ALBERT-LUDWIGS-UNIVERSITÄT
FREIBURG
INSTITUT FÜR INFORMATIK

Arbeitsgruppe
Grundlagen der Künstlichen Intelligenz

Prof. Dr. Bernhard Nebel

Truthful Feedback for Reputation Mechanisms

Diplomarbeit

Jens Witkowski

ERKLÄRUNG

Hiermit erkläre ich, dass ich diese Arbeit selbständig verfasst und keine anderen als die angegebenen Quellen oder Hilfsmittel verwendet habe. Alle Stellen, die wörtlich oder sinngemäß aus veröffentlichten Schriften entnommen wurden, habe ich als solche kenntlich gemacht. Darüber hinaus erkläre ich, dass diese Arbeit nicht, auch nicht auszugsweise, bereits für eine andere Prüfung angefertigt wurde.

Freiburg, den 29.05.2009

Jens Witkowski

Acknowledgements

A special thanks goes to Malte Helmert for the excellent supervision of both my study project and this Diploma thesis. I want to thank him not only for his professional advice but for his positive attitude towards me, my work and ideas.

I also want to thank Prof. Dr. Bernhard Nebel for agreeing to be my supervisor and, in particular, for his support regarding the publication of my study project's results.

Finally, I want to thank Sven Seuken for proofreading and his valuable (and presumably truthful) feedback.

Contents

Chapter 1

Introduction

E-commerce web sites enjoy huge popularity. Virtually every type of good is traded online: books and electronic devices are ordered from Amazon[1], used items are traded on eBay[2] and, more recently, services such as translation and web design jobs are offered at service auction sites, such as Elance[3]. Almost every e-commerce site employs a so-called reputation mechanism. Reputation mechanisms collect ratings from market participants and make them publicly available. These ratings then allow other market participants to make better-informed choices.

1.1 Reputation Mechanisms

It is instructive to distinguish between different kinds of reputation mechanisms in accordance with the problems they address. Those that are employed by online opinion forums such as Amazon Reviews or Ciao[4] are built to eliminate asymmetric information while those at online auction sites are primarily intended to induce cooperation and trust between market participants. For a general overview of research on reputation mechanisms, see the excellent survey of Dellarocas [2006].

1.1.1 Adverse Selection

The dynamics entailed in markets with asymmetric information were first studied by Akerlof in his seminal paper "The Market for Lemons" [1970]. He presents the example of a market for used cars with two types of cars: cars of bad quality

---

[1] www.amazon.com

[2] www.ebay.com

[3] www.elance.com

[4] www.ciao.com

(which in America are called “lemons”) and cars of good quality. The seller of a car knows whether his[5] car is a lemon or not while potential buyers can only estimate the average quality of all cars taken together. A buyer’s valuation for a car, however, depends on its quality. As she cannot observe the quality of a specific car, all cars look the same to her and choosing one at random gives her average quality in expectation. Consequently, she refuses to pay more than the price that corresponds to average quality. The sellers anticipate this reasoning and while sellers that own a lemon would happily agree, those with a good car would run a loss selling it at the average price which is why they withdraw from the market. The buyers on their part anticipate the withdrawal of the good sellers and—now confronted with the reduced set of cars—lower the price they are willing to pay. Due to the hidden quality of the cars, all but the cars of worst quality are driven out of the market. This effect is called *adverse selection* and it is present in all markets where a product’s type, i. e. its inherent quality, is unknown to the customer before the purchase. The objective of a reputation mechanism in this setting is to reveal the inherent quality of the products to future customers.

1.1.2 Moral Hazard

The primary objective of reputation mechanisms at online auction sites, such as eBay, is different. Instead of providing information about the products that have been traded, the goal is to disseminate information about the trustworthiness of the market participants. The usual procedure at eBay is that the buyer of an auction transfers the payment to the seller before the latter sends the good. Once the seller has received the payment, however, he has an incentive to either not send the good at all or in a quality that is lower than the one he had advertised. The assumption is that for goods of low value, the probability of legal action is sufficiently small as the costs that are entailed for the buyer exceed the good’s value. Anticipating this, the buyer would not send the money in the first place, so that no trade takes place. The problem that eBay’s reputation mechanism addresses is therefore one of opportunistic behavior. The technical term for this type of market failure is *moral hazard* and the objective of a reputation mechanism situated in these settings is therefore to induce cooperation and trust between market participants. Since future buyers resort to the seller’s public feedback history, seller cooperation is achieved through the threat that non-cooperative behavior is sanctioned.

The central distinction between settings with pure *adverse selection* and settings with pure *moral hazard* is the type of seller behavior: in settings with

---

[5] We refer to sellers and buyers as male and female, respectively.

pure *adverse selection*, sellers differ in their abilities. That is, some sellers are of higher quality than others. In settings with pure *moral hazard*, all sellers are equally able but are tempted to exert an effort that is below their abilities because it involves smaller costs.

1.1.3 Mixed Settings

In addition to settings with pure *adverse selection* and pure *moral hazard*, many reputation mechanisms are situated in settings where both are present simultaneously. An example for such a reputation mechanism is the one employed by Elance. Elance may be regarded as the equivalent of eBay for services: potential customers can post a project and providers can place their bids on finishing it. Note that in contrast to eBay, Elance employs a *reverse* auction, i. e. the *sellers* are bidding. Another difference to eBay is that the buyers are not deemed to take the lowest offer but can consider both the posted bid and the published feedback about the seller. Consider the example of a web designer who offers his service via Elance. Clearly, different web designers have different abilities as virtually everybody with an Internet connection can open up an account and offer his services. Yet while unskilled web designers can only produce low quality, skilled web designers may choose between high and low quality and the production of low quality involves less time and money than the production of high quality. The objective of the reputation mechanism in mixed settings, such as Elance, is therefore both to reveal the seller's abilities and to induce cooperative behavior.

1.2 Truthful Feedback

A common feature of almost all reputation mechanisms in e-commerce environments is the dependency on honest buyer feedback. Most mechanisms in the literature simply *assume* that feedback is reported honestly. However, as the outcomes are private information of the agents, this assumption is rather strong. From a game-theoretic point of view, there are two issues in particular: the first is the agents' motivation to participate at all. The feedback procedure requires the user to register an account, to log in and to fill out forms describing the experiences. While this is time consuming and thus costly, the reported information benefits other customers but not the posting agent herself, so that standard economic theory predicts an under-provision of feedback. The second difficulty is honesty. External interests, i. e. biases towards dishonest reporting, come from a variety of motivations. Imagine, for example, two companies competing for the same group of customers. Either company has an incentive to badmouth its competitor, to praise its own products or to pay the rating agents to do so. Another potential reason for biased reports are externalities, i. e. an agent's utility for a good changes if other agents consume it as well. An example for a good with positive externalities is a voice-over-IP service for which an agent's utility is higher the more other agents she can call with it. The agent could therefore be tempted to report a quality level that is higher than the perceived quality in order to lure more agents into using the service. The analogous holds true for settings with negative externalities, such as a web service that evenly divides its bandwidth among its users. A particularly common issue at online auction sites with bi-directional feedback, where both the seller and the buyer can rate the transaction, is retaliatory feedback. Empirical data of eBay's pre-2008 reputation mechanism suggests that sellers waited until the buyers had posted their feedback and matched it thereafter. That is, sellers posted the same rating they had received from the buyers and, in particular, retaliated against negative feedback [Resnick and Zeckhauser, 2002; Bolton *et al.*, 2009]. The truthful elicitation of reputation feedback is thus crucial to incorporate into the design of a reputation mechanism.

For a product rating environment with pure adverse selection, there is a method to elicit truthful feedback due to Miller, Resnick and Zeckhauser [2005] (henceforth, MRZ). Their so-called "peer prediction method" pays a buyer for her feedback depending on the feedback that was given about the same product by another buyer. The intuition behind this method is that the quality experiences of two buyers that have bought the same product should be "essentially" identical. Differences in experiences may occur but can be captured in a noise parameter. Take a digital camera bought via Amazon as an example: while different customers may experience different quality due to noise in the production process or due to different opinions on what constitutes a "good" camera, all buyers receive the identical model. We elaborate on this method in Section 2.2.

The only work we are aware of that is addressing the elicitation of truthful feedback in moral hazard environments is that of Jurca and Faltings [2007b]. The mechanism they propose, however, is only applicable to settings where the buyer and the seller *frequently* interact with one another. It does not extend to the usual e-commerce setup, at online auctions for example, in which a buyer is interacting with a specific seller only once.

1.3 Contributions

The largest part of the literature studies moral hazard reputation mechanisms under the assumption that they are given honest feedback. The authors then usually study how to improve seller cooperation. For example, Dellarocas [2005] studies the optimal length of published feedback history with regard to seller cooperation.

In this thesis, we take the complementary view and assume that *we are given such a reputation mechanism that induces some degree of seller cooperation under the assumption that feedback is honest.* Using this reputation mechanism, we then study whether truthful feedback can be elicited. That is, we are using the reputation mechanism as a "black box" and investigate whether one can design a feedback mechanism for it. More specifically, we study whether the aforementioned "peer prediction method" can be modified for this purpose. It is important to note that in settings with moral hazard, the problem of inducing seller cooperation and that of truthful feedback elicitation can be viewed largely in isolation. In theory, it is thus perfectly possible that a setting allows for a reputation mechanism that induces full cooperation with assumed honest feedback while it is impossible to design a feedback mechanism for this setting, and vice versa. Please note that when the words *reputation mechanism* and *feedback mechanism* can be confused, we sometimes refer to the latter as a "feedback plug-in" indicating that it builds on an external reputation mechanism.

We show that in a pure moral hazard setting, there is no peer based feedback mechanism that elicits truthful feedback. For a mixed setting, however, we retrieve a positive result and construct a feedback mechanism that can be used as a plug-in by reputation mechanisms. Furthermore, the plug-in can be used for pure moral hazard settings with "cheating types", i. e. sellers who do not value the published feedback high enough so that they find it optimal to always cheat. Note that the mixed setting is a *strict* mixed setting, i. e. there needs to be adverse selection to escape the mentioned impossibility result. We also give a first result with regard to the elimination of non-truthful equilibria that are unavoidable in purely peer based feedback mechanisms. That is, we exemplify a method for a simple example setting with perfect monitoring to construct a payment scheme that is guaranteed to induce truthfulness once a single buyer is truthful.

Our experimental findings show that both computational complexity and expected budget are feasible in practice. Even for large signal sets with 30 signals, our feedback plug-in can be computed in less than 250 milliseconds.

1.4 Outline

The remainder of the thesis is organized as follows. In Chapter 2, we introduce the pure adverse selection setting which models the situation in online opinion forums such as Amazon Reviews. We furthermore elaborate on MRZ's feedback mechanism for this setting as well as its budget-optimal formulation as a Linear

Program which is due to Jurca and Faltings [2006]. In Chapter 3 we analyze the strategic implications of online auction settings with pure moral hazard and distinguish them from those at opinion forums. We prove that the pure adverse selection mechanism cannot be applied to these online auction sites and discuss possible escapes from this impossibility. In Chapter 4 we introduce a setting that combines the characteristics of the pure adverse selection setting from Chapter 2 with those of the pure moral hazard setting from Chapter 3. We retrieve a positive result and show how to construct a “feedback plug-in” that induces a truthful perfect Bayesian equilibrium. We experimentally evaluate this “feedback plug-in” in Chapter 5. In Chapter 6 we exemplify a method to construct a payment scheme that is guaranteed to induce truthfulness once a single buyer is truthful. Finally, we conclude with a brief summary of this thesis and an outlook on future research in Chapter 7.

Chapter 2

Pure Adverse Selection

Akerlof's lemon market that we introduced in Chapter 1 is similar, though not identical, to the setting we are faced with in online opinion forums, such as Amazon Reviews. The main difference is that in Akerlof's setting, every buyer experiences a *different* product while the feedback reports at opinion forums are written about *identical* products, such as the same digital camera. Nevertheless, adverse selection is present in the opinion forum setting as well: if buyers cannot distinguish between the quality levels of different products, a buyer's willingness to pay is independent of the quality which again drives out all products of higher value. A number of remedies for this situation have been developed. For example, the seller can give out a "money back guarantee" that allows the buyer to experience the product without the risk of losing her money. Another method is to leave it to a trusted third party to thoroughly test the product and provide information about its quality. In this chapter, we take yet another approach and resolve pure adverse selection through the publication of previous buyers' experiences. More specifically, we will ask the buyers of a product to report their perceived quality and present a mechanism that induces them to be honest.

The remainder of this chapter is organized as follows: in Section 2.1 we describe the pure adverse selection setting as studied by MRZ. In Section 2.2 we present their reputation mechanism for this setting, which is based on the comparison of two quality reports. While MRZ pay the reporting agents according to a strictly proper scoring rule, we also introduce the Linear Program formulation of the mechanism which is due to Jurca and Faltings [2006]. In Section 2.3 we give a small example motivated by the setting at Amazon Reviews and in Section 2.4 we conclude this chapter with a discussion on other application areas of the mechanism.

2.1 The Setting

A group of agents experiences the same product or service. Its quality (henceforth its *type*) is drawn out of a finite set of possible types:

$$\Theta = \{\theta_1, \dots, \theta_{|\Theta|}\} \tag{2.1}$$

Once determined by "nature", a product's type is fixed. All agents share a common prior belief $Pr(\theta)$ that the product is of type θ with

$$\sum_{\theta \in \Theta} Pr(\theta) = 1 \tag{2.2}$$

while $Pr(\theta) > 0$ for all $\theta \in \Theta$. Please note that the type of a product is never revealed.

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 privately receives a signal drawn out of a set of signals:

$$S = \{s_1, \dots, s_M\}. \tag{2.3}$$

Let s^i denote the signal received by agent i and let

$$f(s_m | \theta) = Pr(s^i = s_m | \theta) \tag{2.4}$$

be the probability that agent i receives the signal $s_m \in S$ if the product is of type $\theta \in \Theta$. The signal observations again constitute a probability distribution, i. e. for all $\theta \in \Theta$ we have:

$$\sum_{m=1}^{M} f(s_m | \theta) = 1. \tag{2.5}$$

We assume that different types generate different conditional signal distributions and that all $f(s_m | \theta)$ are common knowledge.

We allow the mechanism (henceforth, the *center*) to pay agents for their feedback. For example, electronic market sites could give away rebates on future sales. However, payments do not have to be monetary as long as the agents associate utility with them. In particular, we assume that utilities are linear in payments. Please see Section 2.4 for examples with non-monetary payments.

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). As mentioned in the introduction of this thesis, these benefits can come from a variety

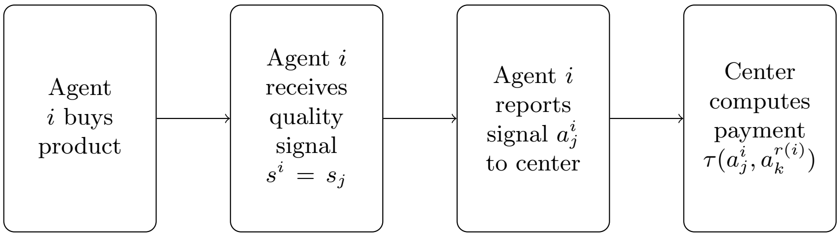

Figure 2.1: The reporting procedure for a single agent i in the pure adverse selection setting.

of backgrounds such as badmouthing competitors. We assume upper bounds

$$C = \max_i C^i \tag{2.6}$$

and

$$\Delta(s_j, s_h) = \max_i \Delta^i(s_j, s_h) \tag{2.7}$$

on the participation costs and external lying benefits, respectively. This way the center does not require knowledge on an individual agent's preferences. Note that by definition $\Delta(s_m, s_m) = 0$ for all $s_m \in S$.

2.2 The Peer Prediction Method

The problem we are facing is difficult because the type, i. e. the "ground truth" is never revealed. That makes the setting different to prediction markets, for example, where a publicly observable event eventually materializes [e. g., Pennock and Sami, 2007; Wolfers and Zitzewitz, 2004]. However, we can use the report of one agent and compare it to that of another agent.

At the core of the reputation mechanism by MRZ is a so-called payment scheme that determines a payment to the reporting agent depending on her signal report and the signal report by another agent, called the reference reporter. Before we elucidate on the construction of the payment scheme, we describe the general procedure (compare Figure 2.1): first, an agent buys a product and receives a quality signal. Thereafter, the center asks the agent for feedback regarding the quality signal she received. As this is private information of the agent, there are three alternatives:

- the agent can choose to report the signal she actually received,
- she can lie, i. e. report some other signal $s_h \neq s^i$,
- or she can choose not to report at all.

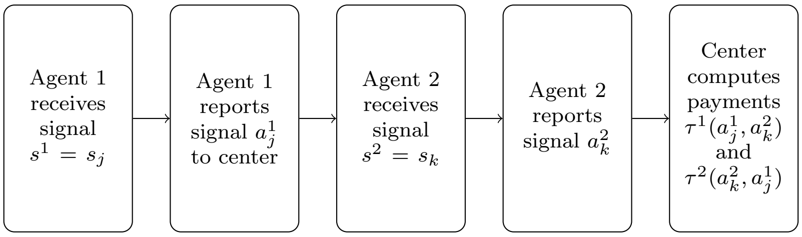

Figure 2.2: The procedure of the pure adverse selection reporting game.

For the moment, we assume that the agent has incentives to report a signal and let

$$a^i = (a_1^i, \ldots, a_M^i) \tag{2.8}$$

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, i. e. always reporting the signal received, is:

$$\bar{a} = (s_1, \ldots, s_M). \tag{2.9}$$

As becomes clear from the definition of the agent strategies, we assume they are independent of the product. That is, an agent can lie about the quality signal she perceived but she cannot announce reports for products she has not bought. This assumption is reasonable if the reputation mechanism is located at an intermediary, such as a booking site. Here, agents cannot claim to have experienced a product that in fact they have not since the center knows from the booking data which product was bought by whom.

After the agent has reported a signal to the center, it is compared to the report of her reference reporter $r(i)$. Let $s^i = s_j$ and $s^{r(i)} = s_k$ be the signals received by agent i and $r(i)$. For her report, she receives a payment

$$\tau(a_j^i, a_k^{r(i)}) \tag{2.10}$$

that depends on both her own report a_j^i and the report of her reference reporter $a_k^{r(i)}$. The central idea of comparing two signal reports is that knowing the signal received by agent i should tell you something about the signal received by her reference reporter. This assumption is called *stochastic relevance* [Miller *et al.*, 2006].

Definition 1. Random variable s^i is *stochastically relevant* for random variable $s^{r(i)}$ if and only if the distribution of $s^{r(i)}$ conditional on s^i is different for all different realizations of s^i.

That is, s^i is stochastically relevant for $s^{r(i)}$ if and only if for any two distinct

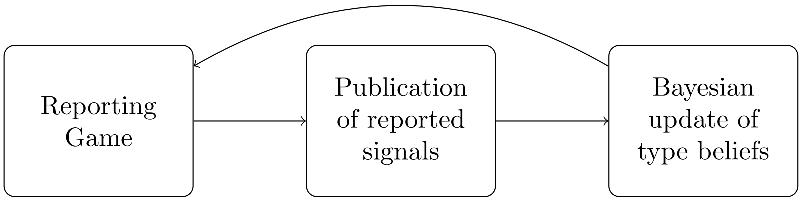

Figure 2.3: The overall procedure of the pure adverse selection setting.

realizations of s^i, call them s_j and s_h, there exists at least one realization of $s^{r(i)}$, call it s_k, such that

$$Pr(s_k|s_j) \neq Pr(s_k|s_h). \tag{2.11}$$

MRZ prove that combinations of $Pr(\theta)$ and $f(\cdot|\cdot)$ that fail stochastic relevance have Lebesgue measure 0 [2005]. This means that small belief perturbations make stochastically irrelevant belief combinations stochastically relevant again. For the remainder of this chapter, we therefore assume that stochastic relevance holds.

The point in time at which the reported signals can be published depends on the rule that is used to choose an agent's reference reporter. MRZ discuss several rules and their advantages but we constrain ourselves to the basic rule which collects feedback from 2 subsequent agents, rates them against each other and then publishes the two reports. For the game between the next two agents, the type beliefs are computed via Bayesian updating (Equation 2.14 on p. 12). See Figure 2.3 for the overall procedure.

2.2.1 Belief Computations

Let again s_k and s_j denote the signals received by $r(i)$ and agent i, respectively. The probability that $r(i)$ received s_k given i received s_j is denoted by:

$$g(s_k|\, s_j) = Pr(s^{r(i)} = s_k|\, s^i = s_j). \tag{2.12}$$

We transform Equation 2.12 until we are left with values that are given with the setting. The first step is to expand the conditional probability into a summation:

$$g(s_k|\, s_j) = \sum_{\theta \in \Theta} f(s_k|\, \theta) \cdot Pr(\theta|\, s^i = s_j). \tag{2.13}$$

Applying Bayes' Theorem to the second part of the summation in Equation 2.13 yields:

$$Pr(\theta|\, s^i = s_j) = \frac{f(s_j|\, \theta) \cdot Pr(\theta)}{Pr(s^i = s_j)}. \tag{2.14}$$

The denominator of Equation 2.14 is the prior signal probability which can be computed with Equation 2.15:

$$Pr(s^i = s_j) = \sum_{\theta \in \Theta} f(s_j|\, \theta) \ \cdot \ Pr(\theta). \tag{2.15}$$

With these, we have all necessary calculations to compute $g(s_k|\, s_j)$ for all $s_k,\ s_j \in S$. We can therefore now turn to the design of the payment scheme.

2.2.2 Proper Scoring Rule Formulation

The payment scheme is an $M \times M$ matrix that is used to determine the payment to an agent depending on the signal announcements by her and her reference agent. As before, let s_j and s_k denote the signal received by agent i and agent $r(i)$, respectively, and let $\tau(a^i_j, a^{r(i)}_k)$ denote the payment agent i receives if she announces a^i_j and her reference reporter announces $a^{r(i)}_k$. We build the payment scheme with the assumption that $r(i)$ honestly reports her signal, i. e.

$$a^{r(i)}_k = s_k \tag{2.16}$$

and show how to find payments for agent i that make honest reporting her best response to the honest report by her reference agent. As the game is symmetric, honest reporting then becomes a best response to honest reporting which makes it a Nash Equilibrium. The expected payment to agent i given her received signal and given an honest report by $r(i)$ is:

$$E(a^i_j, s_j) = \sum_{k=1}^{M} g(s_k|\, s_j) \ \cdot \ \tau(a^i_j, a^{r(i)}_k). \tag{2.17}$$

The idea is that once agent i received a signal herself, she can update her probabilistic belief of $r(i)$'s signal using Equation 2.12 on p. 11.

The question now is how to elicit the agent's updated posterior belief and MRZ employ so-called *proper scoring rules* which can be used to truthfully elicit beliefs or forecasts from agents about the likelihood of mutually exclusive events [e. g., Savage, 1971; Cooke, 1991; Winkler *et al.*, 1996]. Once an event materializes ("happens"), its outcome is publicly observed and the agent receives a numerical score or payment that depends on her announced forecast. Scoring rules that are *strictly proper* are chosen such that the agent maximizes her expected score if and only if she honestly announces her distributional belief. The two following definitions formalize this.

Definition 2. A *scoring rule* is a function $R : \mathcal{P} \times X \rightarrow \mathbb{R}$ that assigns a numerical score to each pair (p, x), where p is a probability distribution and x is the event that eventually materializes.

Definition 3. A scoring rule is said to be *proper* if an agent is maximizing her expected score by truthfully announcing her belief $p \in \mathcal{P}$ and *strictly proper* if the truthful announcement is the only announcement maximizing her expected score.

MRZ apply strictly proper scoring rules to our reputation feedback setting. The event that agent i shall forecast is the signal received by her reference agent. Take a look at the space of updated posterior distributions: agent i receives exactly one out of M signals so that there are only M possible posterior distributions, $g(\cdot|\, s_j)$, all of which the center can compute with the setting's common knowledge data. So if we have a proper scoring rule $R(\cdot,\cdot)$ that truthfully elicits probabilistic beliefs, we can ask the agent to choose one out of these M possible posterior beliefs through the reporting of her signal. Note that, for the moment, we assume that there are no external benefits from lying.

Lemma 1. *$R(p, x) = \log_2(p_x)$ is a strictly proper scoring rule where p_x is the probability that was announced for the event that actually materializes.*

See the work by Savage [1971] for a more detailed discussion of proper scoring rules including a proof for Lemma 1. Note that we could use any strictly proper scoring rule to construct the payment scheme and use the logarithmic rule solely for its notational simplicity.

Proposition 2. *If we assign the payments according to*

$$\tau(a_j^i, a_k^{r(i)}) = \log_2(g(a_k^{r(i)}|a_j^i)), \tag{2.18}$$

honest reporting is a strict Nash Equilibrium in the simultaneous reporting game with no external benefits from lying.

Proof. What needs to be shown is that honest reporting by agent i is the sole best response to an honest report by $r(i)$, and vice versa. Given an honest report by agent $r(i)$, the expected payment is given by Equation 2.17 that—together with the strictly proper log scoring rule—is uniquely maximized by agent i's honest belief. It is thus the single best response to report the signal she received as the scoring rule is then applied to the updated posterior distribution. Due to the game's symmetry, this also holds for scoring agent $r(i)$. $\square$

What remains to be shown is how to incorporate the external benefits from lying as well as the participation constraints. As we assumed utilities that are linear in payments, linear transformations of the payments do not change the strategic properties of the reporting game. That is, we may apply affine transformations. How to proceed is described in the proof of Lemma 3 on p. 17.

2.2.3 Linear Program Formulation

Jurca and Faltings (henceforth JF) [2006] study a setting that is essentially the same as that of MRZ. The mechanism they propose also applies peer prediction, i. e. the use of reference reporters to score agents, but instead of using proper scoring rules, they formulate the payment scheme as a Linear Program (LP). The technique of formulating a mechanism design problem as an optimization problem is called Automated Mechanism Design [Conitzer and Sandholm, 2002; Sandholm, 2003].

The Linear Program formulation has several advantages over the formulation with proper scoring rules: first, the mechanism designer can focus on the formulation of the problem rather than implementing the solution. This lowers the probability of programming mistakes and allows for the convenient formulation of certain properties. In our setting, one can observe this, for example, when it comes to the formulation of the external benefits from lying for which the formulation is much more natural in the LP than it is for the proper scoring rules. An advantage with regard to the actual result is that the optimization problem implements the budget-optimal mechanism for a given problem.

The constraints of the LP 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 $s^i = s_j \in S$, there exist $M - 1$ dishonest announcements $a^i_j \neq \bar{a}_j$. 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 $\Delta(s_j, s_h)$:

$$\sum_{k=1}^{M} g(s_k|\, s_j) \cdot \tau(s_j, s_k) - \sum_{k=1}^{M} g(s_k|\, s_j) \cdot \tau(s_h, s_k) > \Delta(s_j, s_h) \qquad (2.19)$$
$$\forall s_j,\, s_h \in S,\, s_j \neq s_h$$

and, in a more compact notation:

$$\sum_{k=1}^{M} g(s_k|\, s_j)\,(\tau(s_j, s_k) - \tau(s_h, s_k)) > \Delta(s_j, s_h). \qquad \forall s_j,\, s_h \in S,\, s_j \neq s_h \tag{2.20}$$

The second group consists of the participation constraints, also called individual rationality (IR) constraints [e. g., Parkes, 2001, p. 34f]. A rational agent will only give feedback if she is remunerated with at least as much as the rating process costs her. Note that in order to avoid indifference between participation and absence we demand that, at the time of her participation decision, the agent receives an expected payment that is *higher* than C. From a strategic point of view, there are three possible times at which an agent could decide about her participation: the first is before she knows her signal. The expected payment at that stage is the a priori or ex ante payment which is why this type of IR constraint is called *ex ante* individual rationality. For our setting, however, this is not sufficient as the agents decide whether to participate only after they have experienced the product, i. e. after they received their signal. Therefore, participation needs to be better than non-participation given any of the M possible signal observations. This type of IR constraint is called *interim* individual rationality:

$$\sum_{k=1}^{M} g(s_k|\, s_j) \cdot \tau(s_j, s_k) > C, \quad \forall s_j \in S. \tag{2.21}$$

For reasons of completeness, let us mention the third type which is called *ex post* individual rationality and is applicable in settings where agents can withdraw from the mechanism once they have learned its decision. The Vickrey auction [1961], for example, satisfies ex post IR.

As mentioned earlier, the objective of the center is to minimize its required budget. Please note that technically we can only achieve a very good approximation of the required budget as we have strict inequalities in our constraints. The budget B is the expected payment given a certain signal weighted with the signal's prior probability:

$$B = \sum_{j=1}^{M} Pr(s_j) \left(\sum_{k=1}^{M} g(s_k|\, s_j) \,\cdot\, \tau(s_j, s_k) \right). \tag{2.22}$$

Note that given a feasible assignment, it is sufficient to solely regard the expected payment of the honest equilibrium since this is what the agents should play given the honesty constraints. An alternative way to come up with the same objective function is to apply basic probability theory: the probability that $\tau(s_j, s_k)$ is paid is $Pr(s_j \cap s_k)$ which is also the factor $\tau(s_j, s_k)$ is weighted with, so that one obtains:

$$B' = \sum_{j=1}^{M} \sum_{k=1}^{M} Pr(s_j \cap s_k) \, \cdot \, \tau(s_j, s_k). \tag{2.23}$$

From probability theory, we know that:

$$Pr(s_j \cap s_k) = Pr(s_k | \, s_j) \cdot Pr(s_j). \tag{2.24}$$

Inserting this into Equation 2.23, then yields:

$$\begin{aligned} B' &= \sum_{j=1}^{M} \sum_{k=1}^{M} Pr(s_k | \, s_j) \cdot Pr(s_j) \, \cdot \, \tau(s_j, s_k) \\ &= \sum_{j=1}^{M} Pr(s_j) \left(\sum_{k=1}^{M} g(s_k | \, s_j) \, \cdot \, \tau(s_j, s_k) \right) \\ &= B. \end{aligned} \tag{2.25}$$

Together with the assumption that there is no possibility to withdraw credit from the agents, so that all payments are non-negative, the summarized payment scheme formulated as an LP in standard form is LP 1. Note that we have to add a small $\epsilon > 0$ to the right side of both the honesty and participation constraints as the definition for Linear Programs does not include strict inequalities.

LP 1.

$$\min \quad B = \sum_{j=1}^{M} Pr(s_j) \left(\sum_{k=1}^{M} g(s_k | \, s_j) \, \cdot \, \tau(s_j, s_k) \right)$$

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

It is clear that LP 1 is bounded as all factors in the objective function are non-negative. To show that it is also feasible, we give a revised result from our earlier work [Witkowski, 2008] and begin with a lemma that reduces the feasibility of LP 1 to that of the following feasibility program (FP). Note that an FP is an LP without objective function. FP 1 consists of the left side of the honesty constraints with the external benefits from lying set to 0:

FP 1.

$$\sum_{k=1}^{M} g(s_k|\, s_j) \cdot \tau(s_j, s_k) - \sum_{k=1}^{M} g(s_k|\, s_j) \cdot \tau(s_h, s_k) \geq \epsilon$$

$$\forall s_j,\, s_h \in S,\, s_j \neq s_h$$

Lemma 3. *Feasibility of FP 1 implies feasibility of LP 1.*

Proof. We proceed by transforming a feasible solution of FP 1 into a feasible solution of LP 1 and begin with incorporating $\Delta(s_j, s_h)$. Let ϵ' denote the maximal of all possible ϵ which corresponds to FP 1's "honesty margin". Note that taking the maximal ϵ is not necessary but makes it easier to follow the reasoning. As expected utility is invariant to affine transformations, we can multiply $\tau(\cdot, \cdot)$ with a constant factor γ without changing the incentive properties. Let $\Delta' = \max_{j,h} \Delta(s_j, s_h)$ denote the maximal external lying benefit. Then we can choose $\gamma = \frac{\Delta'}{\epsilon'}$, so that after multiplying $\tau(\cdot, \cdot)$ with γ, the resulting feasibility program incorporates all external lying benefits.
If the minimal $\tau(s_j, s_k)$ is negative, let τ' denote its absolute value, i. e. it becomes positive, and 0 otherwise. Incorporating participation costs C and ensuring that all $\tau(s_j, s_k)$ are positive is done through the addition of $\tau' + C + \epsilon$ with $\epsilon > 0$ to all $\tau(s_j, s_k)$. As mentioned, expected utility is invariant to affine transformations, so that adding a constant does not change the agents' incentives. ☐

Proposition 4. *LP 1 is feasible.*

Proof. By Lemma 3 it is sufficient to restrict the feasibility analysis to FP 1. Observe that the two summations in FP 1 are the expected payments to agent i given she reports the honest signal s_j and given she reports some other signal s_h, respectively. If we choose $\tau(\cdot, \cdot)$ according to a strictly proper scoring rule, her expected payment is maximized through her choice for the honest signal's posterior distribution. Stochastic relevance implies that every possible signal observation s^i yields a *different* posterior distribution with regard to $s^{r(i)}$, so that choosing the honest signal's posterior distribution is *strictly* better than any other possible posterior distribution. ☐

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 which we call "reporting games" and for which the following proposition holds.

Proposition 5. *Reporting honestly is a Nash equilibrium in every reporting game induced by an update of the type beliefs.*

If LP 1 is feasible and bounded, the proof follows immediately from the design of the payment scheme.

Jurca and Faltings present numerous extensions to this base model. They show how to further lower the budget by using multiple reports and a filtering technique for reports that are false with high probability (while still paying these reports) [2006]. In order to incorporate prior beliefs that are slightly different from the center's, they built a mechanism robust to small changes in these beliefs [2007c]. An insightful presentation of the expressive abilities entailed in the LP formulation is the work on colluding agents and Sybil attacks, i. e. one agent controlling several accounts [2007a]. In our earlier work [2008; 2009] we show how to incorporate temporal changes of the product's type.

2.3 Example: Amazon

Consider the example of a digital camera bought from Amazon. Let there be two types of cameras, a good type G and a bad type B. Potential customers share a prior belief $Pr(\theta = G) = 0.7$ that the camera is of the good type. After purchase, the customers make pictures and play around with the camera through which they make a noisy observations of its type. Note that the source for this noise can either be due to stochastic variance in the production process as well as an uncertainty on the customer's side, e. g. a misunderstanding of the camera's manual that results in overexposed pictures. Let there thus be two signals, namely a high signal h and a low signal l. The belief that a good camera is followed by a high signal is $f(h|G) = 0.75$ whereas the belief that a bad camera is followed by a high signal is $f(h|B) = 0.15$. We assume that falsely announcing a high signal is costlier than falsely announcing a low signal: $0.6 = \Delta(l, h) > \Delta(h, l) = 0.3$. The idea behind this assumption is that after purchase, a customer does not want to admit having bought a bad camera, so that the threshold for biases in this direction are higher. The participation costs are set to $C = 0.5$.

The center now collects the signal reports of two buyers. In order to score these reports against one another, it needs to compute the reference reporter's signal posterior belief $g(\cdot|\cdot)$. For the respective equations, we refer to Section 2.2.1.

$$\begin{aligned} Pr(s^i = h) &= f(h|\,G) \cdot Pr(\theta = G) + f(h|B) \cdot Pr(\theta = B) \\ &= 0.75 \cdot 0.7 + 0.15 \cdot 0.3 \\ &= 0.57 \end{aligned}$$

$$\begin{aligned} Pr(s^i = l) &= 1 - Pr(s^i = h) \\ &= 0.43 \end{aligned}$$

Bayes' Theorem gives us the probabilities for types conditional on signals:

$$\begin{aligned} Pr(\theta = G|s^i = h) &= \frac{f(h|G) \cdot Pr(\theta = G)}{Pr(s^i = h)} \\ &= \frac{0.75 \cdot 0.7}{0.57} \\ &\simeq 0.92 \end{aligned}$$

$$\begin{aligned} Pr(\theta = B|s^i = h) &= 1 - Pr(\theta = G|s^i = h) \\ &\simeq 0.08 \end{aligned}$$

$$\begin{aligned} Pr(\theta = G|s^i = l) &= \frac{f(l|G) \cdot Pr(\theta = G)}{Pr(s^i = l)} \\ &= \frac{0.25 \cdot 0.7}{0.43} \\ &\simeq 0.41 \end{aligned}$$

$$\begin{aligned} Pr(\theta = B|s^i = l) &= 1 - Pr(\theta = G|s^i = l) \\ &\simeq 0.59 \end{aligned}$$

Eventually, we can calculate the probability of a rater $r(i)$ receiving a certain signal conditional on the signal rater i received. Note that we give higher precision values in brackets:

$$\begin{aligned} g(h|\,l) &= Pr(s^{r(i)} = h|s^i = l) \\ &= f(h|G) \cdot Pr(\theta = G|s^i = l) + f(h|B) \cdot Pr(\theta = B|s^i = l) \end{aligned}$$

	$\tau(l,l)$	$\tau(l,h)$	$\tau(h,l)$	$\tau(h,h)$	
min	0.26	0.17	0.17	0.4	
s. t.	0.6	0.4	-0.6	-0.4	$\geq 0.6 + \epsilon$
	-0.3	-0.7	0.3	0.7	$\geq 0.3 + \epsilon$
	0.6	0.4			$\geq 0.5 + \epsilon$
			0.3	0.7	$\geq 0.5 + \epsilon$
	≥ 0	≥ 0	≥ 0	≥ 0	

Figure 2.4: The Linear Program that is used to compute the payment scheme for the Amazon example.

		$r(i)$	
		l	h
Agent i	l	1.75	0
	h	0	1.17

Payment Scheme

Figure 2.5: The payment scheme for the Amazon example.

$$\begin{aligned} &\simeq 0.75 \cdot 0.41 + 0.15 \cdot 0.59 \\ &\simeq 0.4 \qquad (0.39418605) \end{aligned}$$

$$\begin{aligned} g(l|\,l) &= Pr(s^{r(i)} = l|s^i = l) \\ &= 1 - g(h|\,l) \\ &\simeq 0.6 \qquad (0.60581395) \end{aligned}$$

$$\begin{aligned} g(h|\,h) &= Pr(s^{r(i)} = h|s^i = h) \\ &= f(h|G) \cdot Pr(\theta = G|s^i = h) + f(h|B) \cdot Pr(\theta = B|s^i = h) \\ &\simeq 0.75 \cdot 0.92 + 0.15 \cdot 0.08 \\ &\simeq 0.7 \qquad (0.70263158) \end{aligned}$$

$$\begin{aligned} g(l|\,h) &= Pr(s^{r(i)} = l|s^i = h) \\ &= 1 - g(h|\,h) \\ &\simeq 0.3 \qquad (0.29736842) \end{aligned}$$

The payment scheme is depicted in Figure 2.5. The expected budget which equals the a priori expected payment to an agent is 0.92.

2.4 Discussion

Signaling reputation mechanisms are used for settings with different hidden abilities which would otherwise result in adverse selection as described earlier. Related settings are those with asymmetric information without the characteristic that this information are quality levels and corresponding prices. For example, there exist services that ask users to post positions of radar speed checks[1]. Obviously, position data has nothing to with the quality level of a product but the basic requirements regarding the feedback mechanism are the same. As with reputation mechanisms, users need to be motivated to post the information they learned and it could potentially be elicited by the same techniques. A similar service could elicit and display information about traffic jams, possibly in combination with GPS devices that are already built into cars and mobile phones. Yet another example are so-called "games with a purpose" [von Ahn and Dabbish, 2008]. Arguably the most popular of these, the ESP game, was designed to label the vast amount of images that are found on the web [von Ahn and Dabbish, 2004]. To achieve this, it randomly matches two players and displays the same image to both of them. The players are then asked to write down keywords that describe the image's content. When they agree on a keyword, they receive points and another image is displayed. The ESP game is interesting to our research as it shows that humans can be motivated by a non-monetary point system. At the same time, though, current experiences with the ESP game emphasize the importance of carefully built incentive schemes as players often agree on popular words, such as colors. For a game-theoretic analysis of the ESP game, see the work by Jain and Parkes [2008]. In order to elicit more informative labels, revised rules could weigh the points as to how likely a certain word is a priori using the same techniques as the MRZ mechanism [Weber *et al.*, 2008].

---

[1] www.radalert.de

Chapter 3

Pure Moral Hazard

The term moral hazard is widely used in the economics literature with examples arising in development economics, finance and all kinds of principal agent models. According to The New Palgrave Dictionary of Economics [Kotowitz, 2008, our emphasis],

> Moral hazard may be defined as actions of economic agents in maximizing their own utility to the detriment of others, in situations where they do not bear the full consequences [. . .] of their actions due to *uncertainty and incomplete information* or *restricted contracts* which prevent the assignment of full damages.

In the context of reputation mechanisms, moral hazard is mostly though not exclusively used in the spirit of limited contract enforcement [e. g., Dellarocas, 2006; Bolton and Ockenfels, 2006]. In order to distinguish between the two meanings, we begin with an example for moral hazard as it is introduced by many microeconomics text books with the specific example taken from Varian [2006]. Please note that we use the term *imperfect* information instead of *incomplete* information as the situation corresponds to a game with imperfect rather than incomplete information. Subsequently, we elucidate on the standard type of online auction site and eBay with the latter arguably being the most stressed example for moral hazard in a reputation context.

To illustrate the kind of moral hazard which is due to imperfect information, take the example of an insurance against bicycle theft. In contrast to a situation with adverse selection, we assume that all bicycle owners live in areas with identical probabilities of theft. That is, in contrast to the previous chapter there are no differences in types. Instead, the probability of theft depends on the *effort* or care that is exerted by the owner of the bicycle. For example, he can choose to leave his bicycle unlocked or to use a light lock instead of a more expensive solid lock. This choice or action clearly influences the likelihood of theft. The reasoning is that if the owner of the bicycle has an insurance that covers for the loss, he will not purchase a solid lock since he has to pay for the lock while the costs of theft are borne by the insurance company. It is important to note that this problem is one of imperfect information: if the owner's exerted effort could be observed by the insurance company, the latter could condition their premiums on the degree of risk that is taken and thereby induce the bicycle owner to exert the optimal, i. e. efficient, effort. This is why moral hazard is sometimes also referred to as *hidden action*.

In the context of reputation mechanisms, the prime example for moral hazard are online auction sites, such as eBay [e. g., Dellarocas, 2005, p. 6]. In order to discuss the difference to the bicycle theft example, we first give the standard procedure at these sites (excluding the reputation mechanism):

1. The seller has a good that he wants to sell (e. g., a book).
2. He describes the quality of the good and posts it on the auction site.
3. The site determines the buyer by some kind of auction.
4. The buyer transfers the money to the seller.
5. After reception of the payment, the seller sends the good to the buyer.

The crucial part is that the buyer pays for the good before the seller sends it. To illustrate the moral hazard problem, we can apply backward induction: the last step in the procedure is the seller deciding whether to send the good or not. At that point, however, he already received the payment, so that he is strictly better off not sending it independent of the buyer's action. In step 4, the buyer makes her decision whether to transfer the money to the seller. Anticipating that the seller will withhold the good no matter how she decides, it is only reasonable for her to keep the money for herself so that no trade takes place. Note that both buyer and seller can fully observe the action of the other player and, still, the market fails. That is, in contrast to the bike theft example, the problem is not one of hidden action. Rather, the problem is one of *hidden intention*. Nevertheless, some researchers view certain aspects of eBay's procedure as hidden action. For example, the seller may send the good but delay the sending which negatively influences the date of arrival and hence the utility of the buyer. However, as the sending date is given on the package, it is revealed to the buyer at the time of arrival. Therefore, the action is observed ex-post and could be (and often is) put into the good's description, such as "will be sent within three working days upon receipt of payment". We furthermore believe that the uncertainty about the duration of money transfer is negligible.

If the buyer and the seller are known to repeatedly interact with one another, cooperation can arise due to the threat of future punishment through the other party, so that in these settings no feedback is required [e. g., Fudenberg and Tirole, 1991, p. 192]. Unfortunately, this situation is not given in most electronic markets. Instead, most buyers interact with a specific seller only once and the information regarding the outcome of these games is private knowledge of the participating agents. This chapter is therefore devoted to the exploration of the mechanism design space of truthful feedback elicitation. The remainder of this chapter is organized as follows. We begin with a formal description of the setting at eBay-like online auctions in Section 3.1. In Section 3.2 we give the definition of the peer-based Feedback Game that is played between the seller and two buyers. In Section 3.3 we then show that it is not possible to create a payment scheme similar to that of Chapter 2. In Section 3.4 we conclude with a discussion of potential remedies for this impossibility result.

3.1 The Setting

The setting is that of a prototypical online auction with the procedure given earlier. We make the usual assumption for moral hazard mechanisms that the seller is a long-lived player that meets a sequence of short-lived buyers, i. e. buyers who are interested in the seller's product for only one round and then depart. For the feedback game, we leave out the seller's quality announcement at step 2 and assume an honest description. It is important to note that this is not a simplification of the problem for the following reason: the good's description together with the promise of sending it (possibly within a certain time) is the contract. If—given the assumption of an honest description—we can solve the feedback part so that every buyer honestly reports whether this contract was fulfilled, the seller has no incentive to lie. Announcing a quality level that is lower than the actual product's would result in a lower profit while announcing a quality that is too high would breach the contract, so that he could as well keep the good for himself, i. e. not send it at all. Therefore, if we can solve the feedback part, we also ensure an honest product description. When the product's description is online, potential buyers bid in an auction to receive it. The highest bidder wins and has to pay the amount of the second-highest bid. It is well-known that it is a weakly-dominant strategy for every buyer to bid according to her valuation [e. g., Nisan, 2007].

We restrict our analysis to binary settings with regard to both seller actions and observed signals. That is, the feedback procedure for a single seller-buyer transaction consists of the seller's choice whether to cooperate or cheat, the buyer's binary signal observation and her subsequent signal report. Please note that the negative result of Section 3.3 readily extends to the general case.

To denote the seller action for buyer i, we use the notation

$$e^i \in \{\text{coop}, \text{cheat}\}. \tag{3.1}$$

Note that in contrast to Chapter 2, every seller has the same abilities to offer good service (i. e. to cooperate) with cooperation being costlier than cheating:

$$c(\text{coop}) > c(\text{cheat}). \tag{3.2}$$

Let s^i denote the signal received by buyer i.

$$s^i \in \{s_1, s_2\} \tag{3.3}$$

Please note that for reasons of clarity, we write low (l) and high (h) signals instead of s_1 and s_2 whenever possible. The seller action influences buyer i's signal in that cooperation makes it more probable that the signal is high. Nevertheless, it is perfectly possible that cooperation results in a negative signal. For example, the seller may send the good as described, i. e. he cooperates, but the package is lost in the mail. It is important to notice that the seller cannot observe the outcome of his actions, e. g., whether the package was lost. The probability that buyer i receives a high signal given the respective seller action is:

$$f(h|\, e^i) = Pr(s^i = h|\, e^i), \quad e^i = \text{coop}, \text{cheat} \tag{3.4}$$

Note that while for the eBay example a high outcome is reasonable only given cooperation, the model is more expressive in that it allows settings in which cheating can be followed by a high outcome. These signal emissions constitute a probability distribution, so that

$$f(h|\, e^i) + f(l|\, e^i) = 1, \quad e^i = \text{coop}, \text{cheat} \tag{3.5}$$

We assume that $f(h|\cdot)$ is common knowledge and identical for all buyers.

The buyers have valuations

$$v^i(s^i) \tag{3.6}$$

for the two possible signals. We only assume that a high signal is strictly preferred to a low signal, that is

$$v^i(s^i = h) > v(s^i = l). \tag{3.7}$$

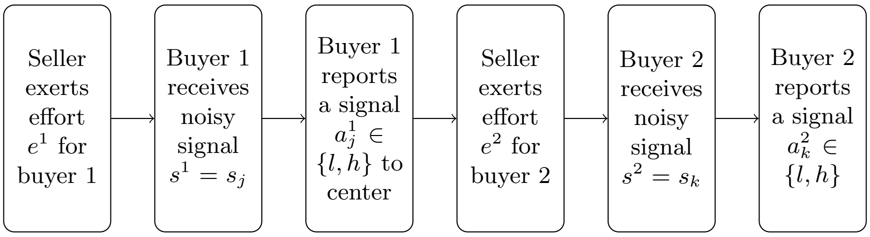

Figure 3.1: The procedure of the pure moral hazard Feedback Game between the seller and the first two buyers.

As in Chapter 2, a payment scheme shall incorporate external benefits from lying. Let $\Delta^i(s_j, s_h)$ be the external benefit buyer i could gain by falsely announcing signal s_h instead of signal s_j, the one actually received. We assume an upper bound

$$\Delta(s_j, s_h) = \max_i \Delta^i(s_j, s_h) \tag{3.8}$$

on the individual external lying benefits. Note that by definition $\Delta(s_m, s_m) = 0$ for all $s_m \in S$. Let C^i be the costs reflecting buyer i's time and effort required for the rating process. Similar to the lying benefits, we assume an upper bound

$$C = \max_i C^i. \tag{3.9}$$

3.2 The Feedback Game

Similar to the MRZ mechanism in Chapter 2, we want to rate the reports of two buyers against one another. Who is rated against whom depends on the reference reporter choice rule. We simply rate the first two buyers against each other, compute the payments and continue with the third and the fourth buyer. The first game that is played is therefore the three player game between the first two buyers and the seller. That is, we compare the first buyer's announcement to that of the second buyer and vice versa. Different from the seller's actions that are given with the setting, the buyers' actions are only introduced by the feedback procedure: let $a^i = (a_1^i, a_2^i)$ be the reporting strategy of buyer i, such that she reports signal $a_j^i \in \{l, h\}$ if she received s_j. The honest strategy is then $\bar{a} = (l, h)$, i. e. always reporting the signal received. See Figure 3.1 for the depicted procedure.

The utility function of the seller consists of three parts: the payment he receives for his service, the costs that are associated with his actions and the valuation of the buyers' signal announcements that are to be published. As mentioned in the introduction, the objective is to design a "feedback plug-in" that can be used by

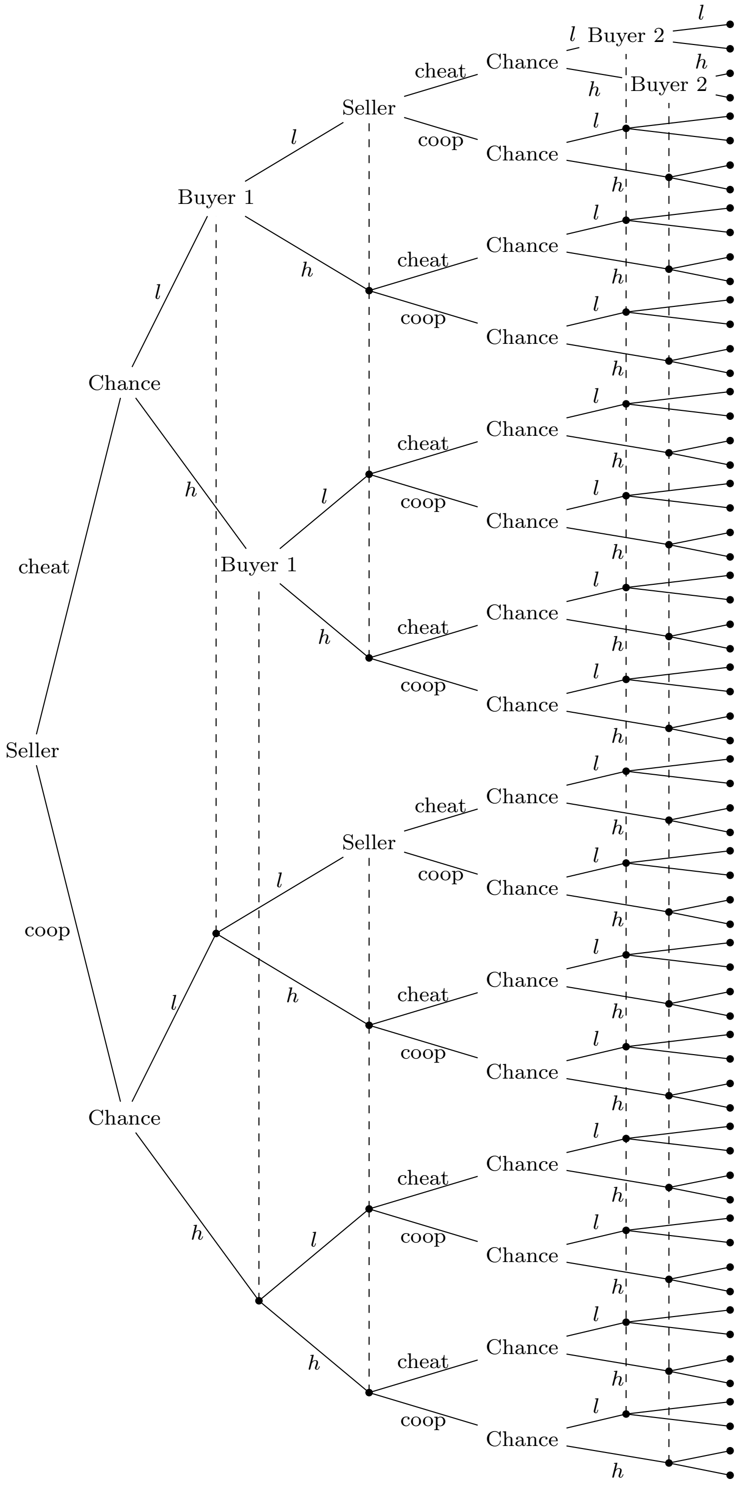

Figure 3.2: Complete game tree for the Feedback Game between the first two buyers and the seller. Note that buyer 2 has actions l and h at both of her information sets.

reputation mechanisms that otherwise would simply *assume* honest feedback. With regard to the seller's behavior we thus make a crucial assumption: *given the center could directly observe the signals* (which corresponds to the assumption of honest feedback made by most reputation mechanisms), *the center could compute the seller's optimal mixed strategy profile* $p(e^i)$. While at first sight this may seem like a rather strong assumption, it simply means that we can compute how the publication of transaction outcomes influences the seller actions. This is done in many reputation mechanism papers where the objective is usually to induce the seller to cooperate as much as possible. Dellarocas, for example, studies how the length of published history influences seller cooperation in a setting similar to ours [2005]. That is, we are taking the complementary view to the usual literature and assume that *we are given a reputation mechanism that induces some degree of seller cooperation and which gives us the seller's mixed strategy profile that he would play under the assumption that feedback is honest.*

The utility function of buyer i also consists of three parts: the price she has paid for the seller's service, the valuation of the received signal $v^i(s^i)$ and the payment associated with her announcement and the announcement of her reference reporter $\tau^i(a^i, a^{r(i)})$. As the price is already determined in the auction, it is independent of the Feedback Game's outcome for all players and we can ignore it. Figure 3.2 shows the game tree's structure of this three player extensive game with imperfect information.

3.3 Necessary Conditions for Truthfulness

In order to have a unifying concept for truthfulness, we include the seller in the definition. We are aware that there are other concepts such as *faithfulness* that capture both information revelation and robustness to rational manipulation [Shneidman and Parkes, 2004]. When the monitoring of the outcomes is imperfect, however, it is not generally possible to achieve full seller cooperation, and thus robustness to rational manipulation, even if feedback is truthful. The seller may cheat occasionally because a bad signal does not immediately reveal his opportunistic nature. It is again important to distinguish between the objectives of the reputation mechanism and the feedback mechanism, respectively: the reputation mechanism that we are given induces some degree of seller cooperation if feedback is truthful. Depending on the actual setting and in particular the level of noise in the setting, it may be possible to achieve only partial cooperation. The objective of the feedback mechanism is to provide the reputation mechanism with truthful buyer feedback that the latter can use to induce as much seller cooperation as possible. That is, it could well be that the reputation mechanism induces only partial cooperation given truthful feedback while the feedback mechanism induces 100% truthful buyer feedback given this reputation mechanism. Since faithfulness would imply that the seller is fully cooperative in equilibrium, it is not better-suited for our purpose and we include the seller in the definition of truthfulness for the remainder of this thesis.

Definition 4 (Truthfulness)**.** An equilibrium is *truthful* if and only if the seller plays according to $p(e^i)$ and both buyers report their signal outcomes honestly.

In extensive games, such as the Feedback Game, Nash equilibrium is no longer a sufficient concept due to incredible threats or promises. For extensive games with perfect information, subgame perfect equilibrium is widely accepted. An equilibrium concept that generalizes subgame perfection to games with imperfect information is perfect Bayesian equilibrium which we apply in the mixed setting of Chapter 4. At this point, however, it is sufficient to note that for every reasonable equilibrium of an extensive game the best response condition must hold. The way forward is therefore to take out a player, fix the truthful strategies of the remaining players and search for payments that make it in the best interest of the left-out player to be truthful as well. Proposition 6 shows that this is impossible.

Proposition 6. *There is no payment scheme for the Feedback Game that makes truthfulness a best response to truthful play by all other players.*

Proof. The proof is by contradiction. Assume there is a payment scheme that makes truthful play a best response to truthful play by all other players. Then, given truthful play by both the seller and buyer 1, there are payments that make honest reporting by the second buyer a best response. The expected payment for buyer 2 is computed by Equation 3.10. Please note that we abbreviate $E(a^2_j = s_h | s^2 = s_j)$ with $E(s^2_h | s^2 = s_j)$.

$$\begin{aligned} E(s^2_h | s^2 = s_j) = v^2(s_j) & \\ &+ \big[p(e^1 = \text{coop}) \cdot f(l |\, \text{coop}) \\ &+ p(e^1 = \text{cheat}) \cdot f(l |\, \text{cheat})\big]\, \tau^2(s_h, l) \\ &+ \big[p(e^1 = \text{coop}) \cdot f(h |\, \text{coop}) \\ &+ p(e^1 = \text{cheat}) \cdot f(h |\, \text{cheat})\big]\, \tau^2(s_h, h) \end{aligned} \tag{3.10}$$

The report of buyer 2 is honest if and only if given a negative signal, she reports l while given a positive signal she reports h. Honestly reporting a negative signal is a best response if given a negative signal the expected utility for reporting l is at least as high as the expected utility for reporting h. The analogous must hold true for the honest reporting of a positive signal. Her utility depends on the payment for her report and the external benefits she may receive as captured by $\Delta(\cdot,\cdot)$. The honesty constraints that must hold are thus:

$$
\begin{aligned}
E(l^2|s^2 = l) - E(h^2|s^2 = l) &\geq \Delta(l,h) && (3.11)\\
E(h^2|s^2 = h) - E(l^2|s^2 = h) &\geq \Delta(h,l) && (3.12)
\end{aligned}
$$

If $\Delta(l,h) > 0$ and $\Delta(h,l) > 0$, necessary conditions for Equation 3.11 and Equation 3.12 are given by Equation 3.13 and 3.14, respectively. Note that these require strict inequalities:

$$
\begin{aligned}
E(l^2|s^2 = l) = {}& v^2(l)\\
&+ \big[p(e^1 = \text{coop}) \cdot f(l|\,\text{coop})\\
&+ p(e^1 = \text{cheat}) \cdot f(l|\,\text{cheat})\big]\,\tau^2(l,l)\\
&+ \big[p(e^1 = \text{coop}) \cdot f(h|\,\text{coop})\\
&+ p(e^1 = \text{cheat}) \cdot f(h|\,\text{cheat})\big]\,\tau^2(l,h)\\
> {}& v^2(l) \qquad (3.13)\\
&+ \big[p(e^1 = \text{coop}) \cdot f(l|\,\text{coop})\\
&+ p(e^1 = \text{cheat}) \cdot f(l|\,\text{cheat})\big]\,\tau^2(h,l)\\
&+ \big[p(e^1 = \text{coop}) \cdot f(h|\,\text{coop})\\
&+ p(e^1 = \text{cheat}) \cdot f(h|\,\text{cheat})\big]\,\tau^2(h,h)\\
= {}& E(h^2|s^2 = l)
\end{aligned}
$$

$$
\begin{aligned}
E(h^2|s^2 = h) = {}& v^2(h)\\
&+ \big[p(e^1 = \text{coop}) \cdot f(h|\,\text{coop})\\
&+ p(e^1 = \text{cheat}) \cdot f(h|\,\text{cheat})\big]\,\tau^2(h,h)\\
&+ \big[p(e^1 = \text{coop}) \cdot f(l|\,\text{coop})\\
&+ p(e^1 = \text{cheat}) \cdot f(l|\,\text{cheat})\big]\,\tau^2(h,l)\\
> {}& v^2(h) \qquad (3.14)\\
&+ \big[p(e^1 = \text{coop}) \cdot f(h|\,\text{coop})\\
&+ p(e^1 = \text{cheat}) \cdot f(h|\,\text{cheat})\big]\,\tau^2(l,h)\\
&+ \big[p(e^1 = \text{coop}) \cdot f(l|\,\text{coop})\\
&+ p(e^1 = \text{cheat}) \cdot f(l|\,\text{cheat})\big]\,\tau^2(l,l)\\
= {}& E(l^2|s^2 = h)
\end{aligned}
$$

Put together, Equation 3.13 and Equation 3.14 imply that

$$\begin{aligned}
&\left[p(\mathrm{coop}^1) \cdot f(l|\,\mathrm{coop}) + p(\mathrm{cheat}^1) \cdot f(l|\,\mathrm{cheat})\right] \;\; \tau^2(l,l) \\
&+ \left[p(\mathrm{coop}^1) \cdot f(h|\,\mathrm{coop}) + p(\mathrm{cheat}^1) \cdot f(h|\,\mathrm{cheat})\right] \tau^2(l,h) \\
> \;\; &\left[p(\mathrm{coop}^1) \cdot f(l|\,\mathrm{coop}) + p(\mathrm{cheat}^1) \cdot f(l|\,\mathrm{cheat})\right] \;\; \tau^2(h,l) \\
&+ \left[p(\mathrm{coop}^1) \cdot f(h|\,\mathrm{coop}) + p(\mathrm{cheat}^1) \cdot f(h|\,\mathrm{cheat})\right] \tau^2(h,h) \\
> \;\; &\left[p(\mathrm{coop}^1) \cdot f(h|\,\mathrm{coop}) + p(\mathrm{cheat}^1) \cdot f(h|\,\mathrm{cheat})\right] \tau^2(l,h) \\
&+ \left[p(\mathrm{coop}^1) \cdot f(l|\,\mathrm{coop}) + p(\mathrm{cheat}^1) \cdot f(l|\,\mathrm{cheat})\right] \;\; \tau^2(l,l).
\end{aligned}$$

Please observe that the first two lines are identical to the left side of Equation 3.13. Line 3 and 4 are identical to the right side of Equation 3.13 and the left side of Equation 3.14. Lines 5 and 6 are identical to the right side of Equation 3.14 and the left side of Equation 3.13. This is a contradiction as we have strict inequalities. ☐

Corollary 7. *There is no payment scheme that makes truthfulness a perfect Bayesian equilibrium in the Feedback Game.*

3.4 Discussion

The intuition of this chapter's negative result is that, in equilibrium, the seller's strategy would be known a priori by both buyers so that the signal observation of buyer i does not lead to a belief update regarding the signal received by buyer $r(i)$. Without a posterior belief update, however, the two signal observations cannot be stochastically relevant for each other and buyer i's utility maximization cannot depend on the signal she received. It is for the same reason that we cannot do better with a change of procedure such that the announcement of a buyer is scored against the announcement of the seller: because there is no uncertainty about the seller's equilibrium behavior, the buyer's optimal announcement is known a priori, i. e. *before* a buyer knows her signal. Since there are positive results for similar settings in contract theory [Laffont and Martimort, 2001, pp. 298–302], however, there is reason to believe that matching the announcement of the seller against that of a single buyer can lead to truthful payment schemes if—after exerting his effort—the seller observes the *outcome* of his action.

For a pure moral hazard setting similar to ours, Dellarocas [2005] finds that full seller cooperation cannot be achieved when monitoring is imperfect. A remedy from the impossibility result from Section 3.3 could thus be to incorporate a Stackelberg ("commitment") type in the Feedback Game's definition.

That is, one could assume that there is a small probability that the seller is of a type that always plays cooperation. Since the normal ("strategic") type plays cooperation with a lower probability, the two seller types are distinguishable and the two signal observations should become stochastically relevant. Please note that the introduction of commitment types fits well in the standard game-theoretic literature on reputation building [Milgrom and Roberts, 1982; Kreps and Wilson, 1982a]. Alternatively, if there is a small prior probability that the seller is of a "cheating" type, truthful feedback elicitation becomes a special case of the next chapter's setting for which we retrieve a positive result.

Chapter 4

A Mixed Setting

Before we elucidate on the setting of this chapter, let us briefly summarize what we know about the mechanism design space of peer-based feedback elicitation: MRZ have shown that honest feedback mechanisms are feasible if different buyers' experiences are essentially identical, i. e. differences in the two buyers' perceptions are only due to stochastic noise. In Chapter 3 we have shown that a purely strategic connection between the two signal observations does not allow a truthful payment scheme. A natural next step is thus to study whether it is possible to elicit truthful feedback in a setting where the buyers' signals are partly stochastic while also controlled by a strategic seller.

Consider, for example, an Internet booking site for the delivery of groceries. After giving both the list of groceries and a postal code to the booking site, it displays all available delivery services together with their respective prices. In addition to price, however, the services differ in their quality, i. e. their respective speed of delivery. Therefore, a reputation mechanism is located at the booking site that—next to the price—displays the expected duration of delivery. A customer may then choose a service depending on how she values price to speed. The speed of delivery is influenced both by the services' inherent abilities and their actions. For example, services differ in the number of cars they own, i. e. in their type, but a service may also delay its delivery until it has a certain number of orders in the same neighborhood. That is, the services are not required to deliver at their highest possible speed. We assume that the number of users that are shopping via the booking site is small relative to the services' total number of customers, so that the impact of the booking site's customers on the services' behavior is negligible. We furthermore assume that a service's inherent ability does not change over time. The latter, however, can be incorporated as described in our earlier work [Witkowski, 2009].

Without a reputation mechanism, the expected quality is the same for all sellers and a buyer would thus choose the cheapest seller. If the sellers were unable to choose a quality level below their abilities, the result would be one of pure adverse selection: good sellers are pushed out of the market since good quality is costlier to produce. Moral hazard is introduced since sellers can choose to deliver a quality level *below* their abilities with lower quality levels corresponding to lower delivery costs. The result of this market situation without a reputation mechanism would therefore be a competition of all seller types for the lowest quality level. A particularly interesting situation arises if one assumes that higher seller types are more efficient in producing the worst quality level than sellers of lower type. In that case, sellers of the highest type produce at the worst possible quality level. Thus, by introducing moral hazard to a setting with pure adverse selection, the result can be the drive-out of all but the best sellers. This is the opposite of adverse selection. (However, as with adverse selection, it is still the worst quality that is produced).

The remainder of this chapter is organized as follows. In Section 4.1, we introduce the mixed setting and formalize the Feedback Game. Section 4.2 is devoted to the "feedback plug-in" for reputation mechanisms: in Section 4.2.1, we briefly compare sequential equilibrium (SE) with perfect Bayesian equilibrium (PBE) and emphasize why the easier to compute PBE is sufficient in our setting. In Section 4.2.2 we show how to compute the required signal posteriors and in Section 4.2.3 we eventually set up the payment scheme. Section 4.3 provides a detailed example for a different kind of auction than the one analyzed in Chapter 3. In Section 4.4 we conclude the chapter with a brief discussion of the particular difficulties reputation mechanisms are faced with in mixed settings as compared to either of the pure settings.

4.1 The Feedback Game

The Feedback Game is an extensive game with both incomplete and imperfect information. For readers unfamiliar with these concepts, we recommend the textbook by Osborne and Rubinstein [1994]. As in Chapter 3, the set of players is

$$N = \{\text{seller}, 1, 2\} \tag{4.1}$$

which denotes the seller, the first and the second buyer. Please note that, when it cannot be confused with a signal, we sometimes use the letter "s" to denote the seller. Figure 4.1 depicts the complete procedure of a single game. The Feedback Game begins with a move by "nature" that chooses the seller's type θ from a finite set of types

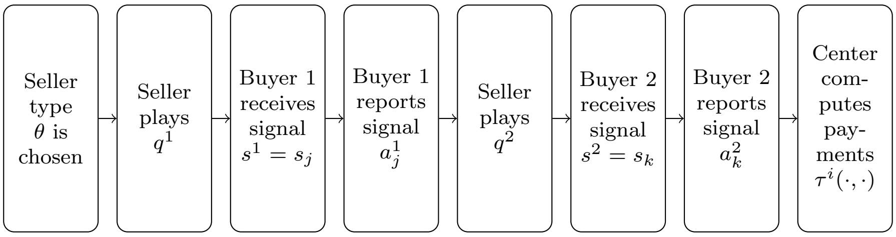

Figure 4.1: The procedure of the mixed setting Feedback Game between the seller and the first two buyers.

$$\Theta = \{\theta_1, \ldots, \theta_T\} \tag{4.2}$$

with $T = |\Theta| \geq 2$. Note that this move by "nature" is not an actual move but used to model the buyers' beliefs regarding the seller's type. That is, all agents share a common prior belief $Pr(\theta_t)$ that the seller is of type θ_t with

$$\sum_{t=1}^{T} Pr(\theta_t) = 1 \tag{4.3}$$

while $Pr(\theta_t) > 0$ for all $\theta_t \in \Theta$. Once determined, the seller learns his own type which stays fixed for the remainder of the game. After the Feedback Game, the type beliefs are updated and used as prior beliefs for the game between the seller, the third and the fourth buyer.

In contrast to the pure adverse selection setting of MRZ, all sellers but those of type θ_1 are players in the actual sense. That is, there is an action set associated with the seller which corresponds to the quality levels he can produce. Different seller types have different quality levels to choose from and the number of quality levels corresponds to the "number" of the type: let θ_t denote the t-th lowest seller type and let q^i denote the quality level the seller plays for buyer i. A seller with type θ_t has t different quality levels to choose from, namely every quality level up to and including the one that corresponds to his type, i. e.

$$q^i \in \{q_1, \ldots, q_t\}. \tag{4.4}$$

We assume perfect recall, i. e. the seller never forgets what he has played earlier in the game. As in Chapter 3, the buyers do not observe q^i directly but some noisy signal s^i. The signal depends on the quality that was produced and different quality levels yield different signal distributions. As before, the signals are drawn out of a set of signals:

$$S = \{s_1, ..., s_M\}. \tag{4.5}$$

Let s^i denote the signal received by buyer i and let

$$f(s_m|\, q_l) = Pr(s^i = s_m|\, q^i = q_l) \tag{4.6}$$

be the probability that buyer i receives the signal $s_m \in S$ given that the played quality level is $q^i = q_l$. These signal emissions constitute a probability distribution:

$$\sum_{m=1}^{M} f(s_m|\, q_l) = 1 \tag{4.7}$$

for all $l \in \{1, \ldots, T\}$. We assume that $f(\cdot|\,\cdot)$ is common knowledge.
Buyers are all of the same type. Analogously to Chapter 3, their actions correspond to possible signal reports. Let

$$a^i = (a^i_1, \ldots, a^i_M) \tag{4.8}$$

be the reporting strategy of buyer i, such that she reports signal $a^i_j \in \{s_1, \ldots, s_M\}$ if she received s_j. The honest strategy is

$$\bar{a} = (s_1, \ldots, s_M), \tag{4.9}$$

i. e. always reporting the signal received. The complete game tree of a small example (excluding utilities) is depicted on p. 46.
As in Chapter 3, the utility of a seller, U^s, depends on three parts. The first is the negative utility associated with the quality levels that he produces:

$$-c_\theta(q^1) - c_\theta(q^2). \tag{4.10}$$

The costs for the seller's actions can be different for different seller types. Furthermore, we assume that the higher the quality level, the higher the costs. For the feedback mechanism, however, we do not need to know these costs explicitly. The second part is the valuation for the given feedback due to the impact a publication has on future buyers' willingness to pay

$$v^s(a^1_j, a^2_k|\, r). \tag{4.11}$$

Note that the seller's valuation for the same reported signals can be different for a different publication history. We use variable r instead of h to prevent confusion with high signals.
The third part is the payment that the seller receives for his service. With

regard to seller utility, we make the analogous assumption to that of Chapter 3: given the center could directly observe the signals, it could compute the seller's optimal, i. e. truthful, strategy

$$p_t(q^i) \tag{4.12}$$

for each seller type θ_t. We assume that when the seller plays according to this strategy, the prior probability for every signal is positive, i. e.

$$Pr(s^i = s_j) > 0 \tag{4.13}$$

for all $s_j \in S, i = 1, 2$. We furthermore assume that at least two seller types differ in their optimal strategies. Note that both assumptions are usually met in real-world applications. The utility of a buyer is defined as in Chapter 3 and consists of the negative utility of the price that she paid for the good, the valuation for the signal she perceived

$$v^i(s^i) \tag{4.14}$$

and the payment that depends on both her own signal announcement and that of the other buyer

$$\tau^i(a^i, a^{r(i)}). \tag{4.15}$$

As before, the price for the service is independent of the game's outcome, so that we can ignore it here.

As in the preceding chapters, the payment scheme shall incorporate external benefits from lying with $\Delta^i(s_j, s_h)$ denoting the external benefit buyer i could gain by falsely announcing signal s_h instead of the honest signal s_j. We assume an upper bound

$$\Delta(s_j, s_h) = \max_i \Delta^i(s_j, s_h) \tag{4.16}$$

on the individual external lying benefits with $\Delta(s_m, s_m) = 0$ for all $s_m \in S$. Furthermore, let C^i be the costs reflecting buyer i's time and effort required for the rating process. Similar to the lying benefits, we assume an upper bound on the individual participation costs and obtain:

$$C = \max_i C^i. \tag{4.17}$$

Once both reports are elicited, they can be published. The type beliefs of the next Feedback Game are then computed using Bayesian updating (compare Figure 4.2). Note that it is unproblematic to publish more than the two most

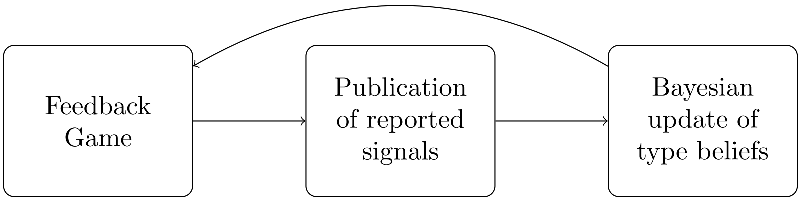

Figure 4.2: The overall procedure of the mixed setting.

recent signals. This is crucial as the optimal length of the publication history in settings with both adverse selection and moral hazard is an open research question [Dellarocas, 2006, p. 20]. We come back to this topic in Section 4.4.

4.2 A "Feedback Plug-in"

The first step in constructing a payment scheme that allows for a truthful equilibrium is to decide on a suitable equilibrium concept. Perfect Bayesian equilibrium is a reasonable choice as it generalizes subgame perfection to games with incomplete information. In Section 4.2.2 we show how to compute a buyer's posterior belief about her reference reporter's signal that is required to compute the buyer's expected utility which we then use to formulate the payment scheme in Section 4.2.3.

4.2.1 Perfect Bayesian Equilibrium

Intuitively an assessment is a perfect Bayesian equilibrium if the strategies are rational given the beliefs and the beliefs are updated by Bayes' law whenever possible.

Definition 5. A *perfect Bayesian equilibrium* is a strategy profile β^* and a belief system μ^* such that

1. At information sets reached with positive probability when β^* is played, beliefs are formed according to β^* and Bayes' law (when necessary). At information sets that are reached with probability zero when β^* is played, beliefs may be arbitrary.

2. At every information set I_i player i's strategy maximizes his/her payoff, given the actions of all the other players β^*_{-i} and player i's beliefs.

Combinations (β', μ') of a strategy profile β' and a belief system μ' are called an *assessment*.

If, in equilibrium, there were information sets that are not reached, sequential equilibrium would reasonably restrict beliefs on these information sets [Kreps and Wilson, 1982b]. Since $Pr(s^i = s_j) > 0$ for all $s_j \in S$ and both buyers, all information sets that belong to buyers are reached with positive probability given truthful play. That means, no buyer is ever "surprised" by an information set due to another player's tremble, i. e. a deviation from truthful play. Every observation a buyer makes will thus be interpreted as the result of truthful play and will therefore not lead to a deviation of her truthful strategies. Note that this is not a feature of the game as much as a feature of the truthful strategies. That is, there can be other (non-truthful) assessments in the Feedback Game where some of the buyers' information sets are not reached. The seller does not make observations other than his type and, technically, the information set that precedes his choice regarding q^2. The "second group" of his information sets, however, solely depends on his earlier choice q^1 and having an individual (behavioral) strategy for each of these is therefore not reasonable. In fact, from an informational point of view, the two seller actions, q^1 and q^2, could be merged into a single action which is placed directly after the initial move by "nature" (he has the same information when choosing q^1 and q^2). The fact that the set of sequential equilibria can change by such a coalescing of moves is a drawback of the sequential equilibrium concept rather than the result of a change in the strategic environment [Osborne and Rubinstein, 1994, pp. 207f, 226f].

4.2.2 Belief Computations

The best response conditions are trivially met for the seller who, given truthful play by the buyers, is truthful by definition. This is different for the buyers and we proceed by fixing truthful play for the seller and buyer $r(i)$ to determine the constraints that necessarily have to be met for buyer i. For the pure adverse selection setting of Chapter 2, MRZ prove that combinations of $Pr(\theta)$ and $f(\cdot|\cdot)$ that fail stochastic relevance (Definition 1) have Lebesgue measure 0 (see Discussion on p. 10). We conjecture that the analogous statement holds for our model. That is, due to the assumption that at least two seller types differ in their optimal strategies, those combinations of $Pr(\theta_t)$, $p_t(q^i)$ and $f(\cdot|\cdot)$ that fail stochastic relevance are stochastically relevant after a slight perturbation of beliefs. In the following, we thus assume stochastic relevance holds.

Let $s^{r(i)} = s_k$ and $s^i = s_j$ denote the signals received by buyer $r(i)$ and buyer i, respectively. The probability that buyer $r(i)$ received s_k given i received s_j is defined as:

$$g^i(s_k|\, s_j) = Pr(s^{r(i)} = s_k|\, s^i = s_j). \tag{4.18}$$

After receiving her signal, buyer i's expected utility is then given by Equation 4.19 (we slightly abuse the notation and incorporate both the external lying benefits $\Delta(\cdot,\cdot)$ and the participation costs C only in the LP):

$$E(a_j^i = s_h | s^i = s_j) = v^i(s_j) + \sum_{k=1}^{M} g^i(s_k|\, s_j)\, \tau^i(s_h, s_k). \tag{4.19}$$

We can compute $g^i(s_k|\, s_j)$ with Equation 4.20:

$$g^i(s_k|\, s_j) = \sum_{l=1}^{T} f(s_k|q_l) \cdot Pr(q^{r(i)} = q_l | s^i = s_j). \tag{4.20}$$

The probability of a specific quality level played for buyer $r(i)$ given the signal perceived by buyer i is obtained by Equation 4.21:

$$Pr(q^{r(i)} = q_l | s^i = s_j) = \frac{Pr(s^i = s_j | q^{r(i)} = q_l) \cdot Pr(q^{r(i)} = q_l)}{Pr(s^i = s_j)}. \tag{4.21}$$

The prior signal probability for buyer i is:

$$Pr(s^i = s_j) = \sum_{l=1}^{T} f(s_j|q_l) \cdot Pr(q^i = q_l). \tag{4.22}$$

The prior probability for a specific quality level played for buyer i that is required for Equations 4.21, 4.22 and 4.25 can be computed with Equation 4.23:

$$Pr(q^i = q_l) = \sum_{t=1}^{T} p_t(q^i = q_l) \cdot Pr(\theta_t). \tag{4.23}$$

For Equation 4.21, we require $Pr(s^i = s_j | q^{r(i)} = q_l)$. At first sight, it may seem that the probability of the quality level for buyer $r(i)$ is independent of the other buyer's signal. This is not true, however, since the materialized $q^{r(i)}$ tells us something about the seller's type which again influences buyer i's signal:

$$Pr(s^i = s_j | q^{r(i)} = q_l) = \sum_{t=1}^{T} Pr(s^i = s_j | \theta_t) \cdot Pr(\theta_t | q^{r(i)} = q_l). \tag{4.24}$$

Computing the type probability knowing the quality level for a certain buyer is a simple Bayesian update:

$$Pr(\theta_t | q^{r(i)} = q_l) = \frac{p_t(q^{3-1} = q_l) \cdot Pr(\theta_t)}{Pr(q^{r(i)} = q_l)}. \tag{4.25}$$

Equation 4.26 shows how to obtain the probability for a signal knowing the seller's type.

$$Pr(s^i = s_j|\theta_t) = \sum_{l=1}^{T} f(s_j|q_l) \cdot p_t(q^i = q_l) \tag{4.26}$$

Finally, the Bayesian update for the type beliefs are computed with Equation 4.27 (compare Figure 4.2 on p. 40):

$$Pr(\theta_t|s^i = s_j) = \frac{Pr(s^i = s_j|\theta_t) \cdot Pr(\theta_t)}{Pr(s^i = s_j)}. \tag{4.27}$$

4.2.3 The Payment Scheme

Of course a buyer will only announce her signal honestly if that maximizes her expected utility. As in Section 2.2.3, we search for $\tau^i(\cdot,\cdot)$ as the solution to a Linear Program with minimized budget. For each of buyer i's possible signal observation $s^i = s_j$, there are $M-1$ dishonest announcements $a^i_j \neq \bar{a}_j$. Given truthful play by both the seller and buyer $r(i)$, we want the expected payment of an honest announcement by buyer i to be larger than the expected utility of any other announcement including the external lying incentives $\Delta(s_j, s_h)$. Note that buyer i's signal valuation $v^i(s_j)$ appears on both sides of the inequality so that it can be canceled out (compare Equation 4.19):

$$\begin{gathered} \sum_{k=1}^{M} g^i(s_k|\, s_j) \cdot \tau^i(s_j, s_k) - \sum_{k=1}^{M} g^i(s_k|\, s_j) \cdot \tau^i(s_h, s_k) > \Delta(s_j, s_h) \\ \forall s_j,\, s_h \in S,\, s_j \neq s_h,\, i \in \{1, 2\} \end{gathered} \tag{4.28}$$

Besides honesty, there is the problem of under-provision of feedback. A buyer 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 buyer's decision whether to participate in the rating is taken after receiving her signal, *interim* individual rationality is appropriate (compare Section 2.2.3):

$$\sum_{k=1}^{M} g^i(s_k|\, s_j) \cdot \tau^i(s_j, s_k) > C \quad \forall s_j \in S,\, i \in \{1, 2\} \tag{4.29}$$

The expected budget the center needs to pay the buyers for their reports is the sum of expected payments to buyer 1 and buyer 2. As we have seen, Equation 4.19 gives us the expected payment to buyer i given that the perceived signal was $s^i = s_j$. By weighting the expected payment with the respective prior signal probability from Equation 4.22, we can then formulate the expected budget. Together with the assumption that there is no possibility to withdraw credit from the buyers (so that all payments are non-negative), the LP for buyer i is:

LP 2.

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

Proposition 8. *Truthfulness is a perfect Bayesian equilibrium of the Feedback Game.*

If LP 2 is feasible, Proposition 8 follows directly from the formulation of the payment scheme and the definition of the seller utilities. Assuming stochastic relevance, the proof of LP 2's feasibility is analogous to that of Proposition 4 on p. 17.

4.3 Example: Elance

Freelance auction sites such as Elance[1] or RentACoder[2] are real-world examples for settings with both adverse selection and moral hazard. One may regard them as the equivalent of eBay for services: a potential customer posts a project and providers (henceforth: sellers) bid on finishing it. Note that in contrast to eBay, these portals employ a *reverse* auction, i. e. the sellers are bidding for projects posted by potential buyers. Another difference to eBay is that the buyers are not deemed to take the best offer but consider both the bid and the published transaction outcomes of the seller. As these transaction outcomes are private information of the respective customers this is where the "feedback plug-in" comes into play.

Consider the example of a web designer who offers his service via Elance. Clearly, different web designers have different abilities as virtually everybody with an Internet connection can open up an account and offer his services. Yet while the worst web designers do a low quality job no matter what they do,

---

[1] www.elance.com

[2] www.rentacoder.com

skilled web designers may choose to produce a quality level that is below their abilities. Obviously, the production of low quality levels involves less time and money than the production of high quality. As the buyers are no web design experts, however (which is why they bought the designer's service in the first place), they observe the job's actual quality with noise. Note that there are customers in the population that only require a low quality job. For example, some customers may find it sufficient to receive a very simple website with contact data.

Let there be two types of web designers, a bad type θ_1 and a good type θ_2 for which we write B and G, respectively. The common knowledge prior beliefs are $Pr(\theta = B) = 0.3$ and $Pr(\theta = G) = 0.7$. A good designer's actions are denoted b and g which corresponds to q_1 and q_2. Let there be two signals, l and h, and let $f(h|b) = 0.1$ and $f(h|g) = 0.8$ be the conditional signal probabilities. Please note that for reasons of clarity, we write low (l) and high (h) signals instead of s_1 and s_2. If the center could observe every of the designer's signal outcomes and publish it accordingly, the good type would usually produce good quality. Occasionally, however, he would do a "sloppy" job which, in case of a bad signal outcome, will be indistinguishable from an occasional bad signal following good quality play. Let the good designer play bad in 5% of his jobs, i. e. $p_2(q^i = g) = 0.95$ for all buyers i. The extensive game form of the respective Feedback Game is depicted in Figure 4.3 on p. 46.

Elance employs a reputation mechanism with bi-directional feedback. As mentioned earlier, these systems are vulnerable to retaliatory feedback so that we choose $0.8 = \Delta(l, h) > \Delta(h, l) = 0.2$. That is, we demand that it is more expensive to falsely announce a high signal than it is to falsely announce a low signal. The costs of participation are set to $C = 0.5$.

As the probability for good play is the same for both buyers, it is sufficient to regard $g^i(\cdot|\cdot)$ for buyer $i = 1$ and set $\tau^2(\cdot,\cdot) = \tau^1(\cdot,\cdot)$. Note that different from the pure adverse selection scheme in Chapter 2, there are a number of possibilities with regard to the sequential order of computations. We therefore always refer to the respective equation from Section 4.2.2 with number and page.

We begin with the computation of the a priori quality level beliefs (Equation 4.23 on p. 42). As bad designers play b with probability 1 and good designers almost always play g, the prior quality beliefs are not far away from the prior type beliefs:

$$\begin{aligned} Pr(q^i = b) &= p_B(q^i = b) \cdot Pr(\theta = B) + p_G(q^i = b) \cdot Pr(\theta = G) \\ &= 1 \cdot 0.3 + 0.05 \cdot 0.7 \end{aligned}$$

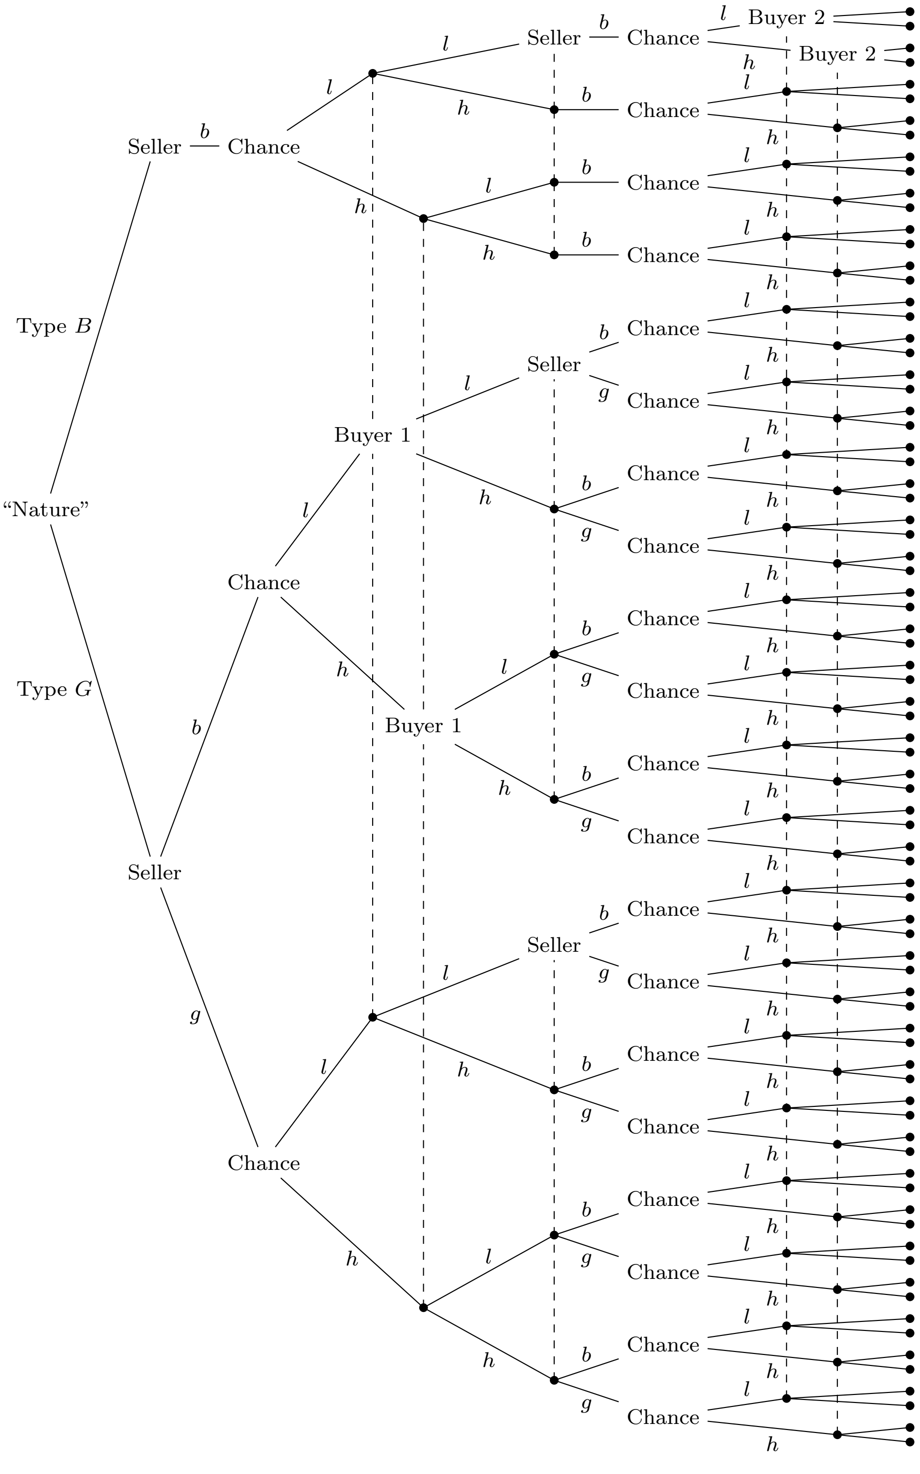

Figure 4.3: The complete game tree for the two-type example from Section 4.3. Note that the identifiers for the announcement of buyer 2 are left out due to lack of space.

$$\begin{aligned} &\simeq 0.34 \end{aligned}$$

$$\begin{aligned} Pr(q^i = g) &= 1 - Pr(q^i = b) \\ &\simeq 0.66 \end{aligned}$$

These quality beliefs are required to compute the prior signal beliefs (Equation 4.22 on p. 42):

$$\begin{aligned} Pr(s^i = l) &= f(l|b) \cdot Pr(q^i = b) + f(l|g) \cdot Pr(q^i = g) \\ &\simeq 0.9 \cdot 0.34 + 0.2 \cdot 0.66 \\ &\simeq 0.44 \end{aligned}$$

$$\begin{aligned} Pr(s^i = h) &= 1 - Pr(s^i = l) \\ &\simeq 0.56 \end{aligned}$$

Another "atomic" computation is that of a signal conditional on a web designer's type (Equation 4.26 on p. 43). Since the bad designer can only play b, the signal probability conditional on bad type equals the signal probability conditional on bad quality level:

$$\begin{aligned} Pr(s^i = l|\theta = B) &= f(l|b) \cdot p_B(q^i = b) + f(l|g) \cdot p_B(q^i = g) \\ &= 0.9 \cdot 1 + 0.2 \cdot 0 \\ &= 0.9 \end{aligned}$$

$$\begin{aligned} Pr(s^i = h|\theta = B) &= 1 - Pr(s^i = l|\theta = B) \\ &= 0.1 \end{aligned}$$

$$\begin{aligned} Pr(s^i = l|\theta = G) &= f(l|b) \cdot p_G(q^i = b) + f(l|g) \cdot p_G(q^i = g) \\ &= 0.9 \cdot 0.05 + 0.2 \cdot 0.95 \\ &\simeq 0.24 \end{aligned}$$

$$\begin{aligned} Pr(s^i = h|\theta = G) &= 1 - Pr(s^i = l|\theta = G) \end{aligned}$$

$$\simeq \quad 0.76$$

For the type beliefs conditional on a quality level (Equation 4.25 on p. 42), there is a similar situation as good quality cannot stem from a bad web designer. Note that it is perfectly possible that buyers *perceive* a high signal nonetheless.

$$\begin{aligned}
Pr(\theta = B|q^{r(i)} = b) &= \frac{p_B(q^{r(i)} = b) \cdot Pr(\theta = B)}{Pr(q^{r(i)} = b)} \\
&\simeq \frac{1 \cdot 0.3}{0.34} \\
&\simeq 0.88
\end{aligned}$$

$$\begin{aligned}
Pr(\theta = G|q^{r(i)} = b) &= 1 - Pr(\theta = B|q^{r(i)} = b) \\
&\simeq 0.12
\end{aligned}$$

$$\begin{aligned}
Pr(\theta = B|q^{r(i)} = g) &= \frac{p_B(q^{r(i)} = g) \cdot Pr(\theta = B)}{Pr(q^{r(i)} = g)} \\
&= 0
\end{aligned}$$

$$\begin{aligned}
Pr(\theta = G|q^{r(i)} = g) &= 1 - Pr(\theta = B|q^{r(i)} = g) \\
&= 1
\end{aligned}$$

The probability for a signal conditional on the other buyer's quality level is computed with Equation 4.24 on p. 42. Please note that although our example is symmetric with regard to the buyers, we cannot simply use buyer i's quality level because knowing the other buyer's quality level only yields information via the designer's type while knowing buyer i's quality level would influence her signal directly.

$$\begin{aligned}
Pr(s^i = l|q^{r(i)} = b) &= \quad Pr(s^i = l|\theta = B) \cdot Pr(\theta = B|q^{r(i)} = b) \\
&\quad + Pr(s^i = l|\theta = G) \cdot Pr(\theta = G|q^{r(i)} = b) \\
&\simeq 0.9 \cdot 0.88 + 0.24 \cdot 0.12 \\
&\simeq 0.82
\end{aligned}$$

$$Pr(s^i = h|q^{r(i)} = b) = 1 - Pr(s^i = l|q^{r(i)} = b)$$

$$\simeq \quad 0.18$$

$$\begin{aligned} Pr(s^i = l|q^{r(i)} = g) &= \quad Pr(s^i = l|\theta = B) \cdot Pr(\theta = B|q^{r(i)} = g) \\ &\quad + Pr(s^i = l|\theta = G) \cdot Pr(\theta = G|q^{r(i)} = g) \\ &\simeq 0.9 \cdot 0 + 0.24 \cdot 1 \\ &\simeq 0.24 \end{aligned}$$

$$\begin{aligned} Pr(s^i = h|q^{r(i)} = g) &= 1 - Pr(s^i = l|q^{r(i)} = g) \\ &\simeq 0.76 \end{aligned}$$

Before we can eventually compute $g^i(\cdot|\cdot)$, we have to calculate buyer $r(i)$'s posterior quality distribution given buyer i's signal (Equation 4.21 on p. 42). Since the good designer plays bad quality with the same probability for either buyer and because it is a prior belief, we can set $Pr(q^{r(i)}) = Pr(q^i)$:

$$\begin{aligned} Pr(q^{r(i)} = b|s^i = l) &= \frac{Pr(s^i = l|q^{r(i)} = b) \cdot Pr(q^{r(i)} = b)}{Pr(s^i = l)} \\ &\simeq \frac{0.82 \cdot 0.34}{0.44} \\ &\simeq 0.63 \end{aligned}$$

$$\begin{aligned} Pr(q^{r(i)} = g|s^i = l) &= 1 - Pr(q^{r(i)} = b|s^i = l) \\ &\simeq 0.37 \end{aligned}$$

$$\begin{aligned} Pr(q^{r(i)} = b|s^i = h) &= \frac{Pr(s^i = h|q^{r(i)} = b) \cdot Pr(q^{r(i)} = b)}{Pr(s^i = h)} \\ &\simeq \frac{0.18 \cdot 0.34}{0.56} \\ &\simeq 0.11 \end{aligned}$$

$$\begin{aligned} Pr(q^{r(i)} = g|s^i = h) &= 1 - Pr(q^{r(i)} = b|s^i = h) \\ &\simeq 0.89 \end{aligned}$$

Eventually, we can compute buyer i's signal posterior belief regarding buyer $r(i)$'s signal (Equation 4.20 on p. 42) which are required for LP 2. Due to

	$\tau^i(l,l)$	$\tau^i(l,h)$	$\tau^i(h,l)$	$\tau^i(h,h)$	
min	0.28	0.15	0.15	0.41	
s. t.	0.64	0.36	-0.64	-0.36	$\geq 0.8 + \epsilon$
	-0.28	-0.72	0.28	0.72	$\geq 0.2 + \epsilon$
	0.64	0.36			$\geq 0.5 + \epsilon$
			0.28	0.72	$\geq 0.5 + \epsilon$
	≥ 0	≥ 0	≥ 0	≥ 0	

Figure 4.4: The Linear Program that is used to compute the payment scheme for the Elance example.

rounding, these values differ slightly from the actual values which is why we give more exact values in brackets:

$$\begin{aligned}
g^i(l|\,l) &= Pr(s^{r(i)} = l|s^i = l) \\
&= f(l|b) \cdot Pr(q^{r(i)} = b|s^i = l) + f(l|g) \cdot Pr(q^{r(i)} = g|s^i = l) \\
&\simeq 0.9 \cdot 0.63 + 0.2 \cdot 0.37 \\
&\simeq 0.64 \qquad (0.6482336)
\end{aligned}$$

$$\begin{aligned}
g^i(h|\,l) &= Pr(s^{r(i)} = h|s^i = l) \\
&= 1 - g(l|\,l) \\
&\simeq 0.36 \qquad (0.3517664)
\end{aligned}$$

$$\begin{aligned}
g^i(l|\,h) &= Pr(s^{r(i)} = l|s^i = h) \\
&= f(l|b) \cdot Pr(q^{r(i)} = b|s^i = h) + f(l|g) \cdot Pr(q^{r(i)} = g|s^i = h) \\
&\simeq 0.9 \cdot 0.11 + 0.2 \cdot 0.89 \\
&\simeq 0.28 \qquad (0.27027851)
\end{aligned}$$

$$\begin{aligned}
g^i(h|\,h) &= Pr(s^{r(i)} = h|s^i = h) \\
&= 1 - g(l|\,h) \\
&\simeq 0.72 \qquad (0.72972149)
\end{aligned}$$

The resulting Linear Program formulation and the payment scheme are depicted in Figure 4.4 and Figure 4.5, respectively. The expected budget for a single buyer which equals the a priori expected payment is 0.87.

		Buyer $r(i)$	
		l	h
Buyer i	l	1.73	0
	h	0	0.92

Payment Scheme

Figure 4.5: The payment scheme with $C = 0.5$, $\Delta(h, l) = 0.8$ and $\Delta(l, h) = 0.2$.

4.4 Discussion

In our model, both types and signals are discrete. In settings where they correspond to intervals this may lead to a situation in which moral hazard plays a role within each quality interval. Sellers are led to produce at a quality close to the bottom of the interval since this involves lower costs than, for example, in the middle of it. Two solutions are worth considering: first, one could investigate continuous formulations of the mechanism or, second, one could use a finer discretization resulting in more types. In the latter situation the effect of sellers being pushed to the bottom of the respective interval becomes less harmful as with more types, the intervals become smaller.

While the focus of this thesis is on truthful feedback, it is important to note that mixed settings with noisy observations are challenging even if honest feedback is not an issue. As mentioned earlier, the reputation mechanism has to fulfill two roles at the same time: first, it shall *signal* the inherent abilities of the sellers to future buyers and, second, it shall *sanction* sellers that do not cooperate. The usual approach is to publish previous transaction outcomes. One of the design issues is whether the center should publish the entire history of transaction outcomes or only part of it. If k is the number of recent outcomes that are published, it is not clear how k should be chosen. For pure adverse selection settings, a larger k is always preferred since it reveals more information of the seller's type. Pure sanctioning mechanisms with imperfect monitoring, however, cannot sustain cooperation if k is too large. The intuition is that in perfect monitoring settings, a single play of the cheating action reveals the seller as a "strategic" type, i. e. a player who is not committed to cooperative play. When monitoring is imperfect, however, cooperative play can result in negative observations so that a single play of the cheating action does not immediately reveal the seller's strategic nature. Cripps, Mailath and Samuelson [2004] show that if the buyers observe the entire history of past transaction outcomes, every rational seller cooperates in the beginning to establish a good reputation only to exploit it thereafter through occasional cheating. This eventually reveals his strategic type and drives him out of the market. It is therefore the objective of a pure moral hazard reputation mechanism to adjust the length of the publication record such that a seller can build a reputation but does not have enough transactions to exploit it. To study this trade-off between the signaling and the sanctioning role of reputation mechanisms in mixed settings is an open question for future research.

Chapter 5

Experimental Evaluation

In this chapter, we describe our experimental results that suggest that the "feedback plug-in" of Chapter 4 is feasible for real-world reputation mechanisms. More specifically, the results for both computational complexity and the expected budget are promising. In addition to practicality considerations, the experiments give us a better understanding of the dynamics that govern peer-based feedback elicitation.

The chapter is organized as follows: in Section 5.1 we motivate a measurement for the degree of seller cooperation that is present in a setting which is important because we are interested in how cooperation influences the required budget. In Section 5.2 we describe the experimental setup. In Section 5.3 we evaluate the running time for the computation of a single buyer's payment scheme and show that even for large signal sets, the computational complexity is not a limiting factor for application. Section 5.4 is devoted to the behavior of the expected budget with a particular focus on the influence of seller cooperation.

5.1 Definition of Cooperation Value

A crucial parameter of the setting are the optimal seller strategies $p_t(q^i)$. Clearly, these externally given strategies determine the degree of seller cooperation of a specific setting. For example, if $p_t(q^i = q_t) = 1$ for all θ_t, there is full seller cooperation in the setting. In order to evaluate how the expected budget changes with regard to the seller's cooperation, we therefore have to define a measurement for this. It is important to note that the $p_t(q^i)$ are given to us by the external reputation mechanism. We are interested in the cooperation value that is present in a population of sellers. We denote this value with

$$\eta. \tag{5.1}$$

We begin with the definition of the cooperation value of a single seller type θ_t for a single buyer i and denote this value with

$$\eta_t^i. \tag{5.2}$$

For $t \geq 2$, we have a number of requirements on η_t^i:

1. We want it to be normalized to 1, i. e. a fully cooperative seller should have a cooperation value of 1.

2. A seller of type θ_t should have a cooperation value of 1 if he plays his highest possible action with probability 1, i. e. $p_t(q^i = q_t) = 1$.

3. For a seller of type θ_2, the cooperation value should equal the probability for playing his "good" quality level.

Note that the lowest seller type θ_1 is left out of the definition as he is a commitment type, i. e. not a player in the actual sense, and should therefore not influence the cooperation value. Requirements 1 and 2 induce the nice property that $\eta_t^i = 1$ coincides with the seller behavior in the pure adverse selection setting from Chapter 2. Requirement 3 is motivated by the interpretation of the two type mixed setting as a pure moral hazard setting with cheating types (compare Section 3.4).

A definition that incorporates all of these requirements is the one that puts equal distances in between two "neighboring" seller actions. Using this rule, the quality actions of a θ_4 seller, for example, have the weights 0, $\frac{1}{3}$, $\frac{2}{3}$ and 1 (in increasing order). For a seller of type θ_t, $t \geq 2$, and buyer i we obtain:

$$\eta_t^i = \sum_{l=1}^{t} \frac{l-1}{t-1} p_t(q^i = q_l). \tag{5.3}$$

Our approach for the seller population is to weigh the cooperation value of a single seller type with the prior probability that it occurs. Since θ_1 is left out of the computation of η, we have to normalize the prior type beliefs as given in Equation 5.4. For reasons of clarity, we do not introduce another variable but denote the normalized type priors with $Pr_{-1}(\theta)$, so that for all $t \geq 2$:

$$Pr_{-1}(\theta_t) = \frac{Pr(\theta_t)}{1 - Pr(\theta_1)} \tag{5.4}$$

For the seller population and a single buyer we thus have:

$$\begin{aligned}\eta^i &= \sum_{t=2}^{T} Pr_{-1}(\theta_t) \cdot \eta_t^i \\ &= \sum_{t=2}^{T} Pr_{-1}(\theta_t) \cdot \left(\sum_{l=2}^{t} \frac{l-1}{t-1} p_t(q^i = q_l) \right).\end{aligned} \tag{5.5}$$

The setting includes two buyers with possibly different cooperation values, so that we weigh these two cooperation values using the arithmetic mean of the respective η^i:

$$\eta = \frac{\eta^1 + \eta^2}{2}. \tag{5.6}$$

In Appendix A we explain our algorithm to compute random values for $p_t(q^i)$ given η.

5.2 Experimental Setup

We consider settings in which every type corresponds to a signal, i. e. $M = T$. The default value for the size of the signal set is $M = 5$. The conditional signal observations are chosen in the same way as by Jurca and Faltings [e. g., Jurca and Faltings, 2006]:

$$f(s_m|q_l) = \begin{cases} 1 - \varepsilon & m = l \\ \varepsilon/(M-1) & m \neq l \end{cases} \tag{5.7}$$

Clearly, higher values of ε correspond to more noise in the buyers' observations. The default value for this noise parameter is $\varepsilon = 0.1$. The prior type beliefs are uniformly distributed and all external lying benefits $\Delta(s_j, s_h)$, $h \neq j$, are set to 0.8. The buyers' costs of participation are set to $C = 0.5$.

The default value for seller cooperation is $\eta = 0.95$. Please see Appendix A for a detailed description of the seller strategies' construction given η. We average over 500 randomly generated settings.

5.3 Running Time

It is well known that the computational complexity of solving Linear Programs is polynomial. The Simplex method, however, which has exponential worst-case behavior, is frequently more efficient than guaranteed polynomial algorithms such as interior point methods [e. g., Schrijver, 1998]. In order to determine

M	**time (in ms)**	M	**time (in ms)**	M	**time (in ms)**
2	1.58	12	14.20	22	87.89
4	2.20	14	21.66	24	114.50
6	3.28	16	32.70	26	149.84
8	5.44	18	47.42	28	189.17
10	9.00	20	65.14	30	239.05

Table 5.1: Average CPU time for the computation of the payment scheme with different values for the size of the signal set M.

whether the feedback plug-in from Chapter 4 is feasible for real-world application, we empirically evaluate it on a customary computer with a 1.6 GHz CPU. We use a small program written in Python[1] to compute the beliefs and solve the LP with lpsolve[2]. Table 5.1 depicts the running time for the computation of a single payment scheme depending on the size of the signal set. Observe that the size of the LP's matrix is independent of the size of the type set T. Considering the low numbers even for large signal sets, we firmly believe that computational complexity is not a limiting factor for application.

5.4 Expected Budget

An important characteristic of the computed solution is its expected budget. See Figure 5.1 for the behavior of the expected budget for different values of noise in the buyers' observations. At first sight, the figure may suggest that the required budget becomes lower with larger signal sets. To see why this is misleading take a look at the procedure we use to compute the conditional signal probabilities (Equation 5.7):

$$f(s_m|q_l) = \begin{cases} 1-\varepsilon & m = l \\ \varepsilon/(M-1) & m \neq l \end{cases} \tag{5.8}$$

For a constant ε and a growing signal set size M, the amount of noise that is assigned to signal s_m with $m \neq l$ becomes smaller. While the procedure is sufficient to compare the expected budget for a constant signal set size, it is inadequate for experiments that study the budget behavior with regard to M. Nevertheless, Figure 5.1 also shows that the expected budget is independent of the size of the signal set for perfect monitoring, i. e. $\varepsilon = 0$, and we conjecture that this holds true for a more sophisticated procedure that adequately generalizes a certain noise level.

---

[1] www.python.org

[2] lpsolve.sourceforge.net

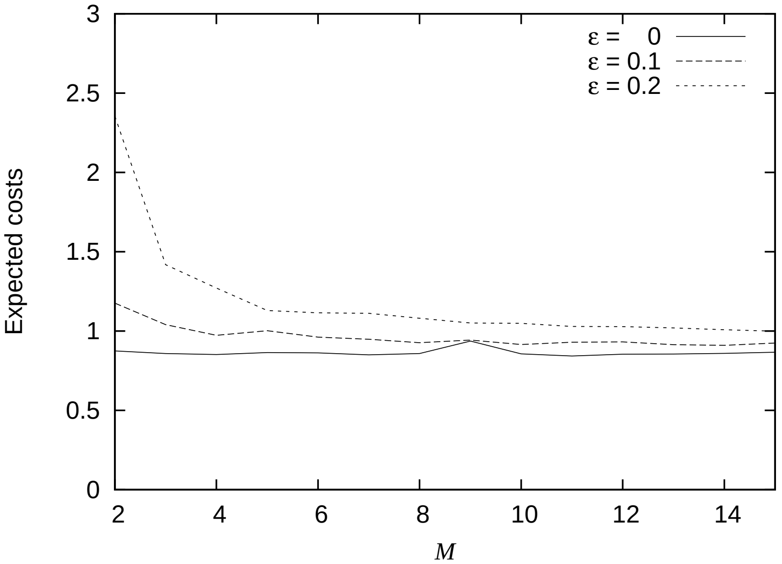

Figure 5.1: The expected budget depending on the size of the signal set M and for different values of noise. We averaged over 200 instantiations. Please note that due to the random generation of type beliefs and seller strategies, the expected budget slightly variates even for $\varepsilon = 0$.

5.4.1 Dependent on Seller Cooperation

Figure 5.2 depicts the expected budget over the entire space of seller cooperation and different noise levels. Higher seller cooperation is generally cheaper for the feedback mechanism which is good for the mechanism designer because it means that there is no conflict of interest between the external reputation mechanism and the feedback plug-in: both want the seller to cooperate as much as possible. To understand why cooperation values close to 0 lead to an extraordinary large budget, consider the extreme case of 0 seller cooperation: $\eta = 0$. The only possible configuration for this situation is that all seller types play their lowest actions q_1. For this case, however, the two buyers' signal observations cannot be stochastically relevant for each other because the buyers' posterior signal distributions are known a priori. Put differently, every seller plays q_1, so that buyer i cannot infer any information about the seller's type once she received her signal. The posterior type beliefs, however, are the information that leads buyer i to update her signal posteriors. Now if η is close to 0, the two signal observations are stochastically relevant again but the signal posteriors are close to each other. That is, $g^i(\cdot | s_j)$ is similar to $g^i(\cdot | s_h)$ for two possible signal observations $s^i = s_j$ and $s^i = s_h$. The budget becomes large for close posterior beliefs because the honesty constraints require that the expected payment for the honest signal announcement is larger by $\Delta(\cdot, \cdot)$. As the expected payment

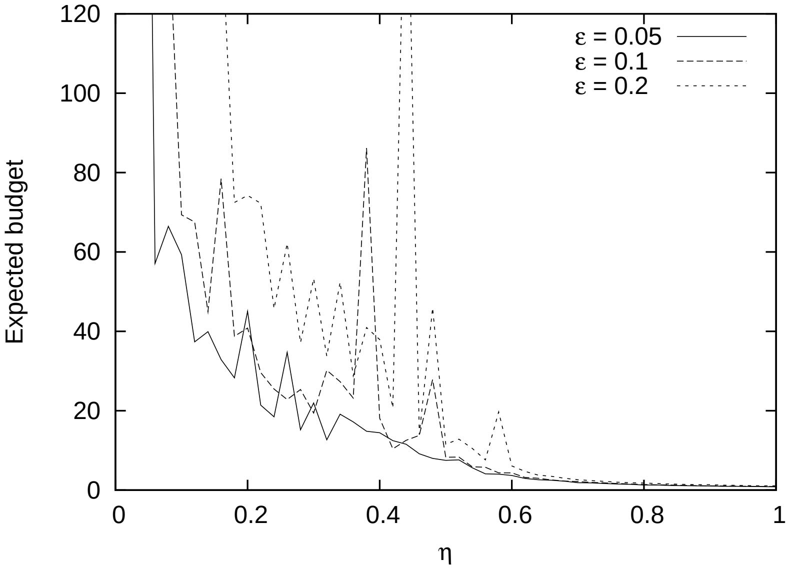

Figure 5.2: Expected budget depending on seller cooperation for different values of noise in the observations.

is composed of $\tau^i(\cdot,\cdot)$ and the signal posteriors, the former must be larger the closer the latter to lift expected payment over $\Delta(\cdot,\cdot)$.

Note that the high variance around $\eta = 0.5$ is due to the way we choose the seller strategy for a given type (compare Appendix A). For $\eta = 0.5$, the seller strategies that are computed for a single seller type are equally distributed over all q_l which again makes the seller types difficult to separate and results in signal posteriors $g^i(\cdot|s_j)$ that are close to each other and therefore costly to separate. For values of η that are larger than 0.6, however, there are no outliers any more. Note that cooperation values as low as 0.6 are unlikely to be an issue in real-world applications. To see why, consider the moral hazard interpretation of the mixed setting with two types, i. e. the "normal" type and the "cheating" type (compare Discussion in Section 3.4). A cooperation value of 0.6 means that the normal type plays his cooperating action with a probability of only 60% and cheats otherwise. In Figure 5.3 we have therefore plotted the feedback mechanism's behavior for values between 0.9 and 1 which we expect to be the relevant interval. Observe that even for $\eta = 0.9$ and a high noise level of $\varepsilon = 0.2$, the required budget is less than 30% larger than in a setting with full cooperation.

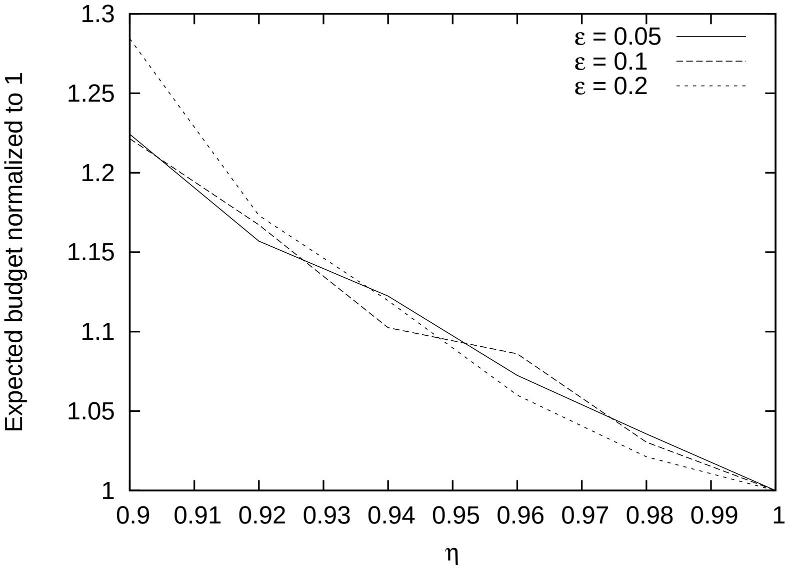

Figure 5.3: Relevant part of the plot for the expected budget depending on seller cooperation. The expected budget is normalized to 1 which is the expected budget of the full cooperation setting.

5.4.2 Depending on Noise

While studying the influence of seller cooperation in isolation provides some insights, in practice it will be coupled with the noise in the observations. For a pure moral hazard setting, Dellarocas finds that the level of cooperation that can be achieved by a reputation mechanism is bounded away from the optimum with the degree of noise in perception [2005]. We conjecture that the same holds true for mixed settings and therefore link the seller's cooperation with the noise in the buyers' perception. More specifically, we set $\eta = 1 - \varepsilon$. While we are aware that the exact relation between noisy observations and seller cooperation can be different, we believe that it gives us a more realistic view on the feedback mechanism's behavior in practice. See Figure 5.4 for a plot of this scenario. Observe that the expected budget rises smoothly for low noise values.

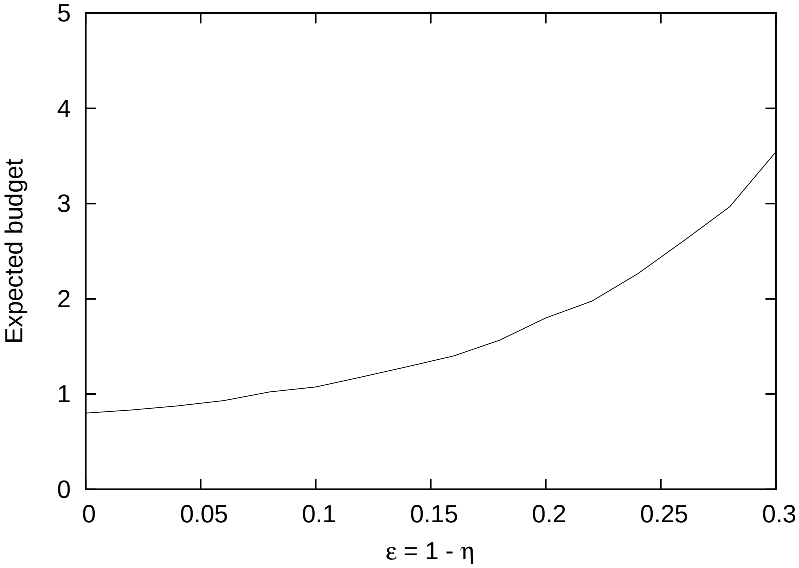

Figure 5.4: The expected budget with imperfect monitoring linked to seller cooperation.

Chapter 6

Multiple Equilibria

In our earlier work [2009] we have shown how, in a pure adverse selection setting that generalizes that of Chapter 2, honest reporting can be made the unique equilibrium of every induced reporting game if the center has the possibility for a single signal acquisition. Unfortunately, non-truthful equilibria are essentially unavoidable in purely peer-based feedback elicitation [Jurca and Faltings, 2005], i. e. if the signal acquisition on behalf of the center is not possible.

In the presence of multiple equilibria, agents have to coordinate on one of these. Schelling [1960] argued that, abstracted away from the strategic definition of the game, there can be so-called "focal points". That is, certain equilibria are more likely to be chosen because they appear special or relevant. For the reporting game in Chapter 2, for example, this can be asserted for the honest equilibrium. However, the results of Chapter 4 can be regarded as being somewhat weaker than those of Chapter 2: through the introduction of the seller as a strategic player, the best response conditions no longer demand that given one buyer is honest, the other one is honest as well. We want to know whether it is possible to achieve this for the mixed setting of Chapter 4 as well. That is, the objective could be informally stated as: "if one buyer is truthful, both the other buyer and the seller should be truthful". In the following, we exemplify how this objective could be reached. The remainder of this section uses a simple example with only two types and perfect private monitoring to construct a payment scheme that fulfills the objective.

Say there are two types B and G. The seller's actions are $q_1 = b$ and $q_2 = g$. The respective costs are denoted $c_G(g) = c_2(q_2)$, $c_G(b) = c_2(q_1)$ and $c_B(b) = c_1(q_1)$. The buyers' actions are reporting high h and a low l quality, respectively. A buyer's report is truthful if and only if she announces l after reception of b and h after reception of g. The prior type beliefs are $Pr(\theta = B) = 0.3$ and $Pr(\theta = G) = 0.7$. $\Delta(h, l) = \Delta(l, h) = 0.1$ and $C = 1$.

6.1 Efficiency Conditions

The assumption regarding the seller valuations are the same as in Chapter 4. That is, given the center could directly observe the transaction outcomes, all sellers would play their optimal actions $p_t(q^i)$. As monitoring is perfect, we assume that $p_G(q^i = g) = 1$, i. e. the seller finds it optimal to play his highest action given truthful feedback was not an issue. Furthermore, we neglect the discount factor for this simple example and, in particular, assume that $v^s(l,h|r) = v^s(h,l|r)$. Since in our example, only the good type G has an actual choice, the following holds by assumption (we assume the seller is *strictly* better off playing his high action given truthful feedback):

$$\begin{aligned} -2c_2(g) + v^s(h,h|r) &> -c_2(b) - c_2(g) + v^s(l,h|r) \\ -2c_2(g) + v^s(h,h|r) &> -2c_2(b) + v^s(l,l|r) \end{aligned} \tag{6.1}$$

Equation 6.1 can be simplified to:

$$\begin{aligned} v^s(h,h|r) - v^s(l,h|r) &> c_2(g) - c_2(b) \\ v^s(h,h|r) - v^s(l,l|r) &> 2\left(c_2(g) - c_2(b)\right) \end{aligned} \tag{6.2}$$

6.2 Reducing the Game Tree

In contrast to Chapter 4, we do not go forward by assuming that *all* other players are truthful, fix their strategies and construct payments that make truthfulness by the left-out player a best response. Instead, we assume truthfulness by only *one* buyer and construct payments that induce truthfulness by the *two* left-out players. Note that the constraints of LP 2 (p. 44) have to hold nonetheless. LP 2 demands that for both buyers i:

$$\begin{aligned} \tau^i(h,h) - \tau^i(l,h) &\geq 0.1 + \epsilon \\ \tau^i(l,l) - \tau^i(h,l) &\geq 0.1 + \epsilon \\ \tau^i(l,l) &\geq 1 + \epsilon \\ \tau^i(h,h) &\geq 1 + \epsilon \end{aligned} \tag{6.3}$$

Without limitation of generality, we take the first buyer to report honestly and reduce the game tree accordingly. See Figure 6.1 for the reduced game tree that results if buyer 1 is committed to honest play .

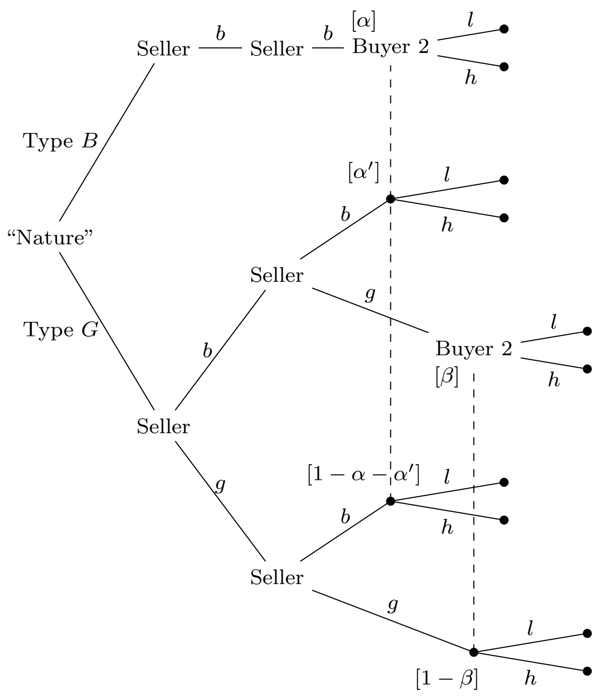

Figure 6.1: Reduced game tree for the simple example with perfect monitoring.

6.2.1 Buyer 2 received a b

We begin with the rational considerations of the second buyer when she received a b. We denote the beliefs at this information set with α, α' and $\alpha'' = 1-\alpha-\alpha'$. From the observation of a b alone, buyer 2 does not know for sure whether the seller is of type B—and could not play anything else—or of type G—and decided to play his low action. What she knows, however, is that a priori the seller is of type B with probability $Pr(\theta = B) = 0.3$. As she received a b, 0.3 became a lower bound on her belief that the seller is of type B:

$$\alpha = Pr(\theta = B|b) \geq 0.3. \tag{6.4}$$

We want that buyer 2 announces an l given she received a b what she will do if her expected payment for an l announcement is strictly larger than announcing h. The expected utility for l and h is:

$$\begin{aligned} E(l^2|\, b^2) &= v^2(b) + \alpha\tau^2(l,l) + \alpha'\tau^2(l,l) + (1-\alpha-\alpha')\tau^2(l,h) \\ &= v^2(b) + \underbrace{(\alpha+\alpha')}_{x}\tau^2(l,l) + \underbrace{(1-\alpha-\alpha')}_{y}\tau^2(l,h) \\ &= v^2(b) + x\tau^2(l,l) + y\tau^2(l,h) \end{aligned} \tag{6.5}$$

and

$$\begin{aligned} E(h^2|\, b^2) &= v^2(b) + \alpha\tau^2(h,l) + \alpha'\tau^2(h,l) + (1-\alpha-\alpha')\tau^2(h,h) \\ &= v^2(b) + \underbrace{(\alpha+\alpha')}_{x}\tau^2(h,l) + \underbrace{(1-\alpha-\alpha')}_{y}\tau^2(h,h) \\ &= v^2(b) + x\tau^2(h,l) + y\tau^2(h,h) \end{aligned} \tag{6.6}$$

We can now set $E(l^2|\, b^2) - E(h^2|\, b^2) \geq \Delta(l,h) + \epsilon$:

$$\begin{gathered} E(l^2|\, b^2) - E(h^2|\, b^2) \geq 0.1 + \epsilon \\ \Leftrightarrow \\ x\tau^2(l,l) + y\tau^2(l,h) - x\tau^2(h,l) - y\tau^2(h,h) \geq 0.1 + \epsilon \\ \Leftrightarrow \\ x(\tau^2(l,l) - \tau^2(h,l)) - y(\tau^2(h,h) - \tau^2(l,h)) \geq 0.1 + \epsilon \\ \Leftrightarrow \\ \underbrace{x}_{\geq Pr(\theta=B)=0.3} \underbrace{\big(\tau^2(l,l) - \tau^2(h,l)\big)}_{>0} - \underbrace{y}_{\leq 1-Pr(\theta=B)=0.7} \underbrace{\big(\tau^2(h,h) - \tau^2(l,h)\big)}_{>0} \geq 0.1 + \epsilon \end{gathered} \tag{6.7}$$

Therefore, if we can ensure that Equation 6.8 holds, every b is honestly announced.

$$0.3\,\tau^2(l,l) - 0.3\,\tau^2(h,l) - 0.7\,\tau^2(h,h) + 0.7\,\tau^2(l,h) \geq 0.1 + \epsilon \tag{6.8}$$

Observe that we are very pessimistic in our estimation. Equation 6.8 ensures an honest announcement of b even in those cases in which sellers of the good type played g with probability 1 for the first buyer followed by a b with probability 1. One can possibly improve on these bounds.

6.2.2 Buyer 2 received a g

We can now further prune the game tree such that there is only one information set left for buyer 2, namely the one when she receives a g (compare Figure 6.2). In contrast to the situation where buyer 2 received a b, she can now be certain that the seller is of type G. However, this also means that we cannot use the prior type beliefs as lower bounds. The way forward is thus a different one than in the b case. Instead of reasoning over buyer 2, we show that if the seller played g for the second buyer, it is only reasonable to assert that he did play g for the first buyer, as well. In order to show this, we slightly abuse the notation and denote $a^2(h|g)$ the probability, i. e. behavioral strategy, that buyer 2 announces a high signal given she received a g. Furthermore, to denote that the seller plays b and g for the first and second buyer, respectively, we simply concatenate the two letters. In order to be part of a perfect Bayesian equilibrium, playing bg must be at least as good for the seller as playing gb. Otherwise, bg cannot be sequentially rational. We begin by writing down the respective seller utilities for these two actions. Note that, as before, we denote the players with a superscript. The expected utility for the seller given he plays gb is:

$$U^s(gb) = -c_2(b) - c_2(b) + v^s(l, h|r) \tag{6.9}$$

For the seller's utility when he plays bg, we need to incorporate the behavioral strategy of buyer 2:

$$\begin{aligned}
U^s(bg) =& a^2(h|g) \cdot (-c_2(b) - c_2(g) + v^s(l, h|r)) \\
& + (1 - a^2(h|g)) \cdot (-c_2(b) - c_2(g) + v^s(l, l|r)) \\
=& - c_2(b) - c_2(g) + v^s(l, l|r) \\
& + a^2(h|g) \cdot c_2(b) + a^2(h|g) \cdot c_2(g) - a^2(h|g) \cdot v^s(l, l|r) \\
& - a^2(h|g) \cdot c_2(b) - a^2(h|g) \cdot c_2(g) + a^2(h|g) \cdot v^s(l, h|r) \\
=& - c_2(b) - c_2(g) + v^s(l, l|r) + a^2(h|g) \cdot v^s(l, h|r) - a^2(h|g) \cdot v^s(l, l|r)
\end{aligned} \tag{6.10}$$

We can then formulate the set of buyer 2's behavioral strategies under which bg is at least as good as gb:

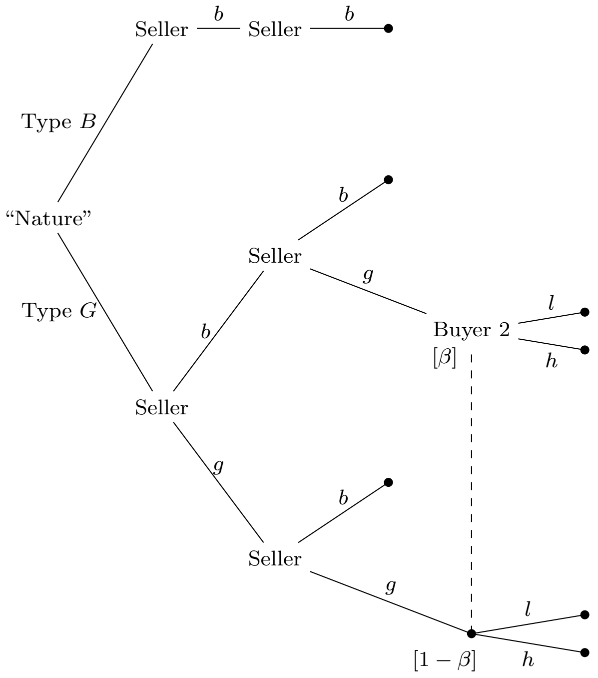

Figure 6.2: Further reduced game tree for the simple example with perfect monitoring.

$$
\begin{aligned}
U^s(bg) &\geq U^s(gb) \\
&\Leftrightarrow \\
-c_2(b) - c_2(g) + v^s(l,l|r) & \\
+a^2(h|g) \cdot v^s(l,h|r) & \\
-a^2(h|g) \cdot v^s(l,l|r) &\geq -c_2(b) - c_2(g) + v^s(l,h|r) \\
&\Leftrightarrow \\
v^s(l,l|r) + a^2(h|g) \cdot v^s(l,h|r) - a^2(h|g) \cdot v^s(l,l|r) &\geq v^s(l,h|r) \\
&\Leftrightarrow \\
a^2(h|g) \cdot (v^s(l,h|r) - v^s(l,l|r)) &\geq v^s(l,h|r) - v^s(l,l|r) \\
&\Leftrightarrow \\
a^2(h|g) &= 1
\end{aligned} \tag{6.11}
$$

This means that from the seller's point of view playing *bg* is only as good as playing *gb* if buyer 2 reports honestly with probability 1. Then, however, playing *gg* is strictly better by assumption as this is the case where both buyers are

		Buyer 1	
		l	h
Buyer 2	l	1	0
	h	0	1

Payment scheme from Chapter 4

		Buyer 1	
		l	h
Buyer 2	l	1	0.71
	h	0	1

Robust payment scheme

Figure 6.3: The payment schemes from Chapters 4 and 6, respectively, for the perfect monitoring example setting.

honest. Receiving a g, buyer 2 can thus infer that the seller played gg, so that she maximizes her utility by honestly announcing h (given $\tau^2(h,h) - \tau^2(l,h) \geq \Delta(h,l) + \epsilon$ which is required from the old scheme). Thus, we have ensured that in this reduced tree, the only reasonable equilibrium is the truthful equilibrium.

6.3 Results

The entire LP for this small example game is thus:

LP 3.

$$\begin{aligned}
\min \quad & B = 0.3\,\tau^2(l,l) + 0.7\,\tau^2(h,h) \\
s.t. \quad & 0.3\,\tau^2(l,l) - 0.3\,\tau^2(h,l) - 0.7\,\tau^2(h,h) + 0.7\,\tau^2(l,h) \geq 0.1 + \epsilon \\
& \tau^2(h,h) - \tau^2(l,h) \geq 0.1 + \epsilon \\
& \tau^2(l,l) - \tau^2(h,l) \geq 0.1 + \epsilon \\
& \tau^2(l,l) \geq 1 + \epsilon \\
& \tau^2(h,h) \geq 1 + \epsilon \\
& \tau^i(s_j, s_k) \geq 0 \quad \forall s_j, s_k \in S
\end{aligned}$$

The resulting payment scheme together with the payment scheme of Chapter 4 is depicted in Figure 6.3. The expected budget is 1 for both the payment scheme from this section and the scheme from Chapter 4. It is clear that the identical budgets are due to perfect monitoring as the $\tau^2(l,h)$ case never occurs in equilibrium. In fact, it can be interpreted as a credible threat to the seller which ensures that it never happens. When there is noise in the setting, these threats obviously implicate higher costs.

Clearly, there is still a lot to do here. First, one needs to study whether the exemplified procedure extends to general settings with noise in perception and sellers that are not fully cooperative. A first result that we have left out due to time constraints is that a similar procedure can be found for settings with noise while we still assume a fully cooperative seller. For the future, we plan to construct a robust payment scheme for the general setting from Chapter 4.

Chapter 7

Conclusion and Outlook

Reputation mechanisms offer an effective way to prevent market failure in online economies. Through the publication of past experiences regarding the quality of products or the trustworthiness of market participants, they enable prospective customers to make better-informed choices. The economic value of these published experiences, however, raises questions regarding the trustworthiness of the mechanisms themselves. While existing systems assume that the privately monitored experiences are honestly reported, there is evidence from both game theory and empirical studies that the agents' reports are biased.

Recently, Miller, Resnick and Zeckhauser [2005] have shown that the truthful elicitation of feedback is possible for reputation mechanisms that address pure adverse selection, such as Amazon Reviews. Their "peer-prediction method" compares an agent's quality announcement with that of another agent, called the agent's reference reporter. As both agents experienced the same product, there exists a payment scheme that pays the agent depending on how well her announcement "predicts the announcement of her reference reporter".

In this thesis we studied whether the peer-prediction framework can be modified such that it elicits truthful feedback in the presence of moral hazard, i. e. if the seller chooses his actions strategically. For a pure moral hazard setting, motivated by the one at eBay, we find that there is no peer-based feedback mechanism that makes honest reporting a best response to truthful play by all other players. For a combined setting, with both adverse selection and moral hazard, we retrieve a positive result and construct a "feedback plug-in" that can be integrated into reputation mechanisms that are situated in mixed settings. Our experimental findings strongly suggest that it is feasible for practical application regarding both computational complexity and expected budget. Furthermore, we exemplify a procedure that can be used to eliminate non-truthful equilibria that are unavoidable in peer based feedback mechanisms.

Other than the setting-specific aspects that we have already mentioned in the discussion sections of the respective chapters, we plan to study whether the sellers in the mixed setting can be paid directly dependent on feedback. So far, we have made two crucial assumptions: we are given an external reputation mechanism that induces sellers to be truthful as long as feedback is guaranteed to be honest and sellers are long-lived. While the latter is the usual assumption for moral hazard mechanisms, it clearly is a simplification of reality. At online auction sites, for example, not every seller can be considered sufficiently long-lived to be incentivized by future gains or losses due solely to the publication of past outcomes. This is further emphasized in the presence of cheap pseudonyms, i. e. if sellers can create new accounts once they have ruined their reputation. We are interested in whether one can drop these assumptions if the center acts as a trusted third party not only for feedback payments but also for the payment from the buyer to the seller of a transaction. Dellarocas [2003] showed that this is feasible under a number of rather strong assumptions with regard to the determination of prices and the assumption that feedback is honest. In the future, we want to study whether we can use our feedback plug-in to create a fully incentivized payment framework.

Appendix A

Construction of Random Seller Strategies

To construct meaningful random values for $p_t(q^i)$ given the value for seller cooperation η is not trivial. We begin with the explanation how we choose each η_t^i given η and, thereafter, we show how to construct the seller strategies for a single seller type given η_t^i. We use the same seller strategies for both buyers, i. e. $\eta^i = \eta$, as it is not reasonable that a seller's strategy changes considerably from one buyer to the other. Note that it is nonetheless perfectly possible that the two quality levels that are played differ.

We begin to explain our procedure to determine all η_t given the overall cooperation value η. The only procedure that is independent of the prior type probabilities is to assign $\eta_t = \eta$ for all θ_t. This procedure, albeit correct, involves no randomness, so that we use a different algorithm that we briefly sketch in the following:

1. Set up a temporary variable "temp" that contains $T - 1$ random values all of which are between η and 1.
2. Compute the cooperation value that would be implemented if "temp" were the distribution of η_t, taking into account the scaled type belief $Pr_{-1}(\theta)$. Let us denote this implemented cooperation value with "coop value". Observe that "coop value" must be larger than η (it equals η if all values in "temp" are η).
3. Divide η by "coop value" and multiply the temporary variable "temp" with this fraction.

The resulting values for η_t are all between 0 and 1 and implement the cooperation value of the overall setting η. That is, they abide to Equation 5.5 on p. 55

which we give in the buyer-symmetric version, i. e. $\eta^i = \eta^{3-i}$:

$$\eta = \sum_{t=2}^{T} Pr_{-1}(\theta_t) \cdot \eta_t. \tag{A.1}$$

With a rule for finding η_t given η and $Pr_{-1}(\theta)$, we can turn to the problem of finding the seller strategies $p_t(q^i = q_l)$ given η_t. The naïve approach is to assign the total value of η_t to the highest possible action, which is weighed with 1, and assign $1 - \eta_t$ to the lowest action which is weighed with 0. If we use this assignment rule, however, it seems we loose too much of the game's structure since the probability of all other quality actions, in between the lowest and the highest, are set to 0. Finding assignments with full support is non-trivial and we search for suitable seller strategies with a small LP. For a given η_t and seller type θ_t, we invoke LP 4:

LP 4.

$$\begin{array}{lll}
\max & u & \\
s.\,t. & p_t(q^i = q_l) \geq u & l = 1, \ldots, t \\
 & \displaystyle\sum_{l=1}^{t} p_t(q^i = q_l) = 1 & \\
 & \displaystyle\sum_{l=2}^{t} \frac{l-1}{t-1}\, p_t(q^i = q_l) = \eta_t &
\end{array}$$

The LP ensures that a strategy with full support is chosen whenever it is possible, i. e. if $\eta_t \neq 0$ and $\eta_t \neq 1$. The solution for θ_4 and $\eta_4 = 0.8$, for example, is

$$p_t(q^i) = \begin{pmatrix} 0.1 \\ 0.1 \\ 0.1 \\ 0.7 \end{pmatrix}$$

Note that finding this type of solution without an LP becomes difficult when η_t is very small so that $p_t(q^i = q_1)$ becomes the largest value.

Appendix B

Zusammenfassung

Reputationsmechanismen bieten eine effiziente Methode zur Verhinderung von Marktversagen auf elektronischen Marktplätzen. Durch die Veröffentlichung von Erfahrugen vorhergehender Kunden mit Produkten oder anderen Marktteilnehmern, ermöglichen sie besser informierte Kaufentscheidungen und damit effizientere Marktgleichgewichte. Der Einfluss und damit ökonomische Wert dieser veröffentlichten Erfahrungen wirft jedoch Fragen bezüglich der Glaubwürdigkeit der Mechanismen selbst auf. Existierende Systeme nehmen an, dass die privaten Erfahrungen der Agenten ehrlich abgegeben werden. Diese Annahme steht im Gegensatz zu Erkenntnissen aus der Spieltheorie sowie empirischen Untersuchungen, die nahelegen, dass das von den Agenten abgegebene Feedback unehrlich beziehungweise verzerrt ist.

Miller, Resnick und Zeckhauser [2005] konnten zeigen, dass es für Reputationsmechanismen, welche ausschließlich negative Auslese bekämpfen (wie beispielsweise die Produktbewertungen bei Amazon), möglich ist, einen Feedback-Mechanismus zu entwerfen, der rationale Agenten dazu bringt ehrliche Bewertungen abzugeben. Ihre Methode basiert auf dem Vergleich der Bewertungen zweier Agenten und einer Bezahlung, die davon abhängt wie gut eine Bewertung die Bewertung des anderen Agent "stochastisch erklärt".

In bezug auf die spieltheoretische Komplexität sind Märkte mit ausschließlich negativer Auslese verhältnismäßig einfach, da der Verkäufer in diesem Fall als Spieler mit nur einer Aktion modelliert werden kann. In dieser Arbeit haben wir untersucht, ob sich ein ähnlicher Mechanismus für solche Reputationsmechanismen entwerfen lässt, die in Märkten mit moralischer Versuchung ("moral hazard") eingesetzt werden. Diese Art von Marktversagen tritt auf wenn der Verkäufer ein "echter" Spieler ist, also strategische Entscheidungen trifft, wobei es einen Interessenskonflikt zwischen dem Verkäufer und dem Kunden gibt. Durch ein von uns auf Basis der Situation bei eBay aufgestelltes Modell mit ausschließlich moralischer Versuchung, konnten wir zeigen, dass es keinen einzig auf den Vergleich von Aussagen basierenden Feedback-Mechanismus gibt. Für ein Marktmodell mit sowohl negativer Auslese als auch moralischer Versuchung erzielen wir hingegen ein positives Resultat und entwickeln ein auf Linearer Programmierung basierendes “Feedback Plug-in”, das in bestehende Reputationsmechanismen eingebunden werden kann. Darüber hin aus zeigen unsere experimentellen Ergebnisse, dass dieses Plug-in auch bezüglich Laufzeit und erwartetem Budget realisierbar ist.

Bibliography

[Akerlof, 1970] George A. Akerlof. The Market for 'Lemons': Quality Uncertainty and the Market Mechanism. *The Quarterly Journal of Economics*, 84(3):488–500, 1970.

[Bolton and Ockenfels, 2006] Gary Bolton and Axel Ockenfels. The Limits of Trust in Economic Transactions – Investigations of Perfect Reputation Systems. Working Paper 33, University of Cologne, 2006.

[Bolton *et al.*, 2009] Gary Bolton, Ben Greiner, and Axel Ockenfels. Engineering Trust – Reciprocity in the Production of Reputation Information. Working Paper 42, University of Cologne, 2009.

[Conitzer and Sandholm, 2002] Vincent Conitzer and Tuomas Sandholm. Complexity of Mechanism Design. In *Proceedings of the 18th Conference on Uncertainty in Artificial Intelligence (UAI 02)*, 2002.

[Cooke, 1991] Roger M. Cooke. *Experts in Uncertainty.* Oxford University Press, 1991.

[Cripps *et al.*, 2004] Martin W. Cripps, George J. Mailath, and Larry Samuelson. Imperfect Monitoring and Impermanent Reputations. *Econometrica*, 72(2):407–432, 2004.

[Dellarocas, 2003] Chrysanthos Dellarocas. Efficiency through Feedback-contingent Fees and Rewards in Auction Marketplaces with Adverse Selection and Moral Hazard. In *Proceedings of the 4th ACM Conference on Electronic Commerce (EC 03)*, pages 11–18, 2003.

[Dellarocas, 2005] Chrysanthos Dellarocas. Reputation Mechanism Design in Online Trading Environments with Pure Moral Hazard. *Information Systems Research*, 16(2):209–230, 2005.

[Dellarocas, 2006] Chrysanthos Dellarocas. Reputation Mechanisms. In Terry Hendershott, editor, *Handbook on Information Systems and Economics.* Elsevier Publishing, 2006.

[Fudenberg and Tirole, 1991] Drew Fudenberg and Jean Tirole. *Game Theory*. MIT Press, 1991.

[Jain and Parkes, 2008] Shaili Jain and David C. Parkes. A Game-Theoretic Analysis of Games with a Purpose. In *Proceedings of the 4th International Workshop on Internet and Network Economics (WINE 08)*, 2008.

[Jurca and Faltings, 2005] Radu Jurca and Boi Faltings. Enforcing Truthful Strategies in Incentive Compatible Reputation Mechanisms. In *Proceedings of the 1st International Workshop on Internet and Network Economics (WINE 05)*, 2005.

[Jurca and Faltings, 2006] Radu Jurca and Boi Faltings. Minimum Payments that Reward Honest Reputation Feedback. In *Proceedings of the 7th ACM Conference on Electronic Commerce (EC 06)*, pages 190–199, 2006.

[Jurca and Faltings, 2007a] Radu Jurca and Boi Faltings. Collusion-resistant, Incentive-compatible Feedback Payments. In *Proceedings of the 8th ACM Conference on Electronic Commerce (EC 07)*, pages 200–209, 2007.

[Jurca and Faltings, 2007b] Radu Jurca and Boi Faltings. Obtaining Reliable Feedback for Sanctioning Reputation Mechanisms. *Journal of Artificial Intelligence Research (JAIR)*, 29:391–419, 2007.

[Jurca and Faltings, 2007c] Radu Jurca and Boi Faltings. Robust Incentive-Compatible Feedback Payments. In *Trust, Reputation and Security: Theories and Practice*, volume 4452 of *LNAI*, pages 204–218. Springer-Verlag, 2007.

[Kotowitz, 2008] Yehuda Kotowitz. Moral Hazard. In Steven N. Durlauf and Lawrence E. Blume, editors, *The New Palgrave Dictionary of Economics*. Palgrave Macmillan, second edition, 2008.

[Kreps and Wilson, 1982a] David M. Kreps and Robert Wilson. Reputation and Imperfect Information. *Journal of Economic Theory*, 27(2):253–279, 1982.

[Kreps and Wilson, 1982b] David M. Kreps and Robert Wilson. Sequential Equilibria. *Econometrica*, 50(4):863–894, 1982.

[Laffont and Martimort, 2001] Jean-Jacques Laffont and David Martimort. *The Theory of Incentives*. Princeton University Press, 2001.

[Milgrom and Roberts, 1982] Paul Milgrom and John Roberts. Predation, Reputation and Entry Deterrence. *Journal of Economic Theory*, 27(2):280–312, 1982.

[Miller *et al.*, 2005] Nolan Miller, Paul Resnick, and Richard Zeckhauser. Eliciting Informative Feedback: The Peer-Prediction Method. *Management Science*, 51(9):1359–1373, 2005.

[Miller *et al.*, 2006] Nolan H. Miller, John W. Pratt, Richard J. Zeckhauser, and Scott Johnson. Mechanism design with multidimensional, continuous types and interdependent valuations. *Journal of Economic Theory*, 136(1):476–496, 2006.

[Nisan, 2007] Noam Nisan. Introduction to Mechanism Design (for Computer Scientists). In Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani, editors, *Algorithmic Game Theory*. Cambridge University Press, 2007.

[Osborne and Rubinstein, 1994] Martin J. Osborne and Ariel Rubinstein. *A Course in Game Theory*. MIT Press, seventh edition, 1994.

[Parkes, 2001] David Parkes. *Iterative Combinatorial Auctions: Achieving Economic and Computational Efficiency*. PhD thesis, Department of Computer and Information Science, University of Pennsylvania, 2001.

[Pennock and Sami, 2007] David Pennock and Rahul Sami. Computational Aspects of Prediction Markets. In Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani, editors, *Algorithmic Game Theory*. Cambridge University Press, 2007.

[Peters, 2008] Hans Peters. *Game Theory – A Multi-Leveled Approach*. Springer-Verlag, 2008.

[Resnick and Zeckhauser, 2002] Paul Resnick and Richard Zeckhauser. Trust Among Strangers in Internet Transactions: Empirical Analysis of eBay's Reputation System. In *The Economics of the Internet and E-Commerce*. Emerald Group Publishing, 2002.

[Sandholm, 2003] Tuomas Sandholm. Automated Mechanism Design: A New Application Area for Search Algorithms. In *Proceedings of the 9th International Conference on Principles and Practice of Constraint Programming (CP 03)*, 2003.

[Savage, 1971] Leonard J. Savage. Elicitation of Personal Probabilities and Expectations. *Journal of the American Statistical Association*, 66:783–801, 1971.

[Schelling, 1960] Thomas Schelling. *The Strategy of Conflict*. Havard University Press, 1960.

[Schrijver, 1998] Alexander Schrijver. *Theory of Linear and Integer Programming.* John Wiley & Sons, Inc., paperback edition, 1998.

[Shneidman and Parkes, 2004] Jeffrey Shneidman and David C. Parkes. Specification Faithfulness in Networks with Rational Nodes. In *Proceedings of the 23rd ACM Symposium on Principles of Distributed Computing (PODC 04)*, 2004.

[Shoham and Leyton-Brown, 2009] Yoav Shoham and Kevin Leyton-Brown. *Multiagent Systems – Algorithmic, Game-Theoretic, and Logical Foundations.* Cambridge University Press, 2009.

[Varian, 2006] Hal Varian. *Intermediate Microeconomics (International Student Edition)*, volume 7. Harvard University Press, 2006.

[Vickrey, 1961] William Vickrey. Counterspeculation, Auctions, and Competitive Sealed Tenders. *Journal of Finance*, 16(1):8–37, 1961.

[von Ahn and Dabbish, 2004] Luis von Ahn and Laura Dabbish. Labeling Images with a Computer Game. In *Proceedings of the 22nd Conference on Human factors in Computing Systems (CHI 04)*, pages 319–326, 2004.

[von Ahn and Dabbish, 2008] Luis von Ahn and Laura Dabbish. Designing games with a purpose. *Communications of the ACM*, 51(8):58–67, 2008.

[Weber *et al.*, 2008] Ingmar Weber, Stephen Robertson, and Milan Vojnović. Rethinking the ESP Game. Technical report, Microsoft Research, 2008.

[Winkler *et al.*, 1996] Robert Winkler, Javier Muñoz, José Cervera, José Bernardo, Gail Blattenberger, Joseph Kadane, Dennis Lindley, Allan Murphy, Robert Oliver, and David Ríos-Insua. Scoring Rules and the Evaluation of Probabilities. *TEST: An Official Journal of the Spanish Society of Statistics and Operations Research*, 5(1):1–60, 1996.

[Witkowski, 2008] Jens Witkowski. Eliciting honest reputation feedback in a Markov setting. Studienarbeit, Albert-Ludwigs-Universität Freiburg, August 2008.

[Witkowski, 2009] Jens Witkowski. Eliciting Honest Reputation Feedback in a Markov Setting. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI 09)*, 2009.

[Wolfers and Zitzewitz, 2004] Justin Wolfers and Eric Zitzewitz. Prediction Markets. *Journal of Economic Perspectives*, 18(2):107–126, 2004.