The Broker Who Would Not Bundle

Kurt DeMaagd and Cliff Lampe
Michigan State University
(kdemaagd)(lampecli)@msu.edu

Abstract

Online reputation systems are a challenge to design. They often require large user-bases to frequently and honestly rate content and users. Although this is a daunting problem, reputation systems can help to alleviate the “cheap pseudonym” problem. One way to make it easier to create and sustain a reputation system is to share information about the identity of users between web sites. We report on some of the obstacles to such a user identity marketplace.

In particular, we find that the brokers needed to would serve as intermediaries in the marketplace have incentives that do not align with the firms that want to trade identity information. More specifically, the individual firms would benefit if the brokers created bundles, but these bundles would reduce the broker’s profitability. Therefore, the brokers will not bundle the information goods.

1 Introduction

The ability of users to easily change their identities in online interactions, the "cheap pseudonym" problem [7], has been tied to adverse effects ranging from degraded online discussions to duplicitous e-commerce transactions. One remedy to the cheap pseudonym problem has been the use of reputation systems, which use information systems to collect and disseminate evaluations of users in online interactions for the use of future, potential participants. Being able to share pseudonym reputations across different online sites would increase the efficacy of reputation systems.

Currently, separate sites need to develop individual reputation systems, making it costly for users to accrue reputation in multiple systems. Also, having reputation systems individually tailored to sites makes it possible for users to more easily abandon pseudonyms that have accrued negative reputations. Creating mechanisms by which reputation attached to pseudonymous identities can be shared across different systems will be more efficient for sites as they attract users, and will reduce the cost for users as they build a history of positive interactions. This paper reports on one of the design problems with such pseudonym sharing mechanisms.

This challenge of developing such a mechanism stands in marked contrast to a large portion of the current studies on information sharing online. Privacy and security have been major points of interest in the recent literature as identity theft is a growing problem and spam continues to clog inboxes. The privacy literature focuses on ensuring that firms do not share information that the consumer wants to keep private. Yet there is a converse problem where firms do not share the information that an individual may want public. Pseudonym sharing often falls into this category.

For example, consider eBay's reputation system. Studies show that a higher eBay reputation translates into higher auction revenues [14]. Yet information about a given user's reputation is controlled by eBay. If a user becomes frustrated with eBay's policies and switches to Yahoo Auctions, she must abandon her previously established positive identity on eBay. She loses the potential revenue benefits of a strong positive identity until she can reestablish a reputation. This poses a switching cost for the user. Furthermore, because an established identity reduces the chances of uncooperative behaviors [7], Yahoo Auctions would benefit if the user's identity could be transferred. Yahoo would also benefit if users could carry their identity between web sites.

From a business perspective, eBay's policy is logical. Their control of the user's information is a form of competitive advantage. Although it may be socially advantageous for eBay to share this information, they have no individual incentive to share their information. In other words, eBay lacks the incentives to act in the general social interest. One solution is to create an identity marketplace where eBay is financially compensated for sharing their information. This paper studies the incentives and obstacles to creating such a system and contrasts the results with several classical assumptions of information economics.

Note that the problem with identity sharing is not addressed by systems such as the OpenID project [13] and other single sign-on systems. Single sign-on systems grant a user a single single sign-on account for use in many sites. They are currently limited to only authentication services. They do not include the additional information about a user necessary to establish a reputation and to reap the benefits of socially cooperative behavior on many sites.

In this paper, we make two contributions to the literature. First, we apply the previous work on mechanism design to the cheap-psuedonym problem. Second, we show how this specific problem presents a challenge to one of the common assumptions of information economics. In particular, while bundling can help to reduce uncertainty over the value of information goods, we show that the agents involved in a transaction may have incentives not to bundle or at least to minimize the size of their bundles.

Before covering the specific problem of bundling, we first review the role of reputation systems and why they are such desirable tools for the Internet. Having shown that reputation online is important, we then discuss some of the constraints and design features of a market for sharing identity information. Finally, we discuss why agents may not choose to create bundles of information goods even though some prior work predicts that they should.

2 Reputation

Before we discuss the user identity marketplace, we first examine the motivation for such a tool. As discussed below, reputation system are an important part of many successful online communities, but there are large barriers to adoption. In the following literature review, we discuss some of the barriers to implementing a successful online reputation system and then we identify ways that an identity marketplace could resolve some of these issues.

2.1 Barriers in recommendation and reputation

Reputation and recommender systems are inherently difficult to design. For them to function properly, they usually require a large user base that actively and honestly provides feedback [2, 8]. As a general rule, the more ratings, the more accurate the product recommendations, document filters, user reputations, etc. can be. In addition, large user bases are usually necessary to ensure that minority opinions are properly considered. Without a sufficiently large user base over which to aggregate the minority views, the minority ratings get de-valued as random error.

Besides needing a requisite number of users, recommender systems depend on large amounts of feedback to operate. In the MovieLens case, unrated movies cannot be recommended as they cannot be placed into a user profile of similar ratings. Assuming that a group is large enough to have the capacity to provide feedback, theories of social loafing, as described above, predict that few moderators would choose to spend their points. In Lampe and Resnick [9], the authors found that only 26% of comments on Slashdot.org received any sort of moderation, however exact numbers of moderators who expend their points still need to be determined.

Although it may be difficult for any single web site generate the traffic necessary to sustain a reputation or recommender system, it may be possible to resolve this problem by aggregating data across many web sites. This can provide the necessary critical mass even in smaller communities that lack the necessary population to use a reputation or recommender system.

2.2 Information cascades and herding

Sometimes, instead of making an evaluation based on the user's own preferences and private information, the user makes a choice based on the choices already made by prior users. The canonical example is making a choice of restaurants to patronize based on the presence of other diners. If the restaurant is busy, the person may choose to eat there purely because others are. If the restaurant is empty, then the person may move on to another location.

"An informational cascade occurs when it is optimal for an individual, having observed the actions of those ahead of him, to follow the behavior of the preceding individual without regard to his own information" [4]. In other words, information cascades occur when the a person makes their decision based not on their own preferences and information but on observations of the choices of others. Although this may be a convenient and quick method for decision making, it can have several consequences. Within an online context, system designers often assume that moderators make independent decisions. Therefore, aggregating those decisions should create an optimized rating. But, if the ratings are part of a cascade, then the first rater can create false positives and false negatives which are reinforced instead of corrected by subsequent raters. Fortunately, research has found that cascades are fragile. In the event that more careful decision makers also participate in the rating, the cascade can be interrupted [4].

Information cascades can lead to "herding" behavior. Herd behavior is "everyone doing what everyone else is doing, even when their private information suggests doing something quite different" [3]. This has several implications for designing recommendation systems. Banerjee finds that the probability that nobody in the population reaches the right decision is bounded away from zero, in contrast to the situation where decision makers are acting independently and someone is likely to find the correct decision [3]. One of the conclusions of Banerjee is that it might be better if early decision makers are not allowed to observe the choices made by other decision makers. That would mean that each decision maker is acting on their own information, rather than reacting to previous decisions.

Most of this previous work has looked at the role of cascades and herding on just one site. One potential solution to the problem would be to aggregate information from many sites. For example, a cascade on one site that claims that a bad user is good could conflict with information from another web site that correctly identified that the user should have a negative reputation. By combining the information from the two sites together, it may provide the independent information necessary to halt the information cascade.

2.3 Cheap Pseudonyms

Recall from the introduction that this paper was initially motivated by Friedman and Resnicks paper on the cheap pseudonym problem. This problem predictings how new users would be treated in an online environment where negative reputations might encourage users to change pseudonyms often [7]. Dellarocas builds on Friedman and Resnick's paper showing that setting newcomers reputations as low as possible increases the efficiency of determining user reputations [6]. This analysis was limited to cases where there is binary (up/down) feedback and where mediators only publish the sum of recent ratings.

The implications of these findings is that in many environments, where pseudonyms that accrue negative reputations are easily dropped, the user base will be naturally suspicious of newcomers. One outcome of this suspicion might be lower rating of newcomer contributions, which can be derived from data logs, or higher overall levels of distrust in fellow members, which can be asked about in survey questions.

2.4 Honest Feedback

In many cases, it is difficult to elicit honest feedback. There are many reasons why users may not provide honest ratings of their experiences [6, 11]. Users may have a conflict of interest, in that they are rating someone who is a friend or a foe, or may even be rating themselves using multiple accounts. They may not have actually experienced the item being rated, but have some incentive to rate the item regardless. The users may receive rewards for doing ratings, but nothing guarantees they actually read the content being moderated or they may not have the correct incentive to take the time necessary to create a proper rating. Raters may simply prefer to be nice and not rate content or other users negatively.

As a result, sites must balance between two goals. On the one hand, they want a large quantity of rating and reputation information. On the other hand, the site doesn't want to push quantity so hard that it sacrifices quality. To balance these two goals, it may be possible for each individual site to create mechanisms that emphasize the quality of the reputation and rating system. Then to fill in for the lack of quantity, the different sites could share the relevant reputation and rating information.

2.5 Literature summary

There are four main points that this review illustrates. First, reputation systems require a critical mass to be successful. Not all sites can generate this mass. Although the site may lack this critical mass, it may be possible to purchase the necessary information for the reputation system to operate correctly. Second, cascades and herding can threaten the accuracy of a reputation or recommender system. If, however, a firm can obtain information from a new and independent source, it may be possible to interrupt the cascade. Third, online identities are easy to dump if they get a negative reputation. If, however, it is more costly to dump an identity than to create a good reputation, then the users have an incentive to pursue cooperative behavior online. One way to increase the cost of dumping an online identity is to make it span across many web sites. Finally, to help elicit honest feedback, individual sites can focus on quality over quantity, compensating for the lack of quantity by sharing the rating information across many sites.

3 Constraints

We now turn our attention to designing the identity exchange market. A market based system for pseudonym sharing creates several substantial challenges. As a market based system, we can not rely on the all-knowing, all-powerful social planner to optimally allocate the information among the different actors. Instead, the mechanism must create the proper incentives such that a series of self-interested agents will find it individually rational to act in a way that also happens to be socially optimal.

The first constraint, as briefly alluded to above, is the participation constraint. The firms and individuals must find that it is in their own best interest to participate in a transaction that potentially transfers pseudonym information from one entity to another. Given that there is no all-powerful social planner to compel the parties to trade, if an individual agent would be worse off as a result of a trade, then the agent will simply choose not to trade. This is true even if it might be socially optimal for the parties to trade. Each party must be better off as the result of the trade or no trade will occur.

This idea is closely tied to the concept of Pareto efficiency—also sometimes referred to as Pareto optimality—in welfare economics [16]. Pareto efficiency represents an allocation of goods such that there is no way to make at least one individual better off without making another worse off. Pareto efficiency is often criticized because it can support a status quo with an inequitable distribution of wealth and there are a limited number of mechanisms for overcoming this distribution of wealth [15]. Two methods are commonly proposed for resolving this problem. The first is for a disproportionately strong agent to force a reallocation. The second is to develop a system of transfer payments to compensate the individual who would otherwise be worse off.

Likewise, the allocation of information across the current web sites controlling pseudonyms and reputation is Pareto optimal. The sites could lose their competitive advantage if they shared this information. We assumed away the existence of a powerful third party, such as a government entity, to compel the web sites to share their private information. This leaves the other alternative of arranging transfer payments between the different parties to create the necessary incentives to trade.

This leads to the next constraint. We assume that the agents will act strategically when attempting to negotiate the transfer payments. The buyer will want to pay as little as possible and the seller will want to charge as much as possible. As a result, the firms may have an incentive to lie about their true value of the information in an effort to maximize their individual profits. Yet unless we have accurate information about how much the different firms value the information goods, it is impossible to achieve an efficient allocation of the goods.

Take, for example, a seemingly equitable allocation mechanism. Let the buyer announce her value for the information and the seller announce his value for the information. Then have the buyer pay an amount half way between the two valuations. For example, if the buyer values the information at 20 and the seller values it at 10, then the buyer should pay 15. Yet such a mechanism would be easily manipulated. For example, the seller could announce that his value is 14, thereby driving the selling price up to 17. Hence, the seller can lie and profit.

Granted, both agents are still better off in the above scenario. As long as the seller believes that there is a sufficiently high probability that the buyer values the goods at more that 14, it is rational for him to lie. But then the seller runs the risk that the buyer has a value higher than his true value but lower than his quoted price. As a result, it is possible that it is in the buyer's self interest to lie even though it may be socially disadvantageous to do so. The problems of an algorithm like the one described above are discussed in more detail by Chatterjee and Samuelson [5].

Recall that the goal is to align the buyer's individual actions with the social optimum. Hence, the mechanism must ensure that the buyers and sellers have the incentive to truthfully report their valuations. This is often referred to as incentive compatibility. Whatever mechanism determines the transfer prices must ensure that the individual agents do not have an incentive to lie. Although this is often a difficult goal, mechanisms such as the second-price sealed-bid auction are examples where the individual agents have the incentive to tell the truth [10].

To reflect real market circumstances, we assume that the individuals and firms do not know how much the other individuals and firms value the information. They have beliefs of the range of values, but they lack exact information. Because it is rare for buyers and sellers to to have perfect information about other parties' valuations, this assumption helps to align the model with actual market dynamics. In the next section, we mathematically define how such a market would work, the obstacles to implementation, and how the market may behave in a fashion that is contrary to common assumptions from information economics.

4 Identity Market

Given the constraints to an identity market identified in the previous section, we now describe a market to trade reputation and recommender system information between web sites. Although there are social advantages to identity sharing, individual firms often lack the incentive to share this information. Because the information represents a source of competitive advantage, the firm will not share this information unless it is duly compensated. This section details a marketplace for the exchange of identities such that firms are compensated for the loss of their competitive advantage. We begin by assuming a bilateral trade of information. Later, we extend the model to selling bundles across many firms.

Assume that the value of information comes from two sources: conventional and strategic. The conventional value is the firm's ability to use information to improve its business processes. In the context of this paper, an established reputation creates an incentive to cooperate. Because cooperation—or at least non-malicious behavior—is fundamental to the efficient operation of many markets, the firm benefits from established user reputations. Sharing this information with another firm in no way reduces its value to the firm.

In contrast, information has an additional strategic value if one firm possesses the information but another does not. For example, if users can not move their reputation to another site, it creates a form of lock-in. This gives the firm a competitive advantage. If the firm were to share this information with other firms, it would lose this source of competitive advantage. This is one of the primary impediments to online identity sharing and is one of the primary determinants of the buyer's price.

Now consider a potential transaction between two firms. The first firm owns information on the identity of a user. The second firm is interested in purchasing this information. Each firm knows its own value of the information, but it does not know the valuation of the other firm. Because the first firm has exclusive ownership, it currently derives both conventional and strategic value from the information. The second firm wants the conventional value of the information and must therefore pay the first firm for the loss of its strategic value.

Expressing this mathematically, the buyer and seller have values of the information equal to v_0 and v_1, respectively. In this case, v_0 represents the buyer's conventional value from the information and v_1 represents the seller's strategic value of the information. Note that we do not model the seller's conventional value. Because the seller will not surrender access to the information after the trade, it maintains the conventional value. As such, the buyer will not need to compensate the seller for this value. The firms do not know each other's valuations, but they have beliefs regarding the distribution of those values. Let $f_0()$ be the probability distribution of firm buyer's value. Let $f_1()$ be the probability distribution of the seller's value. Let F_0 and F_1 be the corresponding densities.

For the sake of simplicity in the explanations, for the remainder of this paper we will assume that the random variables are distributed according to a uniform distribution. The results, however, also hold under other probability distributions. This means that the buyer's values are distributed uniformly between a_0 and b_0. Likewise, the seller's values are distributed between a_1 and b_1.

The next step is for each firm to announce their value for the good so it is possible to determine the price for the exchange. Recall from above that the mechanism needs to ensure that the firms have the incentive to truthfully report their value and that it should be individually rational for them to announce this price. Yet, according to prior research, no appropriate mechanism can exist [12]. If the agents have imperfect

Figure 1: An example of an "impossible" transaction. Notice the overlap of the buyer and seller beliefs.

information about the other agents, if we want them to truthfully reveal all of their information, and if they can not be compelled to participate in the transaction, then there is no mechanism to ensure that all socially optimal transactions will take place. This is the negative result of the Mayerson-Satterwite theorem.

To see why from a more intuitive level, consider Figure 1. The left bar represents the distribution of the seller's valuations. The right bar represents the distribution of the buyer's valuations. Notice the overlapping region. If there was no overlap, then the uncertainty would not create a problem. With no overlap, it is still possible that the buyer and seller could agree upon a price such that they are both guaranteed to benefit from the transaction.

The problem arises because there is an overlap. It is no longer obvious as to whether the two parties can agree upon a price. It is possible that the buyer's value is really less than the sellers so that it is not socially optimal to trade. Given that they two parties may not profit, it may not be individually rational for them to participate in the transaction. For a more complete discussion, including a more mathematical treatment of the problem, see [10].

Although it is impossible to guarantee that all socially optimal transactions will occur, subsequent research has focused on using two techniques so that many socially optimal transaction occur. The first option is to use a broker to mediate between the parties. The second is to use bundling to reduce the uncertainty, potentially eliminating—or at least reducing—the overlap between the beliefs of valuations.

Begin by considering the role of brokers. This is the solution originally proposed by Myerson and Satterthwaite [12]. In such a market, we assume the existence of the two agents described above plus a broker to negotiate the price between the two. Both the buyer and seller report their valuations of the identity information to the broker. The broker then uses these values to compute the buying and selling price, as follows. Note that with the Myerson-Satterthwaite broker, the broker chooses the prices to maximize its own profits and not the social welfare.

Amount given to the seller:

$$p_1 = v_0 - \frac{b_0 - a_1}{2} \tag{1}$$

Amount charged to the buyer:

$$p_0 = v_1 + \frac{b_0 - a_1}{2} \tag{2}$$

The broker profits by keeping the difference between the p_0 and p_1. The broker is not guaranteed to profit on any individual transactions. Instead, the broker agrees to participate in the transaction regardless of whether it will profit on the individual transaction. Nonetheless, prices are determined such that the broker will maximize its profits even if it has to subsidize some transactions.

Interestingly, note that the sellers price is not a function of the seller's value and the buyer's price is not a function of the buyer's value. This is one of the common tricks used to ensure incentive compatibility. Although the values reported by the agents is used to compute the price for other agents and to determine whether it would be socially optimal to trade, the actual price quoted to an agent is not a function of the information that agent reports. As a result, the agent will not get a better price by lying.

Myerson and Satterthwaite's system satisfies the requirements of individual rationality and incentive compatibility [12]. Yet it comes at a cost. Some socially optimal transactions are not closed. This left a substantial hole for additional tools to resolve this problem.

To see how this works in action, consider the simplified case where both probability distributions are uniform over the interval $[0, 1]$. To express that in variables for the above equations, this means that $a_0 = a_1 = 0$ and $b_0 = b_1 = 1$. Plugging the values into the equations, the broker charges the buyer $v_1 + \frac{1}{2}$ and then pays $v_0 - \frac{1}{2}$ to the seller. For example, consider the case where the buyer values the good at $v_0 = 0.8$ and the seller values the good at $v_1 = 0.2$. The broker charges the buyer $0.2 + 0.5 = 0.7$. The broker then pays the seller $0.8 - 0.5 = 0.3$. The broker profits in the amount of $0.7 - 0.3 = 0.4$.

Notice in the above example that a large range of socially optimal transactions will not occur. For example, if the seller instead valued the good at 0.31, the broker would try to charge the buyer 0.81, and the buyer would walk away from the transaction. (Recall that the buyer's value is only 0.8.) It is possible for many such socially optimal trades not to be transacted.

More recently, many researchers have been interested in the role of bundling information goods to reduce uncertainty. Uncertainty is a great impediment to trade. Because the firms do not know the other party's valuation, they can not determine an individual rational and incentive compatible strategy. If, however, the firms could reduce their uncertainty over the values of the information goods, then it may be easier to conduct a trade. One way to reduce this uncertainty is to create a bundle of many information goods [1].

In the case of firm's sharing their reputation information, it is hard for any one firm to create a bundle. Nonetheless, the broker may be able to create a bundle of the information from many different web sites. This requires a slight modification of the above model. Instead of having just two firms, the new model includes one buyer and N different firms with information to sell. Let f_n represent the probability distribution of the nth firm's value. Likewise, let F_n be the density function. Assume that firm 0 is the buying firm and that firms 1 to N are the sellers.

With this set of firms it is now possible to create a bundle of their information goods. Because all of the sellers represent independent firms, it is not possible for any single firm to create the information bundle. Yet, a broker could create a bundle for the buyer. As a result, instead of the seller looking at N separate and independently sold information goods with corresponding, the seller now has the following belief of the value of the information bundle:

$$f_{\text{bundle}} = \sum_{i=1}^{N} f_i \tag{3}$$

As has been shown in the earlier literature, variance of this distribution decreases as the size of the bundle increases. To reduce confusion in the explanation, we assume that f_{bundle} can still be approximately by a uniform distribution. This is a reasonable assumption for small enough N. Of course, as N increases, the distribution begins to converge to a normal distribution. Even under a normal distribution, the following results will still hold.

The actual pricing equations would work as before. The broker would pay the individual sellers as described before. The buyer's equation would change slightly, with the bundled distribution replacing the individual distributions.

The broker charges the buyer:

$$p_0 = v_0 + \frac{b_0 - a_{\text{bundle}}}{2} \tag{4}$$

When uncertainty—as defined by variance in the probability distribution—decreases, notice what happens to the above equation. If uncertainty decreases, that is mathematically represented by an increase in a_{bundle}. As a_{bundle} increases, the numerator of the fraction decreases. As the numerator to the fraction decreases, the whole equation decreases. To say that more mathematically:

$$\frac{\delta p_0}{\delta a_{\text{bundle}}} = -\frac{1}{2} \tag{5}$$

This means that the amount charged to the buyer decreases as the uncertainty decreases. This is certainly good for the buyer, but what about the broker. Recall that the broker profits from the difference between the amount charged to the buyer and the amount paid to the seller. If the amount charged to the buyer decreases as uncertainty decreases, the broker's profits also declines. Hence, a broker acting in its own self-interest would choose not to bundle.

5 Conclusion

We began the paper with a problem: how to design a mechanism such that web sites can exchange user identity information. This includes both information about the user's reputation and various other data necessary to run a recommender system. Although it is in the user's interest for this information to have some portability, the web sites themselves gain a competitive advantage from keeping the information a secret. The sites will only trade information if they are financially induced.

Next, we considered some of the constraints when designing a marketplace for exchanging identity information. It should be individually rational for the participants to trade, and they should have incentives to truthfully report their values for the information goods. Finally, there is uncertainty over how much the other party values the goods.

In such a scenario, the prior literature shows that there is no mechanism to meet our desired goals. Instead, various techniques are used to try to approximate the goals, recognizing that some socially optimal exchanges still will not happen. The two most common approaches are using brokers and bundling. Although the broker can create the necessary incentives for trade, they still tend to loose a fairly large amount of socially optimal trades. Fortunately, when dealing with information goods, other research has shown that bundling the information goods can help reduce uncertainty.

But what happens when this is applied to the problem of an identity marketplace? No single firm controls all of the sites necessary to create a bundle of the information. As a result, the marketplace must rely on a broker to serve as the intermediary. The broker can bundle the goods from many different sites and then sell that bundle to a buyer. This bundle improves the price offered to the buyer, creating a greater incentive to trade.

Although this looks like a positive outcome at first glance, would the broker have the incentive to create the bundle? The more uncertainty there is over the value of the goods, the more the broker profits. Reducing this uncertainty reduces the broker's profits. The broker looses profits when it bundles the goods. Hence, the broker will not bundle.

Extending this beyond the specific case of an identity marketplace, it has further implications for the plausibility of bundling information goods. The results to not challenge the case where a single firm can bundle their own information goods for sale. For example, it is not a challenge to the canonical example of Microsoft bundling Word and Excel to maximize profits. Yet it does present a challenge to any system that may rely on a broker to bundle multiple information goods from several independent firms. In this case, the broker profits from uncertainty. Bundling reduces that uncertainty. Hence, an independently operating broker will choose not to create a bundle.

These results create opportunities for future research. In the above discussion, we have only shown that the result holds when the values are uniformly distributed. Although it is intuitively appealing to assume that the result holds for all probability distributions, this result should be more rigorously established.

Second, the model ignores the externalities of the information sharing. If one firm shares its information with another, that helps to more fully establish the user's online identity, which creates value for the firm. The firm will loose the strategic value of the information, but it will make it more costly for the user to switch pseudonyms. This creates value for the firm. Hence only looking at the strategic costs fails to capture some of the benefits of the exchange.

Which leads to the third extension. The broker can also serve as a demand aggregator. As described above, the more firms that posses the identity information, the more it is worth. Although it may be unprofitable to the broker to create bundles of information goods to sell, it may profitable for the broker establish a large group of buyers to increase the value of the goods and therefore the price it can charge. Hence, although the broker may not have the incentive to bundle to reduce uncertainty, it may still play an important role in an identity exchange market.

References

[1] Yannis Bakos and Erik Brynjolfsson. Bundling information goods: Pricing, profits, and efficiency. *Management Science*, 45(12):1613–1630, 1999.

[2] Balabanovic and Shoham. Content-based, collaborative recommendation. *Communications of the ACM*, 40(3):66–72, 1997.

[3] A. Banerjee. A simple model of herd behavior. *Quarterly Journal of Economics*, 107(3):797–818, 1992.

[4] Bikhchandani, Hirshleifer, and Welch. A theory of fads, fashion, custom and cultural change as information cascades. *Journal of Political Economy*, 100(5):992–1026, 1989.

[5] Kaylan Chatterjee and William Samuelson. Bargaining under incomplete information. *Operations Research*, 31(5):835–851, 1983.

[6] Chris Dellarocas. The digitization of word of mouth: Promise and challenges of online feedback mechanisms. *Management Science*, 49(10):1407–1424, 2003.

[7] Erik Friedman and Paul Resnick. The social cost of cheap pseudonyms. In *Telecommunications Policy and Research Conference*, 1999.

[8] N. Good, J. B. Schafer, J. A. Konstan, A. Borchers, B. M. Sarwar, J. L. Herlocker, and J. Riedl. Combining collaborative filtering with personal agents for better recommendations. In *Sixteenth National Conference on Artificial Intelligence (AAAI)*, 1999.

[9] Cliff Lampe and Paul Resnick. Slash(dot) and burn: distributed moderation in a large online conversation space. In *Conference on Human Factors in Computing Systems (CHI), Vienna, Austria*, 2004.

[10] Andreu Mas-Colell, Michael Winston, and Jerry Green. *Microeconomic Theory*. Oxford University Press, 1995.

[11] Nolan Miller, Paul Resnick, and Richard Zeckhauser. Eliciting hoest feedback: The peer prediciton method. *Management Science*, 51(9):1359–1373, 2005.

[12] Roger Myerson and Mark Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29(2):265–281, 1983.

[13] OpenID Project. OpenID: an actually distributed identity system. http://www.openid.net/, 2007.

[14] Paul Resnick, Richard Zeckhauser, John Swanson, and Kate Lockwood. The value of reputation on ebay: A controlled experiement. *Experimental Economics*, 9(2):79–101, 2006.

[15] A.K. Sen. *Collective Choice and Social Welfare*. Holden Day, San Francisco, 1970.

[16] Hal Varian. *Microeconomic Analysis*. W. W. Norton and Company, New York, 1992.