

A Market-Based approach to Managing the Risk of Peer-to-Peer Transactions

Stephanos Androutsellis-Theotokis
Diomidis Spinellis
Panos Louridas
Kostas Stroggylos

Department of Management Science and Technology
Athens University of Economics and Business
Athens, Greece

Abstract

Market-based principles can be used to manage the risk of distributed peer-to-peer transactions. This is demonstrated by PTRIM, a system that builds a transaction default market on top of a main transaction processing system, within which peers offer to underwrite the transaction risk for a slight increase in the transaction cost. The insurance cost, determined through market-based mechanisms, is a way of identifying untrustworthy peers and perilous transactions. The risk of the transactions is contained, and at the same time members of the peer-to-peer network capitalise on their market knowledge by profiting as transaction insurers. We evaluated the approach through trials with the deployed PTRIM prototype, as well as composite experiments involving real online transaction data and real subjects participating in the transaction default market. We examine the efficacy of our approach both from a theoretical and an experimental perspective. Our findings suggest that the PTRIM market layer functions in an efficient manner, and is able to support the transaction processing system through the insurance offers it produces, thus acting as an effective means of reducing the risk of peer-to-peer transactions. In our conclusions we discuss how a system like PTRIM assimilates properties of real world markets, and its potential exposure and possible countermeasures to events such as those witnessed in the recent global financial turmoil.

Keywords Peer-to-peer networks; transaction processing; reputation management; financial markets

1 Introduction

It is being progressively recognised that information systems and applications supporting collaborative tasks, including online transaction processing systems, that currently follow centralized client-server models can also be based on the maturing wave of peer-to-peer architectures [SW05, ATS04]. In order to manage and reduce the risk inherent in peer-to-peer transactions and their decentralised and uncontrolled environment, a variety of approaches have been proposed, with reputation and trust management systems being the most prominent (see Section 2.2). These aim to provide peers with estimates of the risk involved in their transactions, based on the observed past behaviour of their counterparties.

Though reputation management systems (either centralized or distributed) offer a lot in this direction, the information they provide about past behaviour may not be enough to accurately assess the risk involved in a transaction. Open issues, to name but a few, include the fact that not all transactions receive feedback [RZ02], shortcomings in the way reputation information is represented [Del04], perceived, and used by the transacting parties (especially for inexperienced users) [GS06], and dealing with participants lacking past transactions [FPCS04]. In a recent study, for instance, it was suggested that the way in

Computer Networks, 54(5):675–688, April 2010.

This is a machine-readable rendering of a working paper draft that led to a publication. The publication should always be cited in preference to this draft using the reference in the previous footnote. This material is presented to ensure timely dissemination of scholarly and technical work. Copyright and all rights therein are retained by authors or by other copyright holders. All persons copying this information are expected to adhere to the terms and constraints invoked by each author's copyright. In most cases, these works may not be reposted without the explicit permission of the copyright holder.

which the positive and negative feedback is presented to users of the eBay system can make it difficult for buyers to evaluate past illegal behaviour of sellers [GS08]. As a result, there still is considerable concern over the amount of risk involved in online transactions, even with the support of reputation management systems [FTW05, DR03, GT08].

In this work we propose a novel approach, based on the principles and instruments used in different (including commercial and financial) markets for managing, transferring or reducing credit and transaction risk. We demonstrate the practicality of this approach and evaluate its efficacy through the prototype PTRIM (“Peer-to-Peer Transaction Risk Management”) system. On top of a generic transaction processing system, PTRIM creates a peer-to-peer transaction insurance market-like layer that is used to manage the risk of transaction default. The transacting peers have the option to request offers from peers in this layer to underwrite the risk of their transaction, and therefore alleviate themselves from the need to collect, process and evaluate reputation information. A transaction default market is thus built on top of the main transaction processing system, within which the cost of “insuring” a transaction is determined through market-based mechanisms.

The main contribution of our approach is that it proposes a novel way of identifying untrustworthy peers and perilous transactions, based on the insurance offers the market will produce; the risk of peer-to-peer transactions is contained at a small additional cost; and at the same time, our system showcases a way for members of a peer-to-peer network to participate in a new profitable role, as transaction insurers.

It should be noted that although we implement and examine our proposed approach within a peer-to-peer environment, its underlying principles could also apply to different, not completely decentralised network architectures.

In this paper, we first present the key concepts of market-based transaction risk management that our approach is based on, briefly touching on the underlying economics principles, and we examine its relation to other work (see Section 2). We proceed to present the design, architecture and implementation of our proposed system in Section 3, and our methodology for testing and evaluating it both theoretically and practically (see Section 4), which includes studies with a deployed prototype, as well as experiments involving both real subjects and real transaction data crawled from the eBay online auction site. We discuss our evaluation results and outcomes in Section 5, and our system’s current limitations and planned future work in Section 6.

2 Market-based transaction risk management

As was briefly described in our introduction, in our approach we form a market for transaction default insurance, and utilise it to determine and manage the risk involved in transactions. We propose this as an alternative to reputation management systems, although the two could supplement each other.

2.1 Risk management

We require a market-based transaction risk management system that would meet the following requirements:

1. transactions should not need an intermediary; i.e., it should be possible to carry out a transaction involving only the transacting peers, and not a centralised clearing house.
2. risk handling should be based on trading insurance; i.e., participants would offer and buy insurance against a transaction.

The first requirement is met by Over The Counter (OTC) financial instruments: instruments that are carried out by transacting entities negotiating and reaching agreement directly, without an intermediary. The second requirement means that, since insuring participants are trading not on the underlying asset but on an asset derived from it (insurance risk), we will be dealing with a model resembling financial derivatives [Bon04, BOW02, Hul08]. In finance, derivatives are financial instruments whose performance depends on another variable, and are often used to transfer or reduce (hedge) risk involved in assets or transactions. These instruments are typically settled upon the occurrence of a credit default event, and the characterisation of the risk involved dictates their cost.

The design of our proposed system and its market layer revolves around the above two concepts.

2.2 Relation to distributed reputation management

Our approach is proposed as a potential alternative to distributed reputation management systems supporting peer-to-peer transactions.

A considerable amount of groundbreaking work has been carried out in the distributed reputation management field in recent years, aiming at providing an expectation about a peer’s behaviour in a transaction by monitoring, maintaining and distributing information about its behaviour in past transactions.

A variety of solutions have been proposed for addressing either or both of the data modelling problem (how to generate, interpret and process the reputation data), and the data management problem (how to store, retrieve, distribute and secure the reputation data in a scalable and efficient manner) [AD01]. Some notable systems in this area include the EigenTrust system [KSGM03], PeerTrust [XL04], Credence [WGS05], a system proposed by Aberer et al based on the P-grid structured routing algorithm [AD01, Abe01], a Bayesian approach proposed by Buchegger et al [BB04], TrustMe [SL03], XRep [DDCdVP⁺02], a partially centralized mechanism presented in [GJA03], to name but a few. Due to space constraints, we refer the reader to a comprehensive overview by Huaizhi and Singhal [HS07] and the references therein.

We want to stress that distributed reputation management systems and our proposed approach are by no means mutually exclusive, and could potentially complement each other within the same scope, supporting the same transaction management system.

2.3 The PTRIM market-based insurance layer

In our design, a peer-to-peer layer acts as a market offering transaction default insurance to a main, peer-to-peer transaction processing system (that is also part of our implemented system). Any peer can participate in this market: One could be regularly active as a buyer or seller, and sporadically choose to also offer insurance for specific transactions (e.g. involving peers that they know can be trusted such as friends, relatives or peers they have interacted with in the past); or another peer could act as a dedicated “insuring institution”, only involved in the business of providing insurance for transactions between other peers (for example a bank wishing to obtain additional income from its database of customer credit histories and its IT infrastructure for determining a customer’s credit default risk).

The risk default market thus operates as an insurance (or financial derivatives) market. Transacting peers can request offers for insurance protection against their transaction default risks, evaluate such offers and select one (or more, see Section 3) of the most competitive ones. As a result, the peers can significantly reduce their transaction risk, for a small increase in their transaction cost. Furthermore, and in contrast to the arbitrary numerical scale used in many reputation management systems [HS07, DR03], the cost of the insurance offers received by a peer directly reflects the subjective knowledge of the risk involved (knowledge about past outcomes or lack of such data). The outcomes of insurance transactions are made available to the market by the system. They can thus be used directly by transacting peers to decide as to whether to proceed with a transaction or not, or indirectly by the peers in the insurance layer to sell transaction insurance.

The triggering event for settling insurance transactions in our case would be the successful or unsuccessful completion of a transaction.

A feature of our proposed approach is therefore that it does not require the peers to be involved in the collection and management of reputation information, or to engage in the decision-making process of whether to proceed with a transaction based on the available reputation information. As discussed above, although very promising work has been done in this direction, in practical terms such information may be incomplete, ambiguous, or presented in ways that do not clearly indicate what course of action the peer should follow (e.g. should a transaction involving a \$20,000 product be initiated with a peer whose reputation score is 0.86/1.00? What if the product costs \$20 and the user’s reputation is denoted as “three out of five stars”?).

2.4 Market efficiency

In order for the system to effectively support peer-to-peer transactions through the insurance cost offers it provides to the transacting peers, the market layer it creates should be efficient.

In an efficient market no participant should be presented with an opportunity to earn more than a fair return on the riskiness associated with a transaction. An efficient market encourages participation, helps allocate goods or services in the most efficient way, and gives correct signals to be utilized by the market players. Note that market efficiency as used here is distinct from the efficient market hypothesis, which broadly postulates that all information on a product is included in its price valuation [Fam70]; the

efficient market hypothesis has been the basis of finance research for many years, but has come under criticism (e.g., see [Shi05]). Our use of the term is related to perfectly competitive markets. Specifically, we claim that the insurance market described here approximates the main characteristics of perfectly competitive markets. The notion of perfect competition has a long history [Sti57]. For our purposes the following characteristics will do [KWO07]:

Participants are price-takers

A perfectly competitive market comprises a large number of “small” transacting entities whose individual actions can have no impact on others or on the entire market, and none of whom have a large market share. This is the case for the individual peers in our insurance market and their actions, whose scope is limited per transaction; an individual’s actions do not influence the market price of the insurance.

Homogeneity

There is no differentiation in the services offered by the insuring peers. In fact in our market they all offer exactly the same service, so any peer could be substituted by another one. In other words, the insurances on offer are standardized products, that is, commodities.

Equal access and free entry

These are guaranteed by the completely decentralized nature of the peer-to-peer architecture PTRIM is based upon. All peers function in the same way and are exposed to the same information. No obstacles are presented to any entity that would wish to enter the network and provide its services.

In our discussion on methodology (Section 4) we describe different approaches to measuring the degree of efficiency of a market, and how we utilize them as a means of evaluating the performance of our proposed system.

3 System description

We briefly describe the architecture, design, implementation and deployment of the PTRIM system.

3.1 System architecture and design

Our system consists of three peer-to-peer layers, as shown in the schematic representation in Figure 1. The infrastructure layer supports the peer-to-peer nature of the system; the transaction processing layer allows peers to participate in online transactions; and the risk management layer is where the insurance (derivatives) market functions.

Peers can participate in any or all of the layers, even simultaneously, in different roles:

- **Rendezvous peers (R)** offer the necessary infrastructure for relaying messages between other peers and providing the network connectivity and availability.
- **Transacting peers (B,S)** advertise products or services they wish to sell, place offers for buying other peers’ products, and eventually participate in transactions as buyers or sellers.
- **Insuring peers (In)** offer insuring services to the transacting peers, and are remunerated in the form of a percentage on the amount they are insuring.

For each transaction they participate in, the peers maintain state information as well as records of all exchanged messages and other data, and they communicate by exchanging messages asynchronously in a distributed manner. Figure 2 shows a state diagram illustrating the transition between different states, and the associated state flag values for each. For simplicity we initially only show the state transitions for a peer acting as a buyer (the states for a peer acting as a seller are symmetric), for a transaction not involving insurance. The figure also shows a UML sequence diagram illustrating a simple example of a transaction involving four peers. Peers $p1$ and $p2$ engage in a transaction as seller and buyer respectively, and peers $p3$ and $p4$ respond to an insurance request placed by $p1$. Note that this is a relatively straightforward example, with the transaction completing successfully.

Each message is characterised by a message type and contains an ID and timestamp, sender and recipient identification, as well as other context-specific data fields that are parsed by the recipient.

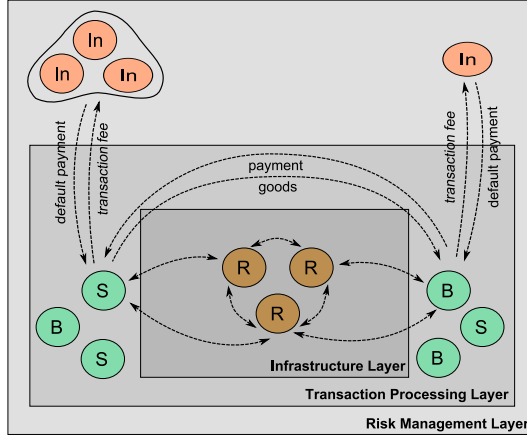


Figure 1: A schematic abstract representation of the PTRIM system. At the core lies the infrastructure layer, based on existing P2P protocols, such as a set of dynamic relay peers (R). The main transaction processing layer consists of buyers (B) and sellers (S) engaged in transactions based on asynchronous message exchanges; above it, the risk management layer includes other peers (In) that offer to insure the transacting peers, either independently or collectively.

Transaction Processing Layer	Risk Management Layer
<ul style="list-style-type: none"> • Maintain, update and advertise in goods descriptions. • Place offers for advertised products. • Review and accept offers, initiate local transaction records, exchange messages to bind their respective records. • Notify application of receipt of products/payments, and complete transaction. • In case of failed transactions, exchange messages notifying the transacting counterpart. 	<ul style="list-style-type: none"> • Transacting peers broadcast requests for insurance offers. • Insuring peers transmit offers to insure transactions . • Transacting peers review insurance offers, accept most favourable ones, initiate insurance transactions. • Transacting peers notify insuring peers of transaction outcome, request reimbursement or pay premium. • Insuring peers applications announce transaction outcome information.

Table 1: The main tasks carried out within the Transaction Processing Layer and the Risk Management Layer of the PTRIM system.

Table 1 briefly itemizes the main tasks that are carried out by the peers (either the users, or the applications themselves), within the scope of the transaction processing layer and the risk management (insurance) layer.

As figure 3 shows, at an abstract level the PTRIM application nodes consist of the following four functionality subsets:

- User interface. This provides the communication with the application user. Allows input of commands, and displays the output of operations.
- Communication module. This is based on the JXTA technology, and allows asynchronous communication with other nodes through exchange of messages. These messages are set up by the peer node, and can be either targeted towards one specific node, or broadcast to the entire network, or a subset of it.
- Data store. Each node locally stores various data regarding transactions, products, offers placed or received, etc.
- Application logic. This is the main processing unit that carries out the application logic based on the input received from the users, messages from other nodes, and locally stored data. This unit

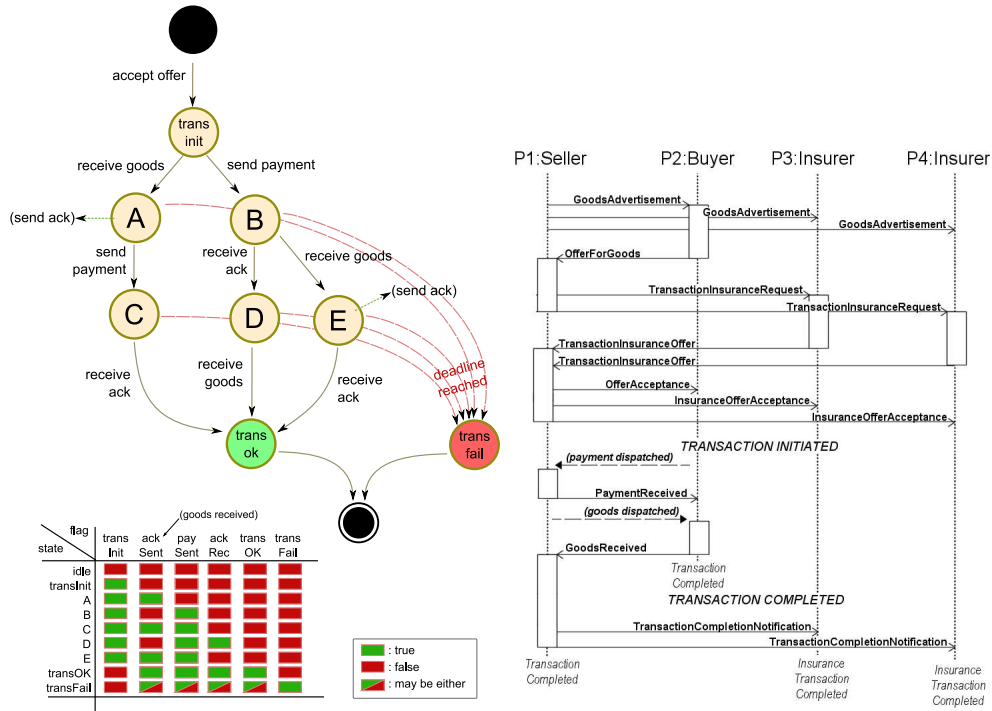


Figure 2: (a) A UML state diagram for a simple transaction (not involving insurance) only from the perspective of a node buying goods. The values of the state flags are shown below; (b) A UML sequence diagram showing two peers ($p1, p2$) engaging in a successful transaction as seller and buyer respectively, with two peers ($p3, p4$) jointly insuring the seller.

executes various “scenarios” in an orchestrated manner, based on exchanges of messages with other peers and input from users.

3.2 Prototype system implementation

The entire system has been implemented using the JXTA[Jxt07] technology—a set of open protocols that allows connected devices on the network to communicate and collaborate in a peer-to-peer manner. JXTA provides the core peer-to-peer infrastructure, allowing nodes to enter or leave the network at any time, and messages to be routed around network or node failures, based on a transient number of “rendezvous” (or relay) peers.

The application, along with documentation, data and experimental outcomes, is available as an open-source project at:

<http://istlab.dmst.aueb.gr/~path/software/ptrim>.

Our system has been deployed and is currently in operation within the domain of our academic institution. More wide-scale deployment is scheduled after we enhance it with more robust non-functional characteristics, mainly related to security (see Section 6.1).

4 Evaluation methodology

In order to devise a comprehensive evaluation methodology, and given that PTRIM adopts an original approach, we looked at some of the most prominent peer-to-peer reputation management systems proposed, and in particular the metrics and methods they used to represent and evaluate their efficiency. Table 2 summarises our findings.

Aiming to provide evaluation metrics and methods analogous to the above (to the extent they are applicable to the PTRIM system), we adopted the following two-stage evaluation methodology:

- 1. Prototype system trials.** Sets of “proof-of-concept” trials were carried out based on our implemented and deployed prototype system. These are described in Section 4.1. This approach was

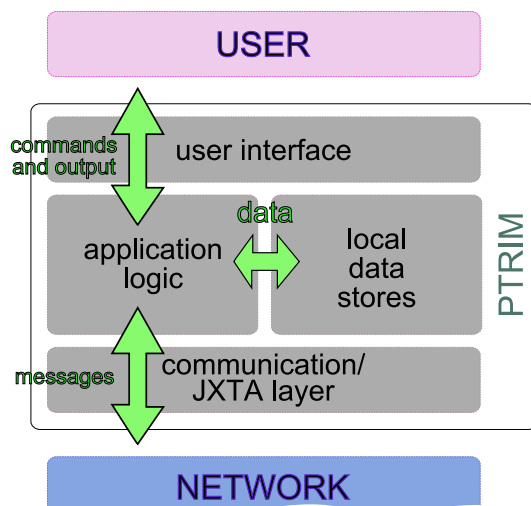


Figure 3: An abstract schematic of the architecture of the PTRIM application nodes and the four functionality subsets they consist of.

	PeerTrust (Nejdl)	PeerTrust (Xiong)	(Datta)	(Gupta)	(Dutta)	Credence (Walsh)	TrustMe (Singh)	(Buchegger)	EigenTrust (Kamvar)	PGrid (Aberer)	PTRIM
Robustness:											
Trust computation error	■	■	■	■	■	■	■	■	■	■	■
Identification of misbehaving peers	■	■	■	■	■	■	■	■	■	■	■
Transaction success rate	■	■	■	■	■	■	■	■	■	■	■
Overall credit loss	■	■	■	■	■	■	■	■	■	■	■
Correlation with peer history	■	■	■	■	■	■	■	■	■	■	■
Coordination of disagreement	■	■	■	■	■	■	■	■	■	■	■
False +ves/-ves	■	■	■	■	■	■	■	■	■	■	■
Cost:											
Bandwidth (messaging)	■	■	■	■	■	■	■	■	■	■	■
Storage	■	■	■	■	■	■	■	■	■	■	■
Computational	■	■	■	■	■	■	■	■	■	■	■
Response time:											
Misbehaving peer detection time	■	■	■	■	■	■	■	■	■	■	■
Time to steady-state convergence	■	■	■	■	■	■	■	■	■	■	■
Threat resistance											
Threat resistance	■	■	■	■	■	■	■	■	■	■	■
Deployment results											
Deployment results	■	■	■	■	■	■	■	■	■	■	■

	PeerTrust (Nejdl)	PeerTrust (Xiong)	(Datta)	(Gupta)	(Dutta)	Credence (Walsh)	TrustMe (Singh)	(Buchegger)	EigenTrust (Kamvar)	PGrid (Aberer)	PTRIM
Simulation											
Simulation	■	■	■	■	■	■	■	■	■	■	■
Use of real data											
Use of real data	■	■	■	■	■	■	■	■	■	■	■
Behaviour modelled											
Behaviour modelled	■	■	■	■	■	■	■	■	■	■	■
System deployment											
System deployment	■	■	■	■	■	■	■	■	■	■	■
Mathematical analysis											
Mathematical analysis	■	■	■	■	■	■	■	■	■	■	■
Discussion of threats											
Discussion of threats	■	■	■	■	■	■	■	■	■	■	■

Table 2: Key metrics used to represent and evaluate system efficiency in ten of the most prominent distributed reputation management systems; and key methods used to measure system efficiency in ten of the most prominent distributed reputation management systems.

only chosen in two out of the ten systems in Table 2.

2. Composite experimental trials. Additional “composite experiments”, described in Section 4.2, were carried out aiming at obtaining larger and more statistically significant results. Our approach here covered both the widely used simulation studies, and the more rare use of real transaction data in Table 2. Modelling different user behaviours was not possible since instead of simulating the insuring peers we used real subjects, but this is in our current future work plans. In terms of metrics, this approach allows us to examine overall credit loss, identification of misbehaving peers, correlation of peer history with insurance cost and other specific measurements discussed in Section 5.

4.1 Prototype system trials

To validate our design and implementation we carried out sets of small-scale trials within controlled groups of subjects (mostly university students), who were instructed to act either as transacting or insuring peers. All nodes also participated as rendezvous peers by default, to support the communication infrastructure.

This approach was targeted toward examining the feasibility of implementing our system based on

specific technologies, and with acceptable bandwidth and storage costs.

The operation of the entire network was followed by a specially designed monitoring peer that regularly requested details about all events and transactions carried out. The results of these trials are discussed in Section 5.1.

4.2 Composite experimental trials

These experiments were aimed at obtaining larger and statistically significant volumes of result data.

4.2.1 Experimental setup

In these experiments, real online transaction data was obtained by crawling the eBay web site (see Section 4.2.2). This data was then utilised as transaction input to our system, and real subjects were invited to participate as insurers. This allowed a much larger number of transactions to be processed in a relatively short period of time. These experiments typically lasted around two weeks, to cover large enough numbers of transactions, so the subjects could have the opportunity to assess their strategies, and the market could mature.

The goal of these experiments was to assess the effectiveness with which our system manages peer-to-peer transaction risk, and the efficiency of its derivatives market layer, through the properties of the insurance offers placed. This was done through a web-based application we developed that presented the crawled transaction details to the subjects, and allowed them to place insurance offers for them and request specific premiums. The main components of the system are outlined in figure 4.

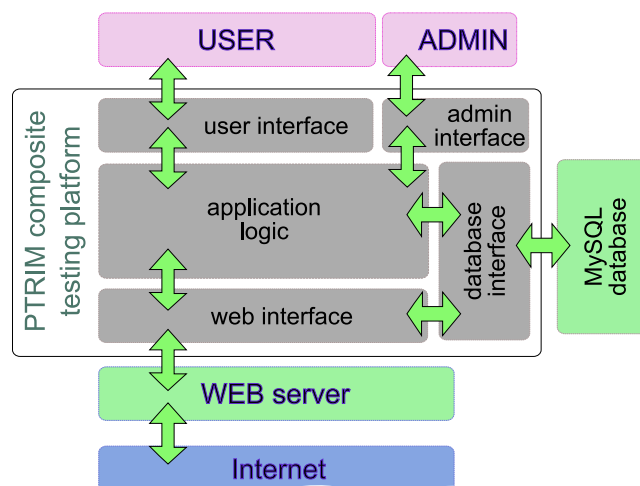


Figure 4: An abstract description of the composite trial platform architecture.

The insurers were presented with a number of transactions for which insurance was requested. Transaction data shown included information about the transacting peer trustworthiness, past transaction outcomes, amounts, etc. The insurers would then choose what transactions to place offers for, and what premiums to request. After a period had elapsed and enough offers were placed, the system accepted the best ones (if any), carried out the transaction producing an outcome statistically based on the peers' trustworthiness, and accordingly rewarded or debited the relevant insurer. This data was collected and analysed and the results are discussed in Section 5.2.

The outcomes and rates requested for all other completed transactions were also made available to the insurers, providing a relatively complete information set about the market within which they were operating.

The subjects (typically around twenty) of the experiment were executives with graduate education, but with no particular experience with such a role in a market (see [Dam02], Chapter 6, for a discussion of the acceptability of such subjects in related studies).

At the end of these experiments, the subjects were also requested to fill in a questionnaire of mostly

qualitative nature to complement the data obtained. Rough descriptions of their applied bidding strategies were also collected and correlated with their outcomes¹.

4.2.2 Transaction data

As mentioned, this experiment was based on real transaction and user data, including past transaction and reputation information, that was crawled from the eBay online auction and shopping website.

The collection of data was performed with the use of CrawlWave, a prototype distributed web crawler that was developed for research purposes [KSS04]. The crawler was configured to collect only pages containing eBay user profiles using a random walk approach within the eBay website only. The required information for each user was then extracted from the collected profile pages using regular expressions. In total, data from approximately ten thousand transactions, and all relevant users, was collected².

Since the user ratings on eBay are predominantly almost perfect (on average 99.8% positive [GS08]), we applied a rare event simulation procedure known as importance sampling [Hei95], typically used in simulating events such as network or queuing system failures, but also in portfolio credit risk studies. The goal was to expose the subjects to considerably more transaction default events than they would regularly encounter, and examine how this would affect their bidding strategies and the insurance market in general.

Importance sampling is based on the idea of making the occurrence of some events (in our case transaction failures) more frequent, by selecting a biased sampling distribution for obtaining the simulation data samples from the data universe (in our case the crawled data).

In our implementation we assume an original sampling distribution (typically homogeneous sampling) $f(x)$ and a property of interest $g(x)$, where the probability of $g(x) > 0$ is very small when taking samples from $f(x)$. With Importance Sampling, instead of $f(x)$ a sampling distribution $f^*(x)$ is used, to significantly increase the probability that $g(x) > 0$. Normally the observations made should then be corrected for the fact that $f^*(x)$ is applied instead of $f(x)$ in the sampling distribution³.

In our case the biased sampling function $f^*(x)$ applied is of the form:

$$C1 + C2 \cdot e^{x/k}$$

where C_1, C_2 and k are constants trimmed to provide the desired bias.

4.2.3 Evaluating market efficiency

The efficiency of the derivatives market is important in ensuring the correct operation of the system. Only through an efficient and competitive market will the rational insurance offers produced reflect the underlying risk in the transactions. There are various approaches to evaluating the efficiency of a market, most of which do not always produce acceptable results in practice.

A standard approach is measuring price dispersion (or price convergence) [Kal94, LG02, KR05], i.e. the variation of prices across sellers (in our case insurers) for similar items (in our case transactions). This is known as the law of one price, stating that identical offerings should have identical prices [Kal94], and can be simply explained by arbitrage arguments. However the observed price dispersion often violates this law, even in online markets, with up to 30% measured dispersion [KR05, Hr03] (game theoretic arguments have been proposed as explanations for this phenomenon [KR05]). Metrics for price dispersion can vary, including price range, percentage difference between highest and lowest price, and standard deviation or variance [BMS06].

¹The questionnaires and results are available online.

²This data is available online.

³To correct for the fact that $f^*(x)$ is applied instead of $f(x)$ in the sampling distribution, if the property of interest is $\gamma = E[g(x)]$, then this can be rewritten as

$$\begin{aligned} \gamma = E_f[g(x)] &= \int g(x)f(x)dx = \int g(x)\frac{f(x)}{f^*(x)}f^*(x)dx = \\ &= E_{f^*}[g(x) \cdot L(x)] \end{aligned}$$

where $L(x) = f(x)/f^*(x)$ is denoted as the likelihood ratio. The expected value of the observations under f is equal to the expected value under f^* corrected for bias by the likelihood ratio, i.e. $E_f[g(x)] = E_{f^*}[g(x) \cdot L(x)]$ [GL03, GLFH].

Another measure is the correlation between price and quality ([Kal94, Dam02]). In this case an ordinal measure of quality is required, and can often be difficult to obtain objectively (sometimes customer reports are used, but this carries other inherent limitations).

Studying the growth of the market in time, as well as the overall wealth accumulated, are additional significant indicators of the effectiveness and efficiency of the market, as is the correlation between the risk taken and the potential reward.

Market efficiency can also be deduced through a study of price changes and their time series properties [Dam02], while price elasticity of demand is also expected to be higher in a more efficient market [Hr03].

In our evaluation we focused on the more measurable of the above indicators, given our experimental setup, which include overall market maturity and growth, correlation between price and quality, correlation between risk and reward, as well as various qualitative observations collected through questionnaires, which help support and explain our findings.

5 Results and discussion

We discuss our results from both the prototype system “proof-of-concept” trials, and the composite experiments, whose goal is to examine the efficiency of the derivatives market and the system’s risk management capabilities.

5.1 Prototype system trials

Table 3 summarises key characteristics of a typical trial ran with the deployed prototype system.

Duration:	4 hours
Transacting peers:	27
Insuring peers:	16
Products advertised:	32
Total transactions:	23
Messages transmitted:	1378
Average premium:	24%

Table 3: Summary of a typical prototype system trial characteristics

Due to the small scale of these first trials no quantitative measurements relevant to the systems ability to reduce transaction risk are made, and we focus mainly on system stability issues.

JXTA proved to be an efficient platform for our application, both in terms of performance, messages exchanged and scalability (this was indeed a reason for selecting it). The largest number of messages exchanged between the peers concerned broadcast advertisements for products offered, and requests for insurance. The overall bandwidth utilised is found to be negligible for trials of this scale, and is known to scale gracefully due to the JXTA infrastructure. This is also due to the fact that each transaction takes a considerable amount of time to complete, with independent peers sporadically performing operations.

Although no statistically significant data is collected, the monitored transaction logs indicate that within a short running period and after a few transactions are performed, insurance offers for transactions involving misbehaving peers have progressively higher rates, as the poor performance of previous insurance transactions is noted. As a result, peers that considered engaging in transactions with misbehaving peers are faced with increasingly more disadvantageous insurance offers, which prove to be a reasonable disincentive for them. The cost of insuring a transaction derived by the insurance market is thus already shown to be a good indication of the risk involved in a particular transaction.

As far as bandwidth, performance, and computational and messaging costs are concerned, although the scale of the experiments we were able to carry out is small, it was clear that there would not be any issues either as a result of the JXTA peer-to-peer infrastructure, or as a result of the application itself. One point to note, however, is that JXTA allows various different communication modes, from full broadcast to direct point-to-point communication. Care needs to be taken in selecting the most appropriate mode, in order to avoid excessive and unnecessary messaging and bandwidth usage.

5.2 Composite experimental trials

Based on the data obtained from the experiment described in Section 4.2, we want to evaluate the degree of efficiency the market was characterised by, and ultimately how effective our system is in providing its users with indications of the risk involved in transactions. We attempt to examine this from different perspectives.

Our first observation concerns the overall trend of the market. Figure 5 shows the overall amount of wealth within the insurance market as a function of time, for two independent trials (the first involving seven insurers and 250 transactions, the second involving 24 insurers and 2300 transactions). We see a similar trend, whereby at first the market players are unsuccessful with their insurance offers and as a result lose more money than they make, however as time passes the trend is reversed, and they begin recovering and accumulating wealth. This could be attributed to the following reasons: First of all, as was discussed in Section 4.2.1, the sampled transactions that the insurers were exposed to on average involved less trustworthy peers than they expected. It turns out that most insurers were surprised by this, and it required some time for them to adjust their strategies and their ability to evaluate the feedback scores of the various players. This was further confirmed in the analysis of the questionnaires, as discussed below. Apart from that, it is an indication that with time, the market matured, the players were able to use the information available, and in time market growth was the predominant trend.

In Figure 6.a, we show the correlation between transaction risk (expressed as the percentage of the transaction amount that is statistically expected to be lost) and insurance offer placed for the transaction. This represents the price-quality correlation, and is an important indicator of market efficiency, as discussed in Section 4.2.3. We see in the figure that as time passes and the market matures, according to our interpretation, the scaled difference between risk and offers placed decreases, and there is stronger price-quality correlation, as one would expect from an efficient market. Note that it was not possible to measure the degree to which fraudulent transacting peers are singled out, since there were almost no repeat transactions in our database.

One metric that we were hoping would yield more conclusive results was the correlation between transaction risk and reward for insurer. As Figure 6.b shows (focusing on transactions with reasonable risk and excluding rates offered that were rejected as too high), there is a degree of correlation, but it is not as prominent as we would expect, yielding a correlation coefficient R value of 0.35. We believe that the main explanation of this is a general immaturity of the market, which was also evident from the questionnaires (see discussion below), and which could be explained by the following arguments:

- The constant availability of new transactions to place offers for acted against the competitive nature of the market.
- The number of market participants was thus not correctly matched to the availability of resources, which is known to affect the overall efficiency [Hr03].
- The subjects may not have been as motivated and focused as they would have if they actually had money at stake. Note that if we ignore the first third of the trial in the calculation of the correlation coefficient, the value of R increases to 0.4, whereas if we only focus on the last third of the trial the value is up to 0.51. This is further evidence of the market maturing.

The questionnaires filled in by the subjects of the experiment allow us to make some additional observations.

On average the subjects seemed to exhibit an unnaturally tolerant view of how trustworthy a transacting peer should be considered, based on their positive feedback percentages. As Table 4 shows, a considerable percentage of them would consider as trustworthy peers that were statistically expected to cheat in one out of four or five times. And yet, in their bidding strategies they would not compensate for this as the rates they requested were not even close to 20–25%. Indeed less than 15% of subjects claimed they were not often surprised by the transaction outcomes, since they failed to appreciate the degree of risk involved in transactions involving peers with such percentages. Of course the high degree of risk was a desired result of the data sampling bias discussed in Section 4.2.2. This particular behaviour was found to improve with time, however it goes a long way to exemplify how difficult it can be for the average buyer or seller to appreciate the implications with respect to risk of the feedback scores assigned to their counterparties.

This almost conveys the impression that, although they could clearly see the percentages, the insurers fixated on their experiences on systems such as eBay. Although our samples were “rigged” through our importance sampling process, the subjects treated it as they would choose ordinary eBay data, to their

Positive feedback	Extremely trust-worthy	Trust-worthy	Neutral	Untrust-worthy	Extremely untrust-worthy
95%	29%	71%			
90%		64%	36%		
80%		7%	57%	36%	
75%			29%	43%	29%

Table 4: Subject’s perception of trustworthiness of transacting peers based on their positive feedback percentages. Note that on average, the rate of positive feedback encountered on eBay is 99.8% [GS08]

chagrin. This is similar to a well-studied error in human decision making, that of *anchoring* a decision from a given starting point (see the work of Kahneman and Tversky [TK74]).

Along the same lines, we were surprised to find that about half of the users would consider the number of prior transactions (regardless of their outcome) to be an equally strong indication of trustworthiness as the percentage positive feedback received for such transactions. These observations led us to confirm a considerable amount of immaturity in the market, but in general there is compelling evidence that humans can be very bad in evaluating percentage probabilities, risks, and making decisions under uncertain conditions (again, the work of Kahneman and Tversky offers a wealth of examples [KST82, KT00]).

A positive sign, on the other hand, was the fact that almost 80% of subjects at some point decided they needed to change their bidding strategies, an indication that the market worked towards imposing discipline through better and more rational strategies, thus increasing overall efficiency, as is discussed with respect to Figure 5.

Another indicator of market efficiency, namely the availability of information, was probed in the questionnaire, and around 75% of subjects considered the information available regarding each transaction to be sufficient, whereas around 90% of subjects found the information that was available about other (competitive) insurers to be very useful.

Finally, almost without exception, all subjects appreciated the importance of the ability to insure transactions, and found the financial perspective of acting as an insurer very appealing.

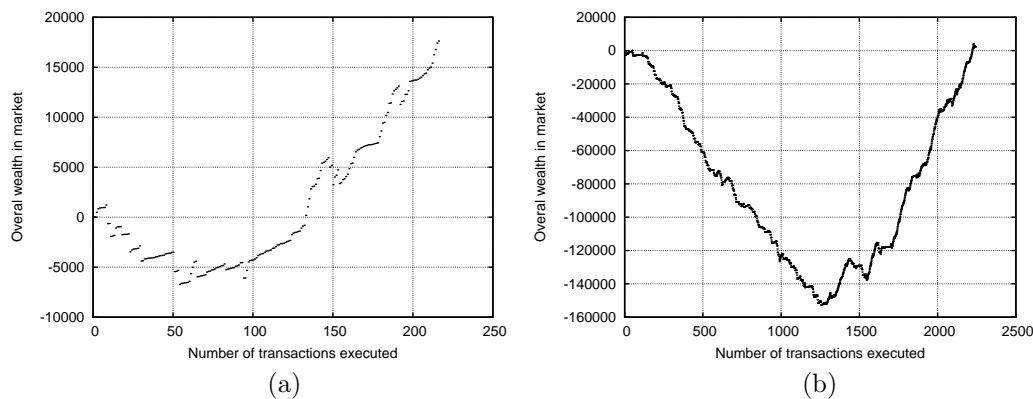


Figure 5: Overall wealth in market as a function of number of transactions executed, for two independent trials. (a) trial with 7 subjects and around 250 transactions; (b) Trial with 24 subjects and around 2300 transactions.

5.3 Insuring insurance transactions

As PTRIM assimilates properties of real world markets and the instruments they use for managing transaction risk, it inevitably also incorporates the need to consider the risk involved in these instruments, namely the insurance transactions themselves.

To mimic what happens in the real world, it could be possible to treat insurance transactions as regular transactions, and allow insurance offers to be placed for them. This means that the entire system would be recursively defined, closely resembling the way the global financial system securitises and repackages risk with all the benefits and (as we have recently witnessed [Eco08]) hazards this entails.

The point is that once we have insurance contracts, these are valuable by themselves, and they can be traded themselves, like banks do in the real world to spread/hedge their risks: they re-parcel their

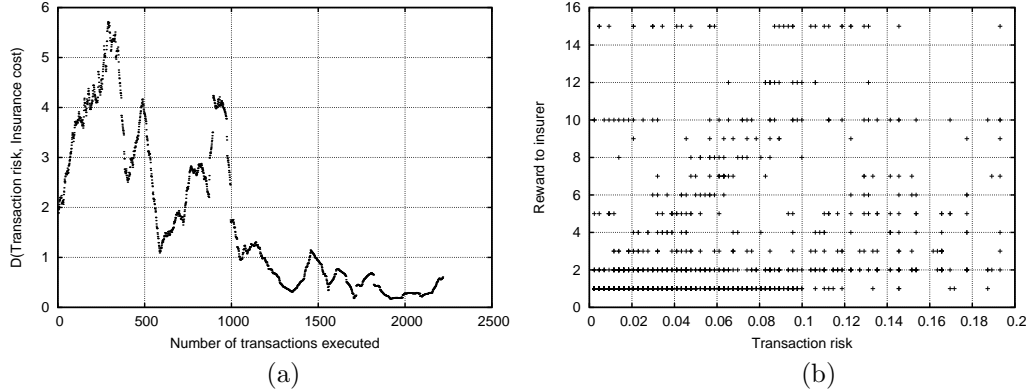


Figure 6: (a) Scaled difference between transaction risk and insurance offer, as a function of number of transactions executed, for a trial with 24 subjects and around 2300 transactions. (b) Correlation between transaction risk (expressed in values between 0 and 1) and reward to insurer, for a trial with 24 subjects and around 2300 transactions.

exposure and sell it. The major drawback of such a system, then, is not individual transaction default, but systemic breakdown [MLS08]. Indeed such systems are prone to systemic shocks exactly because after a couple of recursions, it is hard to know how much exposure there is. This is not a defect of PTRIM in particular, but a feature of credit markets in general; in spring 2008, following the collapse of investment bank Bear Stearns, major banks worldwide were discussing the development of a CDS clearing house [FT08]. More recently, the financial crisis that was sparked from the burst of the sub-prime mortgage bubble has thrown in sharp relief questions over regulation of financial markets. In what regards PTRIM, the “hole” represented by an insurer defaulting on its promise (the “counterparty risk”) is therefore something that must either be accepted, or mitigated by centralized mechanisms or some form of capital requirement guarantee. In general, it would be interesting to explore different regulatory scenarios by experimenting with the PTRIM system.

6 Future work

There are some open issues in our approach, that need to be further elaborated on.

6.1 Security

In our description we have not discussed the issue of security, which is very important owing to the nature of the application and the distributed network architecture. Given the considerable work that has been done in this area, our proposed system can be protected from most security threats by applying known state of the art solutions. Table 5 briefly summarises the main security attack categories for applications of this type, and potential countermeasures for each that have been proposed in the literature, and/or already implemented within other peer-to-peer systems and applications. In our prototype we have incorporated such measures to an extent, to provide authentication and preserve the confidentiality, privacy and integrity of data exchanged in the network.

In terms of peer privacy, a certain amount of information regarding past peer transactions is made available to the network, in the same way that past user history is publically available in online reputation management systems. This is similar to real life situations, where background information is made public when requesting insurance services, or stock exchange listed companies are required to provide data regarding their activities in order to participate in transactions.

6.2 Collusion resistance

Most peer-to-peer networks and reputation management systems are faced with the risk of peers colluding to form malicious collectives or cliques; when such collectives manage to manipulate prices, the market is no longer efficient. A number of approaches to render systems like ours more resistant to collusion have been proposed in the literature and are currently being considered.

Attack category	Countermeasure
Data Integrity Attacks (Both for data stored and data routed between nodes)	Public key infrastructures and encryption (see Cre- dence [WGS06], PeerTrust [XL04]); majority voting mechanisms (see PeerTrust [XL04]); trusted third parties (see EigenTrust [KSGM03], TrustMe[SL03]); other cryptographic algorithms/protocols, such as self-certifying data [CDG ⁺ 02], signed files (see Past [DR01]), information dispersal (see Publius [WAL00], and Mnemosyne [HR02]), secret sharing schemes (see [Moj07]).
Data Confidentiality Attacks (both for data stored and routed)	Public key infrastructures and encryption (see XRep [DDCdVP ⁺ 02], EigenTrust [KSGM03]); Secure rout- ing protocols [CDG ⁺ 02].
Data Availability Attacks	Passive / cache-based / proactive replication (see [CSW00, CS01, Wal02, SGG02], Tapestry[ZKJ01]) and trading of content (see FreeHaven [DFM01]).
Identification-Related Attacks (e.g. pseudospoofing, Sybill at- tack, ID-stealth, decoy and white- washing [WGS06], impersonation attacks [DHA03])	Trusted identity-certifying agents [Dou02]; public key infrastructures and certificates [WGS06].
Bootstrapping Phase Attacks (e.g. malicious node insertions in the network)	Majority quorum approaches [DHA03]; resource- based reputation approaches [DDCdVP ⁺ 02].

Table 5: Main categories of security attacks for peer-to-peer applications, possible countermeasures, and example systems that propose or adopt them. Note that some of the proposed solutions, such as public key infrastructures, are not purely decentralized.

In reference [KSGM03] the notion of pre-trusted peers is introduced, whereas IP clustering is proposed as a solution in reference [DDCdVP⁺02], but both of these approaches may not be practically applicable in our case.

Among the most interesting approaches, Jurka and Faltings [JF07] suggest and analyse an incentive-based collusion resistant approach for online reputation mechanisms, while Feldman et al. [FLSC04] propose the reciprocative decision function, and the use of a maxflow-based reputation management system to achieve optimal levels of cooperation. Zhang et al [ZG⁺04] work on Eigenvector-based reputation systems defining a metric they call amplification factor designed to make them more robust to collusion, and in a relevant subject Marti et al. [MGGM04] propose leveraging the trust associated with social links inherent within peer-to-peer networks for improved DHT routing.

7 Further Experiments and Simulation Studies

It would clearly be beneficial to carry out additional, and different, sets of trials and experiments and simulations of the PTRIM system, and on more wide-scale deployment.

More long-term trials would allow more feedback to be collected and utilised by the participating subjects. We also have considered different experimental set-ups that would allow more rigorous data collection and analysis.

One additional area of particular interest is the modelling of different peer behaviour patterns in simulation studies. For example it would be possible to simulate specific types of threats and attacks to the system, such as individual malicious peers, malicious collectives, spies, impersonation or Sybil attacks, etc.

In these studies, a number of malicious peers would be introduced in the network who would behave individually or collectively according to these threat models. For example individual malicious peers could transmit incorrect information, or malicious collectives could transmit perfect feedback for every other peer within their group. The system’s response to these attacks would then be studied to determine how resistant it is.

8 Conclusions

Our approach for managing the risk of distributed peer-to-peer transactions is based on the principles governing financial markets. PTRIM builds and utilises a transaction default market on top of a main transaction processing system, within which peers offer to underwrite the transaction risk for a slight increase in the transaction cost. The cost is determined through market-based mechanisms, and forms a way of identifying untrustworthy peers and perilous transactions. The risk of the transactions is effectively minimised while the transacting peers are alleviated from the need to collect, process and evaluate reputation information. At the same time our system showcases a way for members of a peer-to-peer network to capitalise on their market knowledge by participating in a new role, as transaction insurers.

The PTRIM approach can be considered as an alternative, but also as complementary to the use of distributed reputation management systems. We evaluated and discussed the performance of our approach through a methodology relevant to the one used for distributed reputation management systems. Our experimental findings suggest that our system functions in an efficient manner, and is able to support the transaction processing layer through the insurance offers it produces. At the same time our findings expose interesting observations regarding the behaviour of the subjects of our trials, and how this behaviour affects the market and its growth towards maturity.

We discussed the system's current limitations and planned future work. One additional concept we feel might be worth exploring concerns the fundamental "tension" between a centrally managed trust system, and the decentralized market-based approach we advocate. How does a system like eBay, for instance, reflect the value of the reputation management services they offer in its cost model, and how can one relate this to our market-based approach from a financial perspective? In theory, is our approach going to be cheaper or more expensive for the participants? This is an interesting open issue beyond the scope of the current work that we feel warrants some more in-depth analysis.

In the meantime, we believe that this work may provide an incentive for other researchers to utilise our concepts in different applications and cases.

A Acknowledgment

The authors wish to thank Athanasios Giannacopoulos, Eric Soderquist, Raphael Markellos, Angeliki Poulymenakou and Nick Nassuphis, for their helpful insights and assistance.

This work is implemented within the framework of the Reinforcement Programme of Human Research Manpower (PENED), and co-financed by National and Community Funds (25% from the Greek Ministry of Development-General Secretariat of Research and Technology and 75% from E.U.-European Social Fund).

References

- [Abe01] Karl Aberer. P-Grid: A self-organizing access structure for P2P information systems. *Lecture Notes in Computer Science*, 2172:179–194, 2001.
- [AD01] Karl Aberer and Zoran Despotovic. Managing trust in a peer-2-peer information system. In Henrique Paques, Ling Liu, and David Grossman, editors, *Proceedings of the Tenth International Conference on Information and Knowledge Management (CIKM01)*, pages 310–317. ACM Press, 2001.
- [ATS04] S. Androutsellis-Theotokis and D. Spinellis. A survey of peer-to-peer content distribution technologies. *ACM Computing Surveys*, 36(4):335–371, December 2004.
- [BB04] Sonja Buchegger and Jean-Yves Le Boudec. A robust reputation system for p2p and mobile ad-hoc networks. In *In Proceedings of the Second Workshop on the Economics of Peer-to-Peer Systems*, 2004.
- [BMS06] Michael R. Baye, John Morgan, and Patrick Scholten. Information, search, and price dispersion. Working Papers 2006-11, Indiana University, Kelley School of Business, Department of Business Economics and Public Policy, December 2006.
- [Bon04] AN Bonfim. *Understanding Credit Derivatives and Related Instruments*. Academic Press, December 2004.

- [BOW02] C Bluhm, L Overbeck, and C Wagner. *An Introduction to Credit Risk Modelling*. CRC Press, September 2002.
- [CDG⁺02] M Castro, P Druschel, A Ganesh, Rowstron A, and DS Wallach. Secure routing for structured peer-to-peer overlay networks. In *Proceedings of the 5th Usenix Symposium on Operating Systems*, Boston, MA, December 2002.
- [CS01] E Cohen and S Shenker. Optimal replication in random search networks. Preprint, 2001.
- [CSW00] I Clarke, O Sandberg, and B Wiley. Freenet: A distributed anonymous information storage and /etrieval system. In *Proceedings of the Workshop on Design Issues in Anonymity and Unobservability*, Berkeley, California, June 2000.
- [Dam02] A Damodaran. *Investment Valuation 2nd edition*. John Wiley & Sons, New York, 2002.
- [DDCdVP⁺02] E Damiani, S De Capitani di Vimercati, S Paraboschi, P Samarati, and F Violante. A reputation-based approach for choosing reliable resources in peer-to-peer networks. In *Proceedings of the In 9th ACM Conf. on Computer and Communications Security*, Washington DC, November 2002.
- [Del04] C. Dellarocas. *Information Society or Information Economy? A Combined Perspective on the Digital Era*, chapter Building Trust On-Line: The Design of Robust Reputation Mechanisms for Online Trading Communities. IDEA Book Publishing, Hershey, PA, 2004.
- [DFM01] R Dingledine, MJ Freedman, and D Molnar. *Peer-to-peer: Harnessing the power of disruptive technology*, chapter 1. A network of peers: Peer-to-peer models through the history of the Internet, pages 3–20. O’Reilly, first edition edition, March 2001.
- [DHA03] A Datta, M Hauswirth, and K Aberer. Beyond ‘web of trust’: Enabling p2p e-commerce. In *Proceedings of the IEEE International Conference on E-Commerce Technology (CEC’03)*, 2003.
- [Dou02] JR Douceur. The Sybill attack. In *Proceedings of the 1st International Workshop on Peer-to-Peer Systems (IPTPS ’02)*, MIT Faculty Club, Cambridge, MA, USA, March 2002.
- [DR01] P Druschel and A Rowstron. Past: A large-scale, persistent peer-to-peer storage utility. In *Proceedings of the Eighth Workshop on Hot Topics in Operating Systems*, May 2001.
- [DR03] C. Dellarocas and P. Resnick. Online reputation mechanisms: A roadmap for future research summary report. In *Proceedings of the First Interdisciplinary Symposium on Online Reputation Mechanisms*, April 26-27 2003.
- [Eco08] The Economist: Clearing the fog, April 17 2008.
- [Fam70] Eugene F. Fama. Efficient capital markets: A review of theory and empirical work. *The Journal of Finance*, 25(2):383–417, 1970.
- [FLSC04] M. Feldman, K. Lai, I. Stoica, and J. Chuang. Robust incentive techniques for peer-to-peer networks. In *Proceedings of the EC04 Conference*, May 2004.
- [FPCS04] M. Feldman, C. Papadimitriou, J. Chuang, and I. Stoica. Free-riding and whitewashing in peer-to-peer systems. In *Proceedings of the 3rd Annual Workshop on Economics and Information Security (WEIS2004)*, May 2004.
- [FT008] Financial Times: Banks eye CDS clearing house, April 18 2008.
- [FTW05] Ming Fan, Yong Tan, and Andrew B. Whinston. Evaluation and design of online cooperative feedback mechanisms for reputation management. *IEEE Trans. on Knowl. and Data Eng.*, 17(2):244–254, 2005.
- [GJA03] M Gupta, P Judge, and M Ammar. A reputation system for peer-to-peer networks. In *Proceedings of the NOSSDAV’03 Conference*, Monterey, CA, June 1-3 2003.

- [GL03] P. Glasserman and J. Li. Importance sampling for portfolio credit risk, 2003.
- [GLFH] Carmelita Gorg, Eugen Lamers, Oliver Fuss, and Poul Heegaard. Rare event simulation.
- [GS06] Dawn G. Gregg and Judy E. Scott. The role of reputation systems in reducing on-line auction fraud. *Int. J. Electron. Commerce*, 10(3):95–120, 2006.
- [GS08] Dawn G. Gregg and Judy E. Scott. A typology of complaints about ebay sellers. *Commun. ACM*, 51(4):69–74, 2008.
- [GT08] Bezalel Gavish and Christopher L. Tucci. Reducing internet auction fraud. *Commun. ACM*, 51(5):89–97, 2008.
- [Hei95] Philip Heidelberger. Fast simulation of rare events in queueing and reliability models. *ACM Trans. Model. Comput. Simul.*, 5(1):43–85, 1995.
- [HR02] S Hand and T Roscoe. Mnemosyne: Peer-to-peer steganographic storage. In *Proceedings of the 1st International Workshop on Peer-to-Peer Systems (IPTPS '02)*, MIT Faculty Club, Cambridge, MA, USA, March 2002.
- [Hr03] Julia Hring. Different prices for identical products? : Market efficiency and the virtual location in b2c e-commerce. ZEW Discussion Papers 03-68, ZEW - Zentrum für Europäische Wirtschaftsforschung / Center for European Economic Research, 2003.
- [HS07] L Huaizhi and M Singhal. Trust management in distributed systems. *IEEE Computer*, 40(2):45–53, February 2007.
- [Hul08] J. C. Hull. *Options, Futures and Other Derivatives*. Prentice-Hall, 2008.
- [JF07] Radu Jurka and Boi Faltings. Collusion-resistant, incentives-compatible feedback payments. In *Proceedings of the EC07 conference*, San Diego, California, June 2007.
- [Jxt07] The project JXTA web site. <http://www.jxta.org>, Accessed on-line 2007.
- [Kal94] KK Kalita. Measuring product market efficiency: A new methodology. *Marketing Letters*, 5(1), January 1994.
- [KR05] S Kinsella and S Raghavendra. Case study of price dispersion in online markets. Technical report, Digital Enterprise Research Institute, Ireland. <http://library.deri.ie>, accessed online 2008, June 2005.
- [KSGM03] Sepandar D. Kamvar, Mario T. Schlosser, and Hector Garcia-Molina. The EigenTrust algorithm for reputation management in p2p networks. In *Proceedings of the twelfth international conference on World Wide Web*, pages 640–651, New York, 2003. ACM Press.
- [KSS04] Apostolos Kritikopoulos, Martha Sideri, and Konstantinos Stroggylos. Crawlwave: A distributed crawler. In *Proceedings of the 3rd Hellenic Conference on Artificial Intelligence (SETN '04)*, May 2004.
- [KST82] D. Kahneman, P. Slovic, and A. Tversky, editors. *Judgement under uncertainty: Heuristics and biases*. Cambridge University Press, New York, NY, 1982.
- [KT00] D. Kahneman and A. Tversky, editors. *Choices, values and frames*. Cambridge University Press, New York, NY, 2000.
- [KWO07] Paul Krugman, Robin Wells, and Martha Olney. *Essentials of Economics*. Worth Publishers, 2007.
- [LG02] Z Lee and S Gosain. A longitudinal price comparison for music cds in electronic and brick-and-mortar markets: Pricing strategies in emergent electronic commerce. *Journal of Business Strategies*, 19(1), 2002.
- [MGGM04] S. Marti, P. Ganesan, and H. Garcia-Molina. DHT routing using social links. In *In Proceedings of the 3rd International Workshop on Peer-to-Peer Systems (IPTPS 2004)*., 2004.

- [MLS08] RM May, SA Levin, and G Sugihara. Complex systems: Ecology for bankers. *Nature*, 451(21):893–895–55, Feb 2008.
- [Moj07] The MojoNation web site. <http://www.mojonation.net>, Accessed on-line 2007.
- [RZ02] P Resnick and R Zeckhauser. *Advances in Applied Microeconomics*, volume 11, chapter Trust Among Strangers in Internet Transactions: Empirical Analysis of eBay’s Reputation System. The Economics of the Internet and e-Commerce. JAI Press, Amsterdam, 2002.
- [SGG02] S Saroiu, PK Gummadi, and SD Gribble. Exploring the design space of distributed peer-to-peer systems: Comparing the web, TRIAD and Chord/CFS. In *Proceedings of the 1st International Workshop on Peer-to-Peer Systems (IPTPS '02)*, MIT Faculty Club, Cambridge, MA, USA, March 2002.
- [Shi05] Robert Shiller. *Irrational Exuberance*. Princeton University Press, 2nd edition, 2005.
- [SL03] Aameek Singh and Ling Liu. Trustme: Anonymous management of trust relationships in decentralized p2p systems. In *Proceedings of the IEEE Intl. Conf. on Peer-to-Peer Computing*, September 2003.
- [Sti57] George J. Stigler. Perfect competition, historically contempated. *The Journal of Political Economy*, 65(1):1–17, Feb 1957.
- [SW05] K Steinmetz and K Wehrle, editors. *Peer-to-Peer Systems and Applications*. Springer Lecture Notes in Computer Science 3485, 2005.
- [TK74] Amos Tversky and Daniel Kahneman. Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157):1124–1131, 1974.
- [WAL00] M Waldman, Rubin AD, and Cranor LF. Publius: A robust, tamper-evident, censorship-resistant web publishing system. In *Proceedings of the 9th USENIX Security Symposium*, August 2000.
- [Wal02] DS Wallach. A survey of peer-to-peer security issues. In *International Symposium on Software Security*, Tokyo, Japan, November 2002.
- [WGS05] K Walsh and E Gun Sirer. Fighting peer-to-peer spam and decoys with object reputation. In *Proceedings of the SIGCOMM'05 Conference Workshops*, Philadelphia, PA, August 2005.
- [WGS06] K Walsh and Emin Gon Sirer. Experience with an object reputation system for peer-to-peer flesharing. In *Proceedings of the Symposium on Networked System Design and Implementation*, San Jose, California, may 2006.
- [XL04] L Xiong and L Liu. Peertrust: Supporting reputation-based trust for peer-to-peer electronic communities. *IEEE Transactions on Knowledge and Data Engineering*, 16(16), July 2004.
- [ZG⁺04] H Zhang, A Goel, et al. Improving eigenvector-based reputation systems against collusion. Technical report, Stanford University, Workshop on Algorithms and Models for the Web Graph (WAW), October 2004.
- [ZKJ01] BY Zhao, J Kubiawicz, and AD Joseph. Tapestry: An infrastructure for fault-tolerant wide-area location and routing. Technical Report UCB/CSD-01-1141, Computer Science Division, University of California, Berkeley, 94720, April 2001.