{Pr(\theta_l^{t_1} | \theta_o^{t_2}) \cdot Pr(\theta_o^{t_2})}{Pr(\theta_l^{t_1})} \quad (9)$$

and $Pr(\theta_l^{t_1} | \theta_o^{t_2})$ is analogous to the $t_2 \geq t_1$ case.

3.2 The Payment Scheme

Let $\tau(a_j^i, a_k^{r(i)})$ be the payment that agent i receives if she announced a_j^i and the reference reporter announced $a_k^{r(i)}$. The expected payment to agent i given her received signal and given an honest report by $r(i)$ is:

$$E(a_j^i, s_j^{t_2}) = \sum_{k=1}^M g(s_k^{t_1} | s_j^{t_2}) \cdot \tau(a_j^i, a_k^{r(i)}). \quad (10)$$

Similar to JF, we formulate the payment scheme as a Linear Program (LP). Its constraints can be divided into two groups. The first group consists of the honesty constraints which require that the honest signal announcement by agent i is the single best response to an honest report by $r(i)$. For every possible signal observation $O^i = s_j \in S$, there exist $M - 1$ dishonest announcements $a_j^i \neq \bar{a}_j$. Given that the reference report is honest, we want the expected payment of an honest announcement by agent i to be larger than the expected payment of any other announcement. More accurately, incorporating external lying incentives, we want it to be larger by a margin greater than $\Delta(s_j, s_h)$:

$$\sum_{k=1}^M g(s_k^{t_1} | s_j^{t_2}) \cdot \tau(s_j, s_k) - \sum_{k=1}^M g(s_k^{t_1} | s_j^{t_2}) \cdot \tau(s_h, s_k) > \Delta(s_j, s_h) \\ \forall s_j, s_h \in S, s_j \neq s_h$$

The second group consist of the participation constraints. An agent will participate in the rating system if and only if she is remunerated with at least as much as the rating process costs her. As the agent's decision whether to participate in the rating is taken after experiencing the good (i. e. she knows her own signal) but without knowing the signals received by the other agents, *interim* individual rationality is appropriate [e. g., Parkes, 2001, p. 34f]:

$$\sum_{k=1}^M g(s_k^{t_1} | s_j^{t_2}) \cdot \tau(s_j, s_k) > C \quad \forall s_j \in S$$

In addition, we assume there is no possibility to withdraw credit from the agents, so that we require that all payments are non-negative. In order to find an assignment of $\tau(a^i, a^{r(i)})$ that minimizes the required budget, the objective function is the expected payment given a certain signal weighted with the signal's prior probability. Summarizing, the payment scheme formulated as an LP in standard form is:

$$\begin{aligned} \min \quad & \sum_{j=1}^M Pr(s_j^{t_2}) \left(\sum_{k=1}^M g(s_k^{t_1} | s_j^{t_2}) \cdot \tau(s_j^{t_2}, s_k^{t_1}) \right) \\ \text{s.t.} \quad & \sum_{k=1}^M g(s_k^{t_1} | s_j^{t_2}) (\tau(s_j^{t_2}, s_k^{t_1}) - \tau(s_h^{t_2}, s_k^{t_1})) > \Delta(s_j^{t_2}, s_h^{t_2}) \\ & \forall s_j^{t_2}, s_h^{t_2} \in S, s_j^{t_2} \neq s_h^{t_2} \\ & \sum_{k=1}^M g(s_k^{t_1} | s_j^{t_2}) \cdot \tau(s_j^{t_2}, s_k^{t_1}) > C \quad \forall s_j^{t_2} \in S \\ & \tau(s_j^{t_2}, s_k^{t_1}) \geq 0; \quad \forall s_j^{t_2}, s_k^{t_1} \in S \end{aligned}$$

Please note that the objective function uses the expected payment of the honest equilibrium since this is what the agents should play given the honesty constraints.

The payment scheme together with the rater choice rule we use induces a temporal order of extensive games with imperfect information. However, the fact that no agent knows the reported signal of her reference reporter make them equivalent to 2-player strategic games for which we can prove the following proposition.

Proposition 1. *Reporting honestly is a Nash equilibrium in every strategic game induced by an update of the type beliefs.*

Proof. The proof applies induction over the game order.

The basis: The first game is played by agent 1 and 2: agent 1 is rated against agent 2 and given the latter reports honestly, the LP's honesty constraints ensure that the honest report is a best response. Likewise, given an honest report by the first agent, the second agent's best response is reporting honestly.

Inductive step: Assume that reporting honestly is an equilibrium in the n -th game. In game $n + 1$, agent $n + 2$ is rated against agent $n + 1$ while the latter's payoff solely depends on the outcome of game n . That is, agent $n + 1$ is indifferent about the outcome of game $n + 1$. Since honest reporting is an equilibrium of game n , agent $n + 1$'s best response was playing honestly, so that honest play by agent $n + 2$ is a best response. Thus, reporting honestly is a Nash equilibrium of game $n + 1$. $\square$




Figure 1: Fixed quality setting with conditional signal probabilities.

3.3 Updating the Type Beliefs

Up to this point, we have constructed a payment scheme that induces agents to give honest feedback about the signals they receive. What is supposed to be published, though, are the present beliefs that the product has a certain type. We publish the type changes inherent in the MP in every time step using Equation 5. This is different from publishing information that is updated using signal observations (hence, *conditional updates*) for when performing the latter we have to ensure that the respective observations are no longer required as reference reports.

As before, let agent n denote the n -th agent to make a signal announcement. Slightly abusing the notation, let t_n and O^n denote her time step and signal, respectively. At t_1 there is no conditional belief update as there is no other announcement than that of agent 1 who is still needed as a reference reporter for agent 2. After announcement of O^n , however, we incorporate the signal observation at t_{n-1} as we no longer rely on agent $n-1$ as a reference reporter. This also means that every conditional belief update incorporates only a single observation. Thus, we compute the conditional type update and “overwrite” the old type priors using Equation 3:

$$Pr(\theta) := Pr(\theta^{t_{n-1}} | O^{n-1}) \quad \forall \theta \in \Theta$$

4 A Web Service Example

We present an example application from a computing context. Imagine a web service that wishes to publish the expected download speed of a file mirrored on different server sites. The users of the service have to choose one of these mirrors to download from. Clearly, this choice determines the download speed as the mirrors vary in both total bandwidth and work load. Every user wishes to download from the mirror that is expected to serve him with highest speed. Unfortunately, this information is only known to the mirrors and they have no interest in sharing it truthfully as they receive advertisement premiums for every access of their site. Users also have to enter a CAPTCHA before downloading which makes it costly for both the operator of the web service and the users to test the present speed of all mirrors beforehand.

Therefore, a reputation mechanism is located at the web service. By logging the user requests, the center can identify both when the download began and which mirror was chosen. The number of users that are directed to the servers by




Figure 2: The dynamic setting allows the underlying types to change over time. For reasons of clarity, the loops are left out.

		$r(i)$				$r(i)$	
		l	h			l	h
Agent i	l	1	0	Agent i	l	0.17	0
	h	0	0.63		h	0	0.14
Fixed				Markov			

Figure 3: The payment schemes of both settings at $t_2 = 0$ with C and $\Delta(s_j, s_h)$ ($h \neq j$) set to 0.1.

the web service is small compared to the servers’ total number of users, so that the impact of the web service’s agents on the servers’ behavior is negligible. Furthermore, we assume that the servers are programmed to serve at their highest possible speed given their bandwidth and work load. That is, the servers’ behavior is entirely stochastic.

Let us first consider the example of Figure 1. There are only two possible types, a good type G and a bad type B . Furthermore, there are only two possible signals, namely a high signal h and a low signal l . Types and signals correspond to the servers’ total and load-dependent bandwidths, respectively. The prior type probabilities are $Pr(G) = 0.7$ and $Pr(B) = 0.3$. All other probabilities are depicted.

The fixed setting is very noisy as the dynamics of the work load cannot be expressed properly. Alternatively, we can model the servers as the hidden Markov model depicted in Figure 2. The prior type probabilities are $Pr(B_l) = 0.18$, $Pr(B_h) = 0.12$, $Pr(G_l) = 0.14$ and $Pr(G_h) = 0.56$. As a server may not distribute its bandwidth entirely equally, it is still possible to receive a high signal given a bad state and vice versa. The parameters of this dynamic setting are chosen such that they yield the same prior signal distribution as the fixed setting of Figure 1, namely $Pr(l) = 0.371$ and $Pr(h) = 0.629$ (please note that the similar values given in Figure 1 are the *conditional* signal probabilities). Nonetheless, the computed payment schemes for $t_2 = 0$ are different (see Figure 3). The expected payments are 0.43 and 0.13 for the fixed and dynamic setting, respectively. Figure 4 shows that with larger t_2 the expected budget of the dynamic scheme converges to the budget of the fixed setting from below. Thus, with introduction of the Markov setting we can both gain expressiveness and further lower the required budget.




Figure 4: Expected budget for the fixed and dynamic setting.

5 Acquisition of Costly Signals

A drawback of the payment scheme from Section 3 is that the honest equilibrium is not unique. In fact, it can even be Pareto-dominated by a lying equilibrium (compare Figure 3). Unfortunately, as long as the mechanism solely depends on announcements by selfish agents, multiple equilibria are essentially unavoidable [Jurca and Faltings, 2005]. Depending on the setting, however, it can be possible for the center to acquire a costly signal himself. Take the web service example: clearly, the costs involved in a signal acquisition by the operator are higher than the reporting costs of an agent. After all, that is why we propose to use a reputation mechanism. Yet, while it is too costly for the operator to learn the speed of each server before *every* download, it might be affordable to acquire *some* signals that can then be used as reference reports. In fact, we will see that acquiring a *single* signal is sufficient to make honest reporting the unique Nash equilibrium in *every* game induced by a type belief update.

The naïve approach is to acquire a new signal for every agent, rate her against it and update the type beliefs. This is, however, very costly. A cheaper approach is to utilize an acquired signal similar to the publicly observable event in a prediction market: the center commits to an acquisition at time t and rates every agent up to this point against it. Yet, thereafter, the center has to acquire another signal as the types continue to change. Highlighting the importance of the reference reporter choice rule, we propose to acquire a signal before any agent announcement, rate the first agent against it and every subsequent agent against her predecessor. This also allows us to drop the assumption that $\mathbf{P}$ is reversible.

Proposition 2. *The acquisition of the first signal is sufficient to make honest reporting the unique Nash equilibrium of every strategic game induced by an update of the type beliefs.*

Proof. The proof is similar to that of Proposition 1 and also applies induction over the game order.

The basis: The first game is now played by the center and agent 1. Agent 1 is rated against the acquired signal which is honest by definition. The honesty constraints ensure that honest reporting is the single best response and thus that honest reporting is the unique Nash equilibrium.

M	time (in ms)	M	time (in ms)
2	2.36	12	28.47
4	3.20	14	44.14
6	5.94	16	68.89
8	10.06	18	100.94
10	17.36	20	137.64

Table 1: Average CPU time for computation of the payment scheme with different values of M .

The **inductive step** is the same as in Proposition 1 except that game n is assumed to have a *unique* equilibrium and that in game $n + 1$, agent $n + 1$ is rated against agent n . $\square$

Please note that it is also possible to rate every agent against the acquired signal directly. This would, however, bring about three other problems: first, it is harder to keep a single signal undisclosed if every agent learns it. Second, even myopic agents cannot be allowed to participate again while our proposed choice rule only prohibits successive ratings. Third, the required budget usually grows with the number of time steps that lie in between the two raters (compare next section).

6 Experimental results

If not stated otherwise, the parameters for the experiments are created as described in Appendix A.

6.1 Running Time

There are polynomial algorithms for solving LPs. In practice, however, the Simplex method with exponential worst-case running time usually performs better. Thus, in order to determine whether the reputation mechanism is feasible for real-world settings, we empirically evaluate it on a customary computer with a 1.6 GHz CPU.

Let $\Delta(t)$ denote $|t_2 - t_1|$. Table 1 shows the CPU time that is required for computation of the payment scheme for $\Delta(t) = 0$ and different values of M . The running time that comes with the MP is depicted in Figure 5. In addition to $\Delta(t)$, the runtime also grows with M . This is due to the matrix exponentiation algorithm we use whose running time is $O(|\Theta|^3 \cdot \log t)$. For larger $\Delta(t)$, the factor $|\Theta|^3 = M^3$ gets more influence although only multiplied by $\log t$. Yet, values of $\Delta(t)$ should rarely be higher than 10 so that the slowdown will make up less than 3% even for settings with large signal set. Taking into account the low numbers of Table 1, we believe computational complexity is not a limiting factor for application.

6.2 Expected Budget

In the example of Section 4 the fixed setting requires a higher budget than an equivalent dynamic setting. Hidden Markov models, however, are more expressive. In particular, it is possible that every type can be reached from every other type. Figure 6 shows the expected costs for different values of stochastic movement inherent in the fully connected $\mathbf{P}$ of the random setting. The more time passes between the two ratings and the larger ε , the higher are the center's costs. Note that this is inherent in the setting as both factors lead to more




Figure 5: Normalized running time depending on $\Delta(t)$ with $\Delta(t) = 0$ corresponding to 100%. For $M = 5$ we took 10000 settings.




Figure 6: Expected budget for different values of ε .

type perturbations through the MP. For larger $\Delta(t)$ the probability of a certain type at t_2 conditional on the types at t_1 becomes more alike as the MP converges to a stationary distribution and so do the signal posteriors $g(s_k^{t_1} | s_j^{t_2})$. As these make up the coefficients of the LP, the solver needs larger τ to separate the honest from the dishonest announcements.

7 Conclusions and Future Work

We have presented a reputation mechanism that elicits honest feedback in a Markov setting and requires a lower budget than the equivalent fixed setting. When solely relying on the announcements of selfish agents, multiple equilibria are unavoidable. We have shown how in settings where costly signals can be acquired, the acquisition of a single signal is sufficient to make honest reporting the unique Nash equilibrium of every induced game.

A limitation of the setting we studied, however, is the high amount of common knowledge. We believe it is an interesting question under which circumstances it is possible to truthfully elicit the signal observations and also use them to learn the probabilistic parameters. For restricted settings such as the example setting of Section 4 this should be possible: using

the fixed payment scheme, one can elicit the first k agents' signals and subsequently expand the model to the equivalent dynamic setting. With the structure and the prior beliefs for the ground types known, one could then learn the parameters with the Baum-Welch algorithm. The choice of k depends on the trade-off between budget and robustness. Please note that we only need to ensure that the center's model is better than that of future reporting agents which is further supported by only publishing the expected quality (e. g., the expected speed) instead of the entire distribution.

Acknowledgments

We thank Malte Helmert, Sven Seuken and three anonymous reviewers for very helpful comments on this work.

A Random Setting

Every type corresponds to a signal, i. e. $M = |\Theta|$. The default is $M = 5$. All $f(s_m | \theta_l)$ are set according to the rule:

$$f(s_m | \theta_l) = \begin{cases} 1 - \varepsilon & m = l \\ \varepsilon / (M - 1) & m \neq l \end{cases}$$

P is generated analogously:

$$Pr(\theta_j^{t+1} | \theta_i^t) = \begin{cases} 1 - \varepsilon & j = i \\ \varepsilon / (|\Theta| - 1) & j \neq i \end{cases}$$

The default is $\varepsilon = 10\%$. The type vector at $t = 0$ is uniformly distributed and we set $t_1 = 0$ with $t_2 \geq t_1$. All $\Delta(s_j, s_h)$ ($h \neq j$) are set to 0.15, C is set to 0.1 and we average over 1000 randomly generated settings.

References

- [Conitzer and Sandholm, 2002] Vincent Conitzer and Tuomas Sandholm. Complexity of Mechanism Design. In *Proceedings of the Uncertainty in Artificial Intelligence Conference*, 2002.
- [Cooke, 1991] Roger M. Cooke. *Experts in Uncertainty*. Oxford University Press, 1991.
- [Dellarocas, 2006] Chrysanthos Dellarocas. Reputation Mechanisms. In *Handbook on Information Systems and Economics*. Elsevier Publishing, 2006.
- [Jurca and Faltings, 2005] Radu Jurca and Boi Faltings. Enforcing Truthful Strategies in Incentive Compatible Reputation Mechanisms. In *Internet and Network Economics*. Springer-Verlag, 2005.
- [Jurca and Faltings, 2006] Radu Jurca and Boi Faltings. Minimum Payments that Reward Honest Reputation Feedback. In *Proceedings of the ACM Conference on Electronic Commerce*, 2006.
- [Jurca and Faltings, 2007] Radu Jurca and Boi Faltings. Collusion-resistant, Incentive-compatible Feedback Payments. In *Proceedings of the ACM Conference on Electronic Commerce*, 2007.
- [Miller et al., 2005] Nolan Miller, Paul Resnick, and Richard Zeckhauser. Eliciting Informative Feedback: The Peer-Prediction Method. *Management Science*, 2005.
- [Parkes, 2001] David Parkes. *Iterative Combinatorial Auctions: Achieving Economic and Computational Efficiency*. PhD thesis, University of Pennsylvania, 2001.
- [Wolfers and Zitzewitz, 2004] Justin Wolfers and Eric Zitzewitz. Prediction Markets. *Journal of Economic Perspectives*, 2004.