How much trust is enough to trust? A market-adaptive trust threshold setting for e-marketplaces

Zeinab Noorian*, **Mohsen Mohkami*** and **Julita Vassileva** [1]

Abstract.
The inherent uncertainties of open marketplaces motivate the design of reputation systems to facilitate buyers in finding honest feedback from other buyers (advisers). Defining the threshold for an acceptable level of honesty of advisers is very important, since inappropriately set thresholds would filter away possibly good advice, or the opposite - allow malicious buyers to badmouth good services. However, currently, there is no systematic approach for setting the honesty threshold. We propose a self-adaptive honesty threshold management mechanism based on PID feedback controller. Experimental results show that adaptively tuning the honesty threshold to the market performance enables honest buyers to obtain higher quality of services in comparison with static threshold values defined by intuition and used in previous work.

1 Introduction

Open electronic marketplaces consist of autonomous and self-interested participants (buyers and sellers). They provide services or goods of varying quality and can misrepresent their offerings to maximize their utility. Thus, to minimize the risk of engaging with low-quality or malicious partners, the participants need to communicate with each other and share their experiences with other market participants, thus providing advice or recommendations to each other. Yet, exactly because the agents are autonomous, intelligent and self-interested, they are not obliged to tell the truth. Some may be malicious and badmouth good participants (e.g. to reduce the competition for their own services or goods). Others may be incompetent to give a fair evaluation. Some others may simply not want to be bothered providing feedback.

Designing reputation systems for open marketplaces seems to be an effective approach to ensure that only participants with satisfactory qualities can prosper [6, 11]. Reputation systems assist buyers in their decision making process by providing them with trustworthiness assessment techniques to thoroughly evaluate the credibility of other buyers (advisers), considering various parameters and environmental circumstances.

Different reputation mechanism have been proposed in the literature which model the trustworthiness of participants via different approaches such as socio-cognitive [16], game theoretical [18], and probabilistic models [4, 15, 19]. Existing reputation systems perform under the assumption of the existence of a credibility threshold, which sets a decision boundary on the behavioral model of advisers and characterized them as honest and malicious. These systems suffer from a lack of a systematic approach for adjusting the honesty threshold to the dynamic environmental conditions. Defining the threshold for acceptable level of honesty of advisers is very important. The foremost drawback of having a static honesty threshold is that an inappropriately set threshold would filter away possibly good advice, or the opposite - allow malicious buyers to badmouth good services. A low threshold will result in a plenty of possible advisers, but the quality of advice may be low. In this situation, deceitful advisers who maintain a minimum level of trustworthiness remain undetected and could actively contribute into a buyer's decision making process. On the other hand, a higher credibility threshold leads to the contribution of a smaller number of advisers and can make it impossible to find advisers. Clearly, adjusting a threshold value is a trade-off between the number of credible advisers and the risk of being misled by deceptive peers.

This paper proposes a method by feedback on the performance of the marketplace in terms of QoS metrics to dynamically determine appropriate value for honesty threshold to optimize the market performance. We built a controller that monitors the quality of e-marketplace and uses a PID feedback controller technique [13] to determine new values for the honesty threshold. Buyers then dynamically re-evaluate their network of trustworthy advisers according to the new recommended value.

Our approach was validated experimentally by integrating our PID-based honesty threshold controller into a simulated e-marketplace with different population tendency. Experimental results show that adaptively tuning the honesty threshold to the market performance enables honest buyers to obtain higher quality of services and more accurately detect malicious advisers in comparison with the static threshold values defined based on designer intuition that are used in previous work.

A credibility evaluation mechanism guided by the PID-based threshold adjustment creates the opportunity of designing self-improving trust and reputation systems which learn from the state of the e-marketplace promoting the acceptance of web-based agent-oriented e-commerce by human users.

2 Credibility Evaluation Mechanism

Our proposed credibility evaluation mechanism adopts a variation of the *Prob-Cog* model [9, 10] and formalizes the credibility degree of advisers in different steps.

In the first step, a buyer agent c sends a query to its neighbours $A = \{a_1, a_2, ..., a_k\}$ requesting information about previously experienced transaction outcomes with common set of sellers $P = \{P_1, P_2, \cdots, P_j\}$. Neighbour a_k responds by providing rating reports for the common set of sellers. Consumer c calculates the differ-

[1] Department of Computer Science, University of Saskatchewan, Canada,
* Authors equally contributed on the paper. email: zen951@mail.usask.ca

ence of trustworthiness evaluation between c and a_k on the common set of sellers (e.g. P_j) as follows:

$$Diff_{(c,a_k)_{P_j}} = E(pr_r, P_j)_c - E(pr_r, P_j)_{a_k} \tag{1}$$

where $E(pr_r, P_j) = \frac{r+1}{r+s+2}$ represents the expected value of the probability of a positive outcome for seller P_j. Noted that r, s indicate the number of successful and unsuccessful interaction outcomes, respectively.

In the second step, buyer c employs a measure indicated in Equation 2 to further adjust the trustworthiness evaluation difference for all the sellers in P with a weighted average, where the corresponding weight, $Conf_{(r,s)_{c,P_j}}$, signifies the confidence level (reliability degree) of buyer c in evaluating a trustworthiness of sellers in P [22] with respect to its personal experiences (presented in Equation 3).

$$\overline{Diff}_{(c,a_k)} = \frac{\sum_{j=1}^{|P|} \mid Diff_{(c,a_k)_{P_j}} \mid * Conf_{(r,s)_{c,P_j}}}{\sum_{j=1}^{|P|} Conf_{(r,s)_{c,P_j}}} \tag{2}$$

where

$$Conf_{(r,s)} = \frac{1}{2}\int_0^1 \mid \frac{x^r(1-x)^s}{\int_0^1 x^r(1-x)^s \, dx} - 1 \mid \; dx \tag{3}$$

Afterwards, an *honesty threshold*, β where $0 \leqslant \beta \leqslant 1$, is used to determine behavioral patterns of advisers. That is, if a_k's experience with P is compatible with those of c, $(1-\overline{Diff}_{(c,a_k)} \geqslant \beta)$, a_k will be counted as a *credible* adviser, with a credibility degree, $CR_{(c,a_k)} = 1 - \overline{Diff}_{(c,a_k)}$.

In contrast, if a_k's experiences significantly deviates from the buyer agent's direct experiences, $(1 - \overline{Diff}_{(c,a_k)} < \beta)$, a_k will be detected as *malicious* adviser $(CR_{(c,a_k)} = 0)$ and would be filtered out from the buyer c's advisers network.

3 PID-based Credibility Threshold Management

Inspired by the existing electronic commerce quality models[2][1, 2, 14], we consider three factors that contribute to performance of e-marketplaces, including, 1) market liquidity (denoted by $Mliq$), 2) information asymmetry, and 3) buyers satisfaction.

Market liquidity describes a marketplace's ability to facilitate trading of the products promptly without transaction cost (i.e., having to considerably reduce their price) [3]. It also denotes the ability of buyers to find products with desirable features, when needed. However, the open nature of e-commerce, the existence of variety of products with competing features, and the lack of honesty enforcement mechanism make buyers uncertain in discovering the best-suited transaction partners (i.e., trust-wise and profit-wise),thus affecting the liquidity of the market.

Information asymmetry measures whether a buyer has sufficient information to make rational purchase decision in the e-marketplace. Higher information asymmetry is particularly salient in online environments. The buyers suffer from the risk of purchasing the low quality products, which differ from the descriptions claimed by sellers. The availability of credible advisers can effectively reduce the information asymmetry [17].

buyer satisfaction can be measured using the ratio of transactions with successful outcome to all the transactions conducted by buyers.

[2] Different from other approaches, we ascribe the performance of the e-commerce system only to the quality of its participants (buyers and sellers) in conducting transaction.

Through the proposed credibility threshold management, each buyer can further adjust her social network of credible advisers by considering the overall performance of the e-marketplace. For example, a marketplace with poor performance might imply that a considerate amount of advisers and sellers might be malicious. In this case, each buyer might want to carefully check other buyers' qualification as her advisers by increasing the credibility threshold β. In other words, when the community is populated with deceitful advisers, buyers would find it difficult to access honest feedback about sellers. Hence, the buyers should require more credible advisers by increasing β. This can help them to detect and exclude more dishonest advisers from their network, and thus obtain opinions of higher quality advisers.

If $\text{SuccessNum}_{(c)}$ denotes the number of successful outcomes achieved by c in a time stamp t, $\text{transactionNum}_{(c)}$ indicates the number of transactions conducted within t, $\text{purchaseNum}_{(c)}$ denotes the number of transactions that c initially *intended* to perform within t as indicated in its purchase mission [3], we can formulate the *transaction success rate* and the *transaction rate* of the buyer c denoted by $tp(c)$ and $tr(c)$ for the time stamp t as follows:

$$tp(c) = \frac{\text{SuccessNum}_{(c)}}{\text{transactionNum}_{(c)}} \tag{4}$$

$$tr(c) = \frac{\text{transactionNum}_{(c)}}{\text{purchaseNum}_{(c)}} \tag{5}$$

To accurately adjust β, the central server should have a global observation of the system performance. Therefore, buyers are asked to periodically share their $tr(c)$ and $tp(c)$ with the *e-marketplace central server* (ECS). The values of $tr(c)$ and $tp(c)$ reflect the behavior of participants in the e-marketplace. For example, having a high transaction rate $tr(c)$ but a low transaction success rate $tp(c)$ signifies the situation in which a buyer c is misled by dishonest advisers in her network; therefore, could not find high quality sellers.

Given these quality metrics, we propose the performance measures for e-commerce systems as follows:

$$Q(t) = \frac{2 * tp(t) * Mliq(t)}{tp(t) + Mliq(t)} \tag{6}$$

Where $Mliq(t) = \frac{\sum_{k=1}^{n} tr(c_k)}{n}$ and $tp(t) = \frac{\sum_{k=1}^{n} tp(c_k)}{n}$ are the average of all $tr(c)$ and $tp(c)$ shared by buyers at time stamp t, and $Q(t)$ is the *harmonic mean* of the e-commerce quality metrics described above. Since the performance of the marketplace is a function of these quality metrics, we use a harmonic mean to balance them by mitigating the impact of the one with a larger value and aggravating the impact of the other with a lower value.

To adjust β accordingly, ECS adopts the idea of feedback controller, specifically ***P**roportional-**I**ntegral-**D**erivative* (**PID**) controller [20]. Given a designated goal in a system, called the reference r, the feedback control system calculates the error by differentiating the actual outcome, called y, and the reference r. PID controllers provide a means to minimize the error in a system based on the received feedback [13].

In e-commerce systems, the ultimate goal is to maximize the performance of marketplaces in terms of buyers' satisfaction degree and market liquidity, achieving $Q(t) = 1$, so we initialize the goal r to $r = 1$. We designate error, $e(t)$, in the e-commerce system as the difference between the actual performance of the system $Q(t)$ and the goal r which is $e(t) = r - Q(t)$.

[3] We assume that buyers have a pre-determined purchase missions such that they enter the market to buy certain products.

In the ideal e-commerce systems in which no malicious buyers exist $Q(t)$ could converge to one. However, in a realistic situation where the marketplace is populated with different participants with various behavioral dispositions, it is not reasonable to expect the perfect performance of the system; therefore, the system will have $Q(t) < 1$.

Given these values, ECS calculates a new value for β that improves $Q(t)$ to reach the idealistic goal $r = 1$. To this end, ECS incorporates PID controller to determine the extent to which it has to change the value of β.

The new recommended value of β for the next time stamp $t + 1$ is formulated as follows:

$$\beta(t+1) = \beta(t) + \beta_0(t+1) \quad (7)$$

in which $\beta_0(t+1)$ is formalized using the PID controller presented as,

$$\beta_0(t+1) = k_p e(t) + k_i \int_0^t e(\tau)d\tau + k_d \frac{de(t)}{dt} \quad (8)$$

Where k_p, k_i, and k_d are the coefficients that leverage the contribution of *Proportional* **P**, which captures the error $e(t)$ calculated in the time stamp t, *Integral* **I**, which accumulates all errors from the start of the e-marketplace, and *Derivative* **D**, which calculates the deviation of current error $e(t)$ from its previous value $e(t-1)$, respectively.

Since in the e-marketplace it is unrealistic to expect $Q(t)$ reaches the value of r (due to the activity of malicious participants), ECS would stop adjusting β if $Q(t)$ reaches a stable point.

More formally, ECS updates the value of β for the next time stamp $t + 1$, given the following conditions:

$$\beta(t+1) = \begin{cases} \beta(t) + \beta_0(t+1) & |Q(t) - Q(t-1)| > \sigma \\ \beta(t) & \text{otherwise} \end{cases} \quad (9)$$

Where σ is a trigger threshold.
The pseudo code summary of adjusting β in the proposed PID-based credibility threshold management is shown in Algorithm 1.

```
Input:   t : starting time of e-marketplace;
         t_n: end of time of e-marketplace;
         A : set of advisers;
         C: set of buyers;
Output:  β(t);
while t ≤ t_n do
    foreach c ∈ C do
        c filters its advisers in A, based on β(t);
        c shares tr(c) and tp(c) with ECS;
    end
    ESC computes mean of transaction success rate, tp(t) ;
    ECS computes mean of transaction rate, Mliq(t);
    ECS computes Q(t) using Equation 6;
    if |Q(t) − Q(t − 1)| > σ then
        ECS computes β_0(t + 1) using Equation 8;
        ECS computes β(t + 1) using Equation 7;
    else
        β(t + 1) := β(t);
    end
    t = t + 1;
end
```

Algorithm 1: PID-based honesty threshold adjustment algorithm

4 Experimental Results

The e-marketplace environment used for experiments is populated with self-interested buyers and sellers, and is operated for 20 days. We initialize the e-marketplace with 100 buyers in total, each of which has a maximum of 5 requests everyday. Buyers (advisers) are divided into two groups: honest buyers (ones with high credibility), and dishonest buyers (ones with low credibility). Honest advisers generate ratings that differs *at most* by 0.2 points from their actual ratings. In contrast, dishonest advisers generate ratings that differs *at least* by 0.2 points from the actual experience. For example, if the seller's QoS value was 0.9, then the honest adviser would generate a value between (0.7 and 0.9), and dishonest adviser would generate a value between (0.1 and 0.69).

We assume there exist 80 sellers and 20 product types and every 4 of the sellers supply products with the same features. Sellers offer same price for products. We further assume the utility of each product is a value randomly distributed within [50,70] for all sellers. Half of sellers, who supply the same kind of product, are high-performance with QoS values in the range (0.8-1.0). On the contrary, low-performance sellers generate QoS value in a range of (0-0.2). For example, if the seller's QoS is 0.3, the utility of its product is 60 and the price is 5, a buyer's *actual profit* of carrying out a transaction with that seller would be $0.3 * 60 - 5 = 12$.

A buyer, e.g. c, calculates the trustworthiness of sellers e.g. P_j through weighted aggregation of advisers ratings, $r_{(a_k)}$, with its own recent experiences $r_{(c)}$, presented as follows:

$$\tau_{(P_j)} = \omega . r_{(c)} + (1-\omega) \frac{\sum_{k=1}^{n} CR_{(c,a_k)} * r_{(a_k)}}{\sum_{k=1}^{n} CR_{(c,a_k)}} \quad (10)$$

Buyers subjectively decide to conduct a transaction if $\tau_{(P_j)} > T$ where T indicates the transaction threshold. We further set the threshold T to be 0.6. Also, ω is determined based on Equation 18 presented in [12].

The buyer c's expected utility of carrying out a transaction with a seller P_j can be formalized as follows:

$$Exp_c^{P_j} = \tau_{(P_j)} * V_{P_j} - p_s \quad (11)$$

where V_{P_j} and p_s indicate the utility of the product promised by P_j and the price of the product, respectively.
In this experiment, the credibility degree of advisers is calculated through the presented *Credibility Evaluation Mechanism*. However, other credibility evaluation approaches can be used instead.

We conduct experiments in two settings where different groups of buyers populate different percentage of e-marketplaces: 1) *balanced environment* where 50% of buyers are honest and 50% of them are malicious, and 2) *dishonest majority* where the number of dishonest buyers exceeds that of honest ones. We set the inequalities in buyer behaviors to be significant (a 75-25 ratio imbalance is used).
We carry out comparative experiments to evaluate the performance of the e-marketplace in different environmental settings, adopting the fixed $\beta = 0.5$ versus the PID-based β.

We first measure the *market liquidity* by examining the transaction rate of different groups of buyers. Upon arrival, buyers randomly select sellers based on their promised utility (up to round 2). After acquiring sufficient experiences they establish their social network of trustworthy advisers, adopting different honesty threshold approaches: 1) the fixed β and 2) the PID-based β, which is initialized to 0.5. Given the initial setting of β, buyers have a similar transaction rate in initial days. However, we observe that as β increases, the transaction rate of the honest buyers increases while the transaction rate of dishonest advisers decreases, (Figure 1 (b)).

From Figure 1, we notice that in both honesty threshold management approaches, honest buyers have higher transaction rates compared to the dishonest ones. However, comparative results indicate

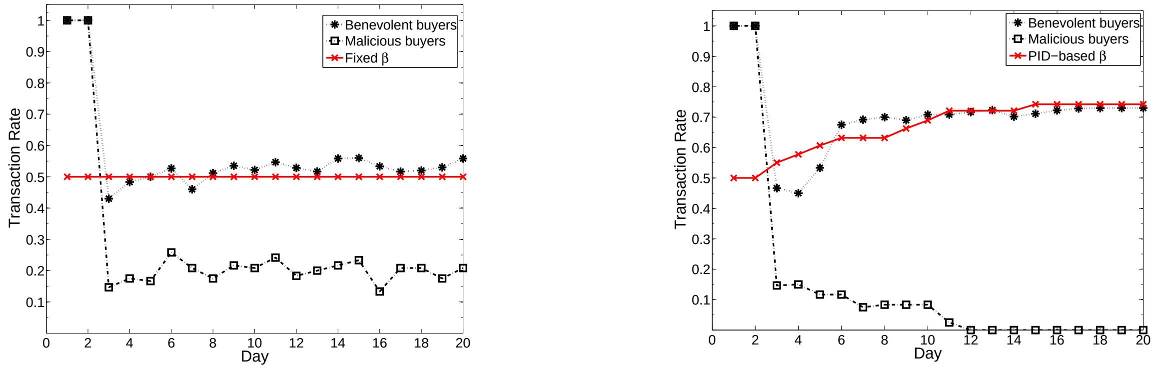

(a) **(b)**

Figure 1: The market liquidity of e-commerce system when number of honest and dishonest buyers are equal: (a) buyers adopt fixed β; (b) buyers adopt PID-based β

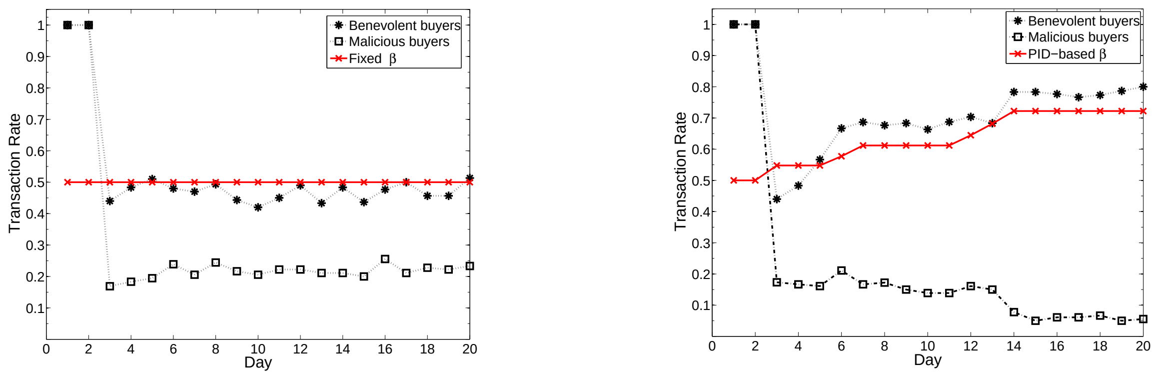

(a) **(b)**

Figure 2: The market liquidity of e-commerce system when dishonest buyers out-number honest buyers : (a) buyers adopt fixed β; (b) buyers adopt PID-based β

that in a PID-based β honest buyers have higher transaction rate than their counterparts in the fixed β approach. The adaptive approach of ECS in adjusting β, based on the quality of marketplaces, results in 1) increase of honest buyers' transaction rate, and 2) detection and isolation of more dishonest advisers.

The adaptive adjustment of β is especially important when the majority of buyers are malicious. From Figure 2 we notice that in this environment, dishonest buyers in the case with fixed β have much higher transaction rate (Figure 2 (a)) than their counterparts in the PID-based β case (Figure 2 (b)). On the contrary, honest buyers have much lower transaction rate in the fixed β than in the PID-based β case. The reason is that in the approach using fixed β a large number of dishonest advisers remain undetected and they continue to mislead buyers in their decision making process, impeding them in finding and conducting transactions with good sellers.

As can be seen in Figures (1 (b), 2 (b)), even though the value of β gradually increases and the dishonest advisers are mostly filtered away, the honest buyers cannot conduct all the transactions they initially intended (i.e., $tr(c) < 1$). This is due to a lack of experience of buyers and advisers with the sellers that they intend to make transactions with.

We measure the level of *information asymmetry* in the e-marketplace by evaluating the accuracy of buyers in classifying their advisers. As shown in Fig 3 (a), the accuracy of buyers in a PID-based β improves consistently and reaches the optimal value as they adaptively re-evaluate their network of advisers based on a new recommended value of β. On the other hand, although the static approach with fixed β shows good classification accuracy, it cannot further adapt to the incurring changes of the environment. This matter is clearly illustrated in Figure 3 (b). When the majority of participants turns to be dishonest, the classification performance of buyers with fixed β drops considerably.
In addition, by comparing the metric of precision for buyers in different environmental conditions we notice that in the fixed β approach, as many dishonest advisers are inaccurately classified as honest ones, buyers significantly rely on dishonest advisers' feedback in their decision making. The problem is aggravated when the number of dishonest buyers exceeds the number of honest ones in the marketplace (Figure 3 (b)). The performance measures (i.e., precision) reflects the ineffectiveness of the credibility evaluation mechanism (with a fixed β) in detecting malicious advisers. Therefore buyers are better off to make a random decision on finding their transaction partners instead of relying on adviser's feedback identified by such credibility evaluation mechanisms. On the contrary, dynamically monitoring and tuning β enables buyers to achieve fairly good precision value hence undermining the impact of dishonest advisers.

Note that high precision and accuracy values describe the situation where buyers can access honest feedback, which implies the e-marketplace with low level of information asymmetry.

Finally, in order to measure *buyers satisfaction* rate we compare the transaction success rate and total profit gained by different buyers[4]. Shown in Figure 4, we conclude that honest buyers provided

[4] Due to the page limit, we only present the results for the balance environment.

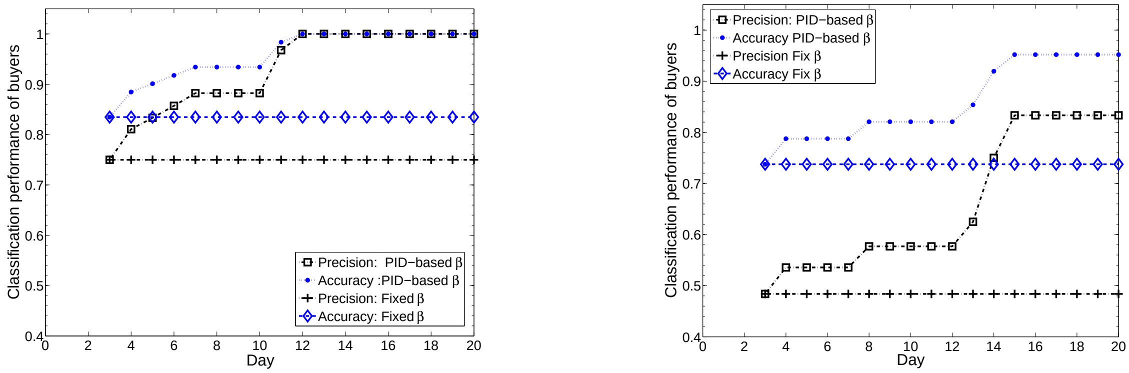

(a) **(b)**

Figure 3: The classification performance of different group of buyers in a: (a) balanced e-marketplace; (b) unbalanced e-marketplace when majority of buyers are dishonest

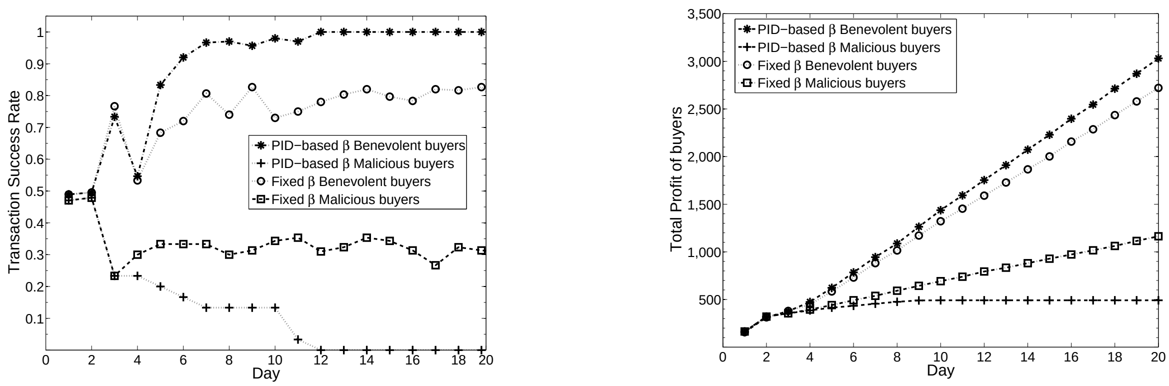

(a) **(b)**

Figure 4: (a) The transaction success rate of buyers in a balanced e-marketplace; (b) The total profit of buyers in a balanced e-marketplace

with the PID-based β conduct more successful transaction (Figure 4(a)), and gain more profit (Figure 4)(b) than other honest buyers in the fixed β approach. Specifically, the profit difference between honest buyers and dishonest buyers with the PID-based β is much larger than that of fixed β. Results indicate that, in the e-marketplaces in which buyers are equipped with a credibility evaluation mechanism with fixed β, dishonest buyers have a good chance of making profit by behaving deceitfully in the environment. This problem is especially important in competitive e-marketplaces where sellers have limited inventories and good sellers are scarce.

5 Related Work

Several credibility evaluation mechanisms have been proposed to dilute the effect of unfair or inconsistent opinions in electronic commerce systems.

In TRAVOS [19], advisers share the history of their interactions with sellers in a tuple that contains the frequency of successful and unsuccessful interaction results. Buyers calculate the probability based on a beta distribution that a particular adviser provides accurate ratings given the adviser's past reports. Once detected dishonest, advisers' rating information would be considered unreliable, and therefore discarded or discounted. Zhang [25] proposed a personalized approach for handling unfair ratings in centralized e-marketplaces. In this model, advisers share their ratings about some sellers. To estimate the credibility of advisers, buyers exploit a probabilistic approach and model advisers' trustworthiness by integrating the public and private reputation components about advisers. Noorian [9] proposed a two-layered cognitive filtering approach to detect and disqualify unfair advisers. The credibility of advisers is evaluated according to the similarity degree of advisers' opinions with those of buyers, as well as their behavioral dispositions in feedback provision. Beta Filtering Feedback [23] and RATEWeb [8] evaluate the ratings based on their deviations from the majority opinion. The basic idea of the proposed method is that if the reported rating agrees with the majority opinion, the raters credibility is increased, otherwise decreased. However, unlike other models, RATEweb does not simply discard the rating, if it disagrees with the majority opinion; instead, RATEWeb decreases the credibility of the rater by a certain degree. Wang [21] proposed super-agent framework for reputation management for service selection environment where agents with more capabilities act as super-agents and become responsible for collecting, managing and providing reputation information. Buyers adopt reinforcement learning approach to model the trustworthiness of super-agents. BLADE [15] provides a model for buyers to interpret evaluations of advisers using a Bayesian learning approach. This model does not discard all unreliable ratings; rather, it learns an evaluation function for advisers who provide ratings similar to their direct experience. BLADE applies a strict judgment on the credibility of feedback providers. For example, BLADE discounts the ratings of advisers with an honesty degree of 0.7.

These existing trust models, however, do not address how they distinguish trustworthy advisers from untrustworthy ones. That is, these models cannot answer the following questions: 1) How to define the acceptable level of honesty, trustworthiness and/or similarity of an adviser?, 2) How to define the credibility adjustment threshold? To the best of our knowledge, in the existing literature, the honesty threshold has been either explicitly initialized by a central server, as in [19, 21, 24, 25] or has been subjectively determined by buyers according to their behavioral characteristics as presented in [8, 9]. The only previous work that addresses these questions is FIRE [5], which defines an adaptive inaccuracy tolerance threshold based on the sellers' performance variation to specify the maximal permitted differences between the actual performance and the provided ratings. This work is different from our approach, however, since in FIRE each buyer filters away advisers based on their local observation on a quality of the sellers, and thus this model suffers from the risk of unfair judgment of advisers.

6 Conclusion and Future Works

This paper pinpoints a common problem of existing trust and reputation systems in electronic commerce systems. Despite the significant advances of the field in detecting and mitigating misbehavior of untruthful participants, these models rely on certain assumptions. One of the most important of these assumptions that has been trivialized in the literature is the existence of a honesty threshold, which serves as decision boundary to separate participants based on their behavioral characteristics. The choice of values for this "magic" threshold is usually left to the designers implementing a particular system.

We address this problem by designing a controller method to adaptively tune the honesty threshold. The proposed controller monitors the quality of e-marketplace and uses a PID feedback controller technique to determine new values for the honesty threshold to adapt to the changing marketplace.

The standalone and context-independent design of the proposed PID-based credibility threshold adjustment makes it well-suited to be incorporated with different credibility evaluation mechanisms and filtering models for electronic marketplaces.

Experimental results show the advantages of adaptive evaluation on the honesty threshold. In particular, we demonstrate that credibility evaluation mechanism guided by PID-based threshold management techniques can increase market liquidity, buyers' satisfaction, and decrease the information asymmetry in the e-marketplace.

Credibility mechanisms using an adaptive honesty threshold to the feedback received from the marketplace provides better accuracy with time, since they have the ability to evolve and dynamically evaluate the changing conditions of the marketplace.

An interesting direction for future work would be to improve the feedback controller method by adopting different dynamic performance metrics supported in the market microstructure literature [7], in addition to those considered here. Furthermore, since the buyers' contribution in providing feedback is an essential elements in the performance monitoring of the marketplace, a useful direction for future work would be the incorporation of an incentive mechanism to promote more participation (in terms of providing honest feedback) from the buyers.

REFERENCES

[1] Stuart J Barnes and Richard Vidgen, 'Measuring web site quality improvements: a case study of the forum on strategic management knowledge exchange', *Industrial Management & Data Systems*, **103**(5), 297–309, (2003).

[2] William H Delone, 'The delone and mclean model of information systems success: a ten-year update', *Journal of management information systems*, **19**(4), 9–30, (2003).

[3] Michael Fleming, 'Measuring treasury market liquidity', *FRB of New York Staff Report*, (133), (2001).

[4] Yasaman Haghpanah and Marie desJardins, 'Prep: a probabilistic reputation model for biased societies', in *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*, AAMAS '12, pp. 315–322, (2012).

[5] T. Dong Huynh, Nicholas R. Jennings, and Nigel R. Shadbolt, 'Fire: An integrated trust and reputation model for open multi-agent systems', in *In Proceedings of the 16th European Conference on Artificial Intelligence (ECAI*, pp. 18–22, (2004).

[6] Athirai Aravazhi Irissappane, Siwei Jiang, and Jie Zhang, 'A framework to choose trust models for different e-marketplace environments', in *IJCAI*, (2013).

[7] Ananth Madhavan, 'Market microstructure: A survey', *Journal of Financial Markets*, **3**(3), 205–258, (2000).

[8] Zaki Malik and Athman Bouguettaya, 'Rateweb: Reputation assessment for trust establishment among web services', *The VLDB JournalThe International Journal on Very Large Data Bases*, **18**(4), 885–911, (2009).

[9] Zeinab Noorian, Stephen Marsh, and Michael Fleming, 'Multi-layer cognitive filtering by behavioral modeling', in *The 10th International Conference on Autonomous Agents and Multiagent Systems - Volume 2*, AAMAS '11, pp. 871–878, (2011).

[10] Zeinab Noorian, Stephen Marsh, and Michael Fleming, 'Multi-layer cognitive filtering by behavioral modeling', in *The 10th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*, pp. 871–878. International Foundation for Autonomous Agents and Multiagent Systems, (2011).

[11] Zeinab Noorian and Mihaela Ulieru, 'The state of the art in trust and reputation systems: a framework for comparison', *J. Theor. Appl. Electron. Commer. Res.*, **5**, 97–117, (August 2010).

[12] Zeinab Noorian, Jie Zhang, Yuan Liu, Stephen Marsh, and Michael Fleming, 'Trust-oriented buyer strategies for seller reporting and selection in competitive electronic marketplaces', *Autonomous Agents and Multi-Agent Systems*, 1–38, (2013).

[13] Hitay Ozbay, *Introduction to feedback control theory*, CRC Press, 2000.

[14] Ananthanarayanan Parasuraman, Valarie A Zeithaml, and Arvind Malhotra, 'Es-qual a multiple-item scale for assessing electronic service quality', *Journal of service research*, **7**(3), 213–233, (2005).

[15] Kevin Regan, Pascal Poupart, and Robin Cohen, 'Bayesian reputation modeling in e-marketplaces sensitive to subjectivity, deception and change', in *Proceedings of the National Conference on Artificial Intelligence*, volume 21, p. 1206. AAAI Press, (2006).

[16] Cristiano Castelfranchi Rino Falcone, 'Socio-cognitive model of trust', in *Encyclopedia of Information Science and Technology (V)*, (2005).

[17] Gregory D Saxton and Ashley E Anker, 'The aggregate effects of decentralized knowledge production: Financial bloggers and information asymmetries in the stock market', *Journal of Communication*, **63**(6), 1054–1069, (2013).

[18] Michael J Smith et al., 'Learning to trust in the competence and commitment of agents', *Autonomous Agents and Multi-Agent Systems*, **18**(1), 36–82, (2009).

[19] WT Luke Teacy, Jigar Patel, Nicholas R Jennings, and Michael Luck, 'Travos: Trust and reputation in the context of inaccurate information sources', *Autonomous Agents and Multi-Agent Systems*, **12**(2), 183–198, (2006).

[20] Antonio Visioli, *Practical PID control*, Springer, 2006.

[21] Yao Wang, Jie Zhang, and Julita Vassileva, 'A super-agent framework for reputation management and cimmunity formation in decentralized systems', *Computational Intelligence*, (2014).

[22] Yonghong Wang, Chung-Wei Hang, and Munindar P Singh, 'A probabilistic approach for maintaining trust based on evidence', *Journal of Artificial Intelligence Research*, **40**(1), 221–267, (2011).

[23] Andrew Whitby, Audun Jøsang, and Jadwiga Indulska, 'Filtering out unfair ratings in bayesian reputation systems', in *Proc. 7th Int. Workshop on Trust in Agent Societies*, (2004).

[24] Li Xiong and Ling Liu, 'Peertrust: Supporting reputation-based trust for peer-to-peer electronic communities', *IEEE Transactions on Knowledge and Data Engineering*, **16**, 843–857, (2004).

[25] Jie Zhang and Robin Cohen, 'Evaluating the trustworthiness of advice about seller agents in e-marketplaces: A personalized approach', *Electronic Commerce Research and Applications*, (2008).