Trust as a Tradable Commodity: A Foundation for Safe Electronic Marketplaces[*]

Reid Kerr and Robin Cohen
David R. Cheriton School of Computer Science
University of Waterloo, Waterloo, Ontario, Canada

Abstract

In large electronic marketplaces populated by buying and selling agents, it is difficult to judge trustworthiness. A variety of systems have been proposed to help traders to find trustworthy partners by learning to discount or disregard disreputable parties. In this paper, we present a novel model for providing safe electronic marketplaces: Commodity Trunits, a system that considers trust as a tradable commodity. In this system, sellers require units of trust (*trunits*) in order to participate in transactions, and risk losing trunits if they act dishonestly. Sellers can purchase trunits when needed, and sell excess quantities. We demonstrate that under Commodity Trunits, rational sellers will choose to be honest, since this is the profit maximizing strategy. We also show that Commodity Trunits provides protection from a number of vulnerabilities common in existing trust and reputation systems, e.g., the important *exit problem*, where sellers can cheat without fear of repercussions if they intend to leave the market. We then present a simulation that validates the system by demonstrating that a market operator can manage the trunit marketplace to ensure sustainability. We conclude with a discussion of the value of Commodity Trunits as a method for promoting trust in electronic marketplaces.

---

[*] This is the pre-peer reviewed version of the following article: (Trust as a Tradable Commodity: A Foundation for Safe Electronic Marketplaces, Reid Kerr and Robin Cohen, Computational Intelligence, Volume 26, Number 2, Copyright 2010, Wiley-Blackwell), which has been published in final form at http://www3.interscience.wiley.com/journal/123387568/abstract.

1 Introduction

An important challenge in the development of effective agent-oriented electronic marketplaces is to model the trustworthiness of the selling agents, in order to provide buying agents with a basis for avoiding disreputable sellers and for selecting the most promising business partners. Much research has been conducted into the notion of trust in a variety of fields, including sociology (e.g., Luhmann, 1979), psychology (e.g., Rotter, 1971), economics (e.g., Williamson, 1993), management (e.g., Mayer et al., 1995), etc; cross-disciplinary surveys include (Gambetta, 1990) and (Rousseau et al., 1998). Here, we limit our focus to the issue of trust in multi-agent scenarios, and work applicable to these scenarios. A variety of such approaches have been proposed for modeling trust and reputation, including the use of reinforcement learning for individual buying agents to determine which sellers are most likely to provide the demanded quality of goods (e.g., Tran and Cohen, 2004), the use of probabilistic reasoning to detect and then disregard disreputable sellers e.g., (e.g., Jøsang and Ismail, 2002) and the use of social networks of advice from other buyers, discounting the advice of less trustworthy agents (e.g., Yu and Singh, 2002; Teacy et al., 2006). Unfortunately, the trust and reputation models developed to date still have inherent vulnerabilities that can be exploited by dishonest selling agents. For example, one important challenge that arises is the *exit problem*, where agents planning to leave the marketplace can cheat, with little regard for the consequences.

In this paper, we advocate the design of electronic marketplaces where buyers feel protected, because sellers are inclined to be honest. Our research builds on the Trunits model for trust management in electronic marketplaces (Kerr and Cohen, 2006), which seeks to ensure the honesty of sellers. Under this model, trust is treated as units that sellers must possess in order to engage in transactions, and which sellers risk losing if they fail to satisfy their buyers. This model addresses several vulnerabilities common to other trust and reputation systems.

We develop an extension to this model known as Commodity Trunits, where trunits can be bought and sold within the marketplace. It can be shown that under this system, the exit problem is also addressed. This is demonstrated through theoretical discussion of the conditions that must be met in the marketplace to guarantee that rational sellers will not cheat buying agents, and through simulations that verify that these conditions can in fact be assured within e-marketplaces. It becomes more profitable for selling agents to be honest, even if they plan to leave the marketplace.

The result is a system for addressing trust in electronic marketplaces that constitutes an important step towards providing safety for buying agents.

2 The Trunits Model

In this section, we present a novel model for trust in multiagent systems, Trunits (Kerr and Cohen, 2006). Trunits differs from most existing trust models in that it attempts not to predict behaviour, but rather to ensure good behaviour by making honesty the most profitable strategy.

Trunits is inspired by the concept of money. Before the advent of money, goods and services were exchanged by bartering. This placed several limitations on trade; of particular interest here, traders had to interact directly, in order to exchange goods. A primary role of money is to overcome these limitations. Money is an abstract ‘substance’, representing quantities of value. Money flows in a transaction, mirroring the flow of value in a barter transaction: the value of the money stands in for the value of a good. Money frees traders from the requirement for a direct two-way relationship where goods move in both directions—value gained from one trader can be ‘spent’ with another (Newlyn, 1971).

In the large electronic marketplaces with which many researchers are concerned, buyers and sellers rarely meet one another directly; eBay is a typical example. The absence of direct relationships between traders in these scenarios prevents trust from forming naturally. Since we seek to overcome the requirement for a direct relationship—to allow trust gained from one trader to be ‘spent’ with another—it seems natural to consider the use of abstract trust units, or *trunits*, to play the same basic role in which money has been so successful.

Trunits was developed with a focus on *advertised price* marketplaces, in which sellers offer products for sale, specifying the nature of the good and the price at which it is offered. Buyers select goods that they wish to purchase, accepting sellers’ offers. When a seller’s offer is accepted by a buyer, the seller is obliged to fulfill the terms of the offer. (This may be viewed as a promise, or more formally as a contract.) We consider a seller’s failure to fulfill his obligation as *cheating*, or *dishonesty*. Cheating would include both the case where a delivered good does not conform to the seller’s promise, and the case where the seller fails to deliver a good at all. This perspective follows that taken by other authors (e.g., Teacy et al., 2006). Conversely, an *honest* transaction is one where a seller fulfills all terms of his commitment.

The movement of money in a transaction mirrors that of the flow of value in a direct bartering situation.

Similarly, the flow of trunits should mirror that of trust in a direct trading relationship. The two 'flows', however, are fundamentally different in nature. The flow of value in a transaction is an exchange process, wherein each trader receives something of value, in exchange for providing the other with something of value. In contrast, we see the 'flow' of trust in a transaction as a *risk* process, focusing on the buyer's trust of the seller as the primary issue:

- Before a buyer will purchase something from a seller, the buyer must have sufficient trust in the seller. The degree of trust required is dependent on a number of factors; the price of the item is likely a major one.
- After purchasing the good, the buyer will evaluate it, relative to her expectations.
  - If the good met her expectations (i.e., it was at least as good as was advertised by the seller), then the seller is likely to gain more of her trust.
  - If the good did not meet her expectations, then the seller is likely to lose some of her trust.

Based on this view, we suggest a model that makes use of abstract units of trust, and in which trust of the seller is not tied to a specific buyer:

- The seller has some quantity of trunits, representing all of the trust gained from all buyers to date. For a buyer to consider purchasing something from a seller, the seller must possess a sufficient degree of trust, i.e., must hold sufficient trunits. The required number of trunits is tied to the price of the good.
- After purchasing the good, the buyer will evaluate it, relative to her expectations (i.e., to the claims of the advertisement).
  - If the good met her expectations, then the seller gains some additional quantity of trunits.
  - If the good did not meet her expectations, then the seller loses some quantity of trunits.

As a seller executes honest transactions, his trunit balance grows, allowing future profitable transactions.[1] In contrast, dishonest sales curtail future transactions. This provides the fundamental incentive for honesty. The number of trunits gained is proportional to the size of the sale. Honest execution of small transactions

---

[1] Note that this relies on buyers being honest in their feedback about sellers. We assume a parallel system is used to ensure buyer honesty. This issue is discussed further in Section 7.

will allow a seller to continue making small sales, and to grow his sales volume, but will not allow him to immediately jump to disproportionately large sales for which he has not demonstrated trustworthiness.

Trunits seeks only to ensure the honesty of sellers, without attempting to address the behaviour of buyers. This focus is rooted in the greater vulnerability of buyers given the common policy that buyers pay before shipment by the seller. Note that existing proposals (e.g., Dellarocas, 2002; Jurca and Faltings, 2003) also focus on solving one of these two substantially different issues, without significantly addressing the other. Here, it is intended that Trunits would be used in parallel with a system to protect sellers against potential dishonesty on the part of buyers.

3 The Basic Trunits Mechanism

With this general model as a foundation, we formalize a basic mechanism that builds on the model. This mechanism uses the notion of a *market operator*, who is responsible for administering the marketplace. The operator maintains accounts and trunit balances for each trader, holds trunits in escrow during the course of a transaction, collects feedback from buyers and updates trunit balances, etc. We focus on marketplaces where the market operator is an identifiable entity (e.g., eBay Inc. operating the eBay marketplace); where such an entity is identifiable, it is considered to be a trusted third party.[2]

3.1 Mechanism overview

When an agent wishes to make a sale, we require him to put up a quantity of trunits to ‘cover’ the sale. These trunits represent the trust that the seller is risking by engaging in a transaction. We require that the number of trunits risked be directly tied to the value of the transaction, using the formula:

[1] $$\tau = \frac{V}{r}$$

where τ is the number of trunits, V is the value (selling price) of the transaction, and r is the required *risk ratio*, a (positive) parameter set by the market operator. The trunits are put into escrow with the market

---

[2] Trunits might be implemented as a decentralized (e.g., peer-to-peer) system. Under such circumstances, the activities of the market operator must still be performed; we may view the ‘market operator’ as the implementer of the system, or simply as the enforcement of rules by the system.

operator, pending completion of the transaction.

Upon completion, if the buyer rates the transaction as unsatisfactory, then the seller loses the τ trunits placed in escrow. If, on the other hand, the buyer rates the transaction as satisfactory, then the τ trunits are returned to the seller, along with some additional quantity of trunits related to the value of the transaction, for a total of:

$$(1+p)\tau = (1+p)\frac{V}{r} \tag{2}$$

where p is a *premium* or *reward* of additional trust for acting in an honest manner, a (positive) parameter set by the market operator. (It is suggested that p be less than 1, in order for trust to be harder to earn than it is to lose (as suggested by Tran and Cohen (2002, 2004)). Where $p < 1$, the $p\tau$ gained for honesty is less than the τ lost for dishonesty.) In the basic mechanism presented here, the same values of r and p are used for all traders and all transactions.

From a buyer's perspective, no evaluation or computation is required prior to purchasing to determine if a seller is trustworthy—if the seller possesses enough trunits for a transaction, then *by definition*, she is trustworthy *for that transaction*. Here, we consider an agent that can be expected to behave honestly to be *trustworthy*.[3] The market operator, who stores the trunits balances, will not allow a transaction to be executed unless the seller has sufficient trunits; if the good is permitted by the operator to be sold, the buyer is assured that the seller has secured it with sufficient trunits. From a seller's perspective, honesty results in a growing trunit balance and the ability to engage in more sales in the future, while dishonesty will reduce the potential for future sales. The intention is that trunits themselves will be valued, deriving from the fact that they enable profitable transactions.

3.1.1 An Example

To illustrate the mechanics of the system, consider an example where a seller has 20 trunits, the required ratio $r = 5$, and the reward ratio $p = 0.2$. By Formula 1, the maximum transaction value that the seller can cover is $V = r\tau = 20 \times 5 = \$100$.

Suppose now that the seller engages in a sale for a price of $50:

---

[3] The expectation of honesty is based on the incentive for honesty, as explained below.

- By Formula 1, the seller must place $\frac{V}{r} = \frac{50}{5} = 10$ trunits in escrow.

- If the buyer is unsatisfied, the seller loses the 10 trunits in escrow. She now has 10 trunits remaining, meaning that she can now cover transactions with a maximum value of \$50.

- If the seller is satisfied, then the seller's trunits are returned to her, and she receives an additional $p\tau = 0.2 \times 10 = 2$ trunits. She now holds a total of 22 trunits, and can cover transactions with a maximum value of $22 \times 5 = \$110$.

3.2 Why a buyer can trust in the system: The incentive for honesty

As noted by Dellarocas (2003a), "sellers care about buyer feedback primarily to the extent that they believe it might affect their future profits." As discussed above, honesty increases sellers' trunit balances, and hence their possible future sales, while dishonesty reduces both. This provides *some* degree of incentive for honesty; the question must be asked, is it a large enough incentive? Will honesty be the most profitable policy for sellers, motivating them to make honest choices, and thus allow buyers to trust in them?

The incentive for honesty provided by Basic Trunits can be demonstrated by considering any single arbitrary transaction.[4] This transaction may coincide or overlap with other transactions executed by the same seller.

A buyer has agreed to purchase a good (for price V, where $V = r\tau$). The τ trunits used to cover the sale are not necessarily the seller's entire balance, but any arbitrary portion of his balance. The seller now has a choice to make: cheat, or be honest? We assume that the seller makes his choice based solely on his own economic motivation—in particular, he is acting alone. Further, we assume that, should the seller engage in an honest sale, he will have the opportunity to engage in future sales using the trunits returned to him. This condition is necessary for honesty to be an attractive choice (i.e., for the trunits received after being honest to have value for future transactions); we believe that this is a reasonable assumption, given the assumed large market size, and the fact that all sellers with sufficient trunits are considered equally trustworthy under the model.

If the seller cheats, the maximum profit he can realize from the sale is $r\tau$—where he fails to ship the good

---

[4] Note that this proof, considering any arbitrary *single* transaction, necessarily ignores issues of collusion—the seller is assumed to be acting alone. The solution to certain forms of collusion is presented later in this paper.

at all, so he incurs no cost. By cheating, he also loses his τ trunits. If the seller is honest, he ships the product and realizes a smaller profit, $(1-c)r\tau$, where c is the cost incurred in selling the item, as a fraction of the sale price. However, after the sale he has $(1+p)\tau$ trunits returned to him, which can be used to engage in an additional sale.[5] Since the seller decides independently for each transaction whether to cheat or be honest, he can cheat this second buyer—in fact, he need not even have items in inventory. Cheating this second buyer returns a profit of $r\tau$, and causes the buyer to lose τ trunits, leaving him with $p\tau$ remaining. The total profit earned (so far) is $(1-c)r\tau + r\tau = (2-c)r\tau > r\tau$ (since $c < 1$). Thus, a rational seller will be honest on the current transaction, since he can make more money by waiting to cheat (and earning honest profits in the interim) than by cheating immediately and destroying the trunits used to cover the transaction.

This is not meant to imply that the rational seller will cheat on the subsequent transaction. Rather, on any given transaction, the seller can make more money in total by delaying his cheating than by cheating immediately. It follows that the rational seller will never cheat, since it always makes sense to engage in 'just one more' honest transaction. Thus, the value that one might earn by cheating on the subsequent transaction is actually a lower bound on the profit from future honest transactions using those trunits. We note this lower bound here because it is direct, and because it would be obvious to any rational seller, even those whose computational capacity is limited.

This examination of the incentive provided by Basic Trunits has been necessarily brief. More detailed proofs are provided by Kerr and Cohen (2007a,b).

3.3 Other properties

Other noteworthy properties of the Basic Trunits mechanism include:

- The mechanism is very simple computationally, particularly for buyers, who do not need to perform any calculations to determine a seller's trustworthiness.
- The storage requirements for the mechanism are minimal, since a single trunit balance is kept for each seller, and a single record for each transaction in progress.
- In many systems, it is clear that trust has an impact on profitability (since being trustworthy allows

---

[5] To simplify discussion here, we assume that this second transaction has the same value as the first.

> more products to be sold, and/or fetches higher prices for products), but it is very difficult to measure that impact. This complicates many issues, including both strategic decisions by sellers, and buyers' interpretation of sellers' histories/reputations. The basic Trunits mechanism allows the value of trust to be determined with some precision, serving as the basis for rational decision making.

Basic Trunits has many desirable properties, but also presents certain limitations. In the next section, we present an extended model that directly addresses some of these limitations.

4 Commodity Trunits

4.1 Motivation

Kerr and Cohen (2006) identified a number of common vulnerabilities in trust and reputation systems, vulnerabilities that could be used by a dishonest trader to overcome the security provided by a system. Basic Trunits provides protection against certain key vulnerabilities, but is still exposed to others. We begin with a discussion of the problems that Basic Trunits does address, followed by a description of certain vulnerabilities that remain unaddressed. This motivates our development of Commodity Trunits as a model that retains the protections offered by Basic Trunits, while resolving other important concerns.

Reputation Lag: Under many systems, there is a delay between the time when a buyer agrees to a purchase, and when the buyer's feedback is ultimately available to other buyers (e.g., reflected in a reputation score). During this lag period, a seller might cheat a virtually unlimited number of buyers before his reputation is updated to reflect his dishonesty.

The Trunits mechanism deals directly with this vulnerability by compelling the seller to place trunits in escrow to cover transactions, forcing him to wait until the trunits have been returned before he can use them in another transaction. The Trunits mechanism regulates the rate at which transactions can occur: if the seller holds τ trunits, then the maximum value of transactions he can engage in during one unit of time is $r\tau$, regardless of timing or circumstances. In effect, the Trunits mechanism prevents the use of the 'same trust' to support multiple simultaneous transactions.

Value Imbalance: Under some systems, the weight of feedback is not related to the size of the transaction. This may allow a buyer to build reputation through a number of honest, small-value sales, then use this good reputation to cheat on sales of much larger value.

The Trunits mechanism deals directly with this vulnerability by basing both the quantity of trunits required to cover a transaction, and the size of the reward, on the value of the transaction.

Re-entry: In the large electronic markets we consider here, identity typically cannot be established. Under such circumstances, an agent who has earned a poor reputation can shed this impediment simply by creating a new user account.

Friedman and Resnick (2001) suggest a solution to dealing with the re-entry problem: new entrants to the market should incur a cost, such that the cost of re-entry exceeds the benefit. Dellarocas (2003b) establishes that in marketplaces with binary feedback mechanisms, the policy of 'optimal social efficiency' is one where new users begin with the worst possible reputation (i.e., the same as very disreputable users).

Under Basic Trunits, new and maximally disreputable users are treated the same, providing no incentive for re-entry, and consistent with the optimal policy.[6]

A new problem: 'Surplus trust' The Trunits mechanism, as described here, suffers from a potential problem: a seller may accumulate trunits beyond what is required to cover their regular transactions. Such surplus trunits could be used to cheat buyers, without having a negative impact on the regular transactions.

For example, consider a seller who has a fixed production capacity, so he can only sell five items per week. Assume that the seller has enough trunits to cover sales at this rate. Also assume that the trust reward $p = 0.2$. For each honest transaction, the seller will get back the trunits used to cover that transaction, plus a 20% reward. The trunits that were returned are all that is required to cover his future sales as well, since his production capacity is fixed; this means that the 20% reward is strictly in excess of his needs. The seller may now use these surplus trunits to cheat buyers; as long as he does not spend his original trunits on dishonest transactions, he will continue to be able to sell his entire production without impediment.

---

[6] Under Trunits, a zero trunit balance does not equate to disreputability. A seller could have zero trunits available, for example, because he is a new seller, or because he has lost his trunits by cheating. Trunits is focused on making honesty the most profitable choice for self-interested traders.

Exit: The analysis provided in Section 3.2 shows that sellers will be honest. It must be noted, however, that this analysis is valid only when the seller actually has inventory. If the seller has no items to sell (or more precisely, will *never* have more items to sell), then he cannot defer cheating until the next transaction—for a transaction to occur, it must be a dishonest one. This is an instance of the 'exit problem', which afflicts most models of trust/reputation.

There was an additional concern about using Basic Trunits in electronic marketplaces: it was unspecified how an agent obtains an initial quantity of trunits with which to begin trading. Further, at any given moment, a seller's trading volume may be constrained by his available trunits, even if he would fulfill every transaction honestly. Motivated by the desire to address these issues, it became clear that a market-based approach might be desirable. Our proposed extention to Basic Trunits, Commodity Trunits, allows agents to buy and sell units of trust. It is clear that, if sellers can buy trunits, the question of how initial quantities are obtained is resolved. (If trunits are unavailable from other sellers (e.g., at market startup), they are created and sold by the market operator. This issue is discussed in further detail in Section 5.2.1).[7] Similarly, the issue of restraint of trade is addressed. Under Basic Trunits, an agent's sales are limited by his available trunits; under Commodity Trunits, an agent can simply purchase trunits when more are required to cover sales. Moreover, allowing trust to be traded provides effective protection against the surplus trust and exit problems.

In the following sections, we verify the safety of this approach, and then discuss its implications for the system.

4.2 Safety

Commodity Trunits allows the sale of trunits. To ensure that this is safe, some conditions must be established. Here, we mean that a market is safe for buyers if they are protected from harm due to the actions of dishonest (rational) sellers.

We denote the selling price of a trunit as b. In this section, we assume that all trunits are bought and sold

---

[7] Another approach might be the use of secured trunit loans. This alternative is just a special case of our proposal, however, where trunits are purchased, then sold back at the same price.

at this price; we relax this constraint later in the paper. For a seller to wish to buy a trunit, he must expect to profit from the purchase (i.e., to increase his net profit from sales). For the seller to make a rational purchase decision, he must understand both the purchase price, and the expected change in profit.

To furnish a good for sale, the seller will incur some cost. We assume that for a given seller, the cost c is a fraction of the selling price, and is constant for all his sales. Each time a quantity of trunits τ is used to cover an honest transaction, the profit on that transaction will be $(1-c)V = (1-c)r\tau$. Upon completion of the transaction, the trunit balance will grow by a factor of p, meaning that $(1+p)\tau$ trunits are available for the next transaction. Thus, the profit from a sequence of h honest sales (where h is referred to as the *horizon*), with initial trunit balance τ_0 can be represented as a geometric series:

$$[3] \qquad P_{S_h} = \sum_{i=0}^{h-1}(1-c)r(1+p)^i\tau_0 = (1-c)r\tau_0\left(\frac{(1+p)^h-1}{p}\right)$$

At the end of the sequence of honest sales, however, the seller still possesses trunits, which have value. To determine a lower bound on the value of the trunits, we base our calculation on the assumption that the seller uses all left-over trunits to cheat at the end of the sequence. This serves as a lower bound because the mechanism permits the seller to cheat at will (if he is willing to sacrifice the trunits), so he can assuredly achieve at least this level of profit. Thus, we modify Formula 3 to incorporate this:

$$[4] \qquad P_{SC_h} = (1-c)r\tau_0\left(\frac{(1+p)^h-1}{p}\right) + r(1+p)^h\tau_0$$

(This evaluation method applies in more general cases as well. For example if a seller were to 'split' a quantity of trunits and use them to secure separate purchases, their value can be determined by evaluating each portion separately.)

Let P_{τ_h} be the profit from one trunit (i.e., P_{SC_h} when $\tau_0 = 1$). For a rational agent to purchase a trunit, then, his anticipated profit from the trunit must exceed the purchase price, i.e., $b < P_{\tau_h}$. A seller can use any sum of trunits τ immediately to cheat, yielding a profit of $r\tau$ (because, by Formula 1, a quantity of τ trunits can be used to cover a sale of value $V = r\tau$). For the sale of trunits to be safe, it must be the case that it costs more to buy them than can be earned by cheating with them; otherwise, the seller could profit by buying trunits, and then cheating buyers with them immediately. This requires that $b\tau > r\tau$, or $b > r$. Together, this requires that $r < b < P_{\tau_h}$.

For example, consider a marketplace in which the required ratio $r = 5$, the reward for honesty $p = 0.2$, the seller's cost to furnish goods $c = 0.5$, and the seller's horizon $h = 3$ transactions[8] . Suppose the seller has the opportunity to engage in a sale with a price of 50, thus requiring 10 trunits (by Formula 1, $\tau = V/r = 50/5 = 10$). Unfortunately, the seller has no trunits available. If the seller acquired enough trunits to execute the sale, she might do so honestly or dishonestly.

If the seller were to cheat (maximally), her profit from the sale (excluding the cost of acquiring the trunits) would be the selling price $V = 50$. For the same 10 trunits, using Formula 4 we can determine the anticipated profit to the seller if she were to purchase them and engage in honest sales to her horizon:

$$P_{SC_h} = (1 - 0.5)5(10)\left(\frac{(1+0.2)^3 - 1}{0.2}\right) + 5(1+0.2)^3 = 177.4$$

Now, consider b, the price paid to purchase trunits. If $b < r$, say $b = 4$, then the seller can purchase the required trunits for $b\tau = 4(10) = 40$. This would allow the seller to earn a profit of 10 by purchasing trunits for 40, then immediately cheating on a transaction and receiving 50. By comparison, if $b > r$, say 6, then the seller can purchase the required trunits for $b\tau = 6(10) = 60$. In this case, the seller loses money by purchasing trunits to cheat, paying 60 but only taking in 50. When $b > r$, the seller will purchase trunits only if they can be used to engage in an honest transaction.

But will the seller purchase the trunits at all? She will only be interested in purchasing the trunits if they cost less than her anticipated profit over her horizon—177.4 in this case, as calculated above. Thus, when $b\tau < P_{SC_h}$, i.e., $b(10) < 177.4$, or $b < 17.74 = P_{\tau_h}$, purchasing trunits is attractive to the seller.

Consider now a seller who has the opportunity to engage in a sale of value V, making use of τ trunits. With these trunits, there are now three choices:

1. Cheat, and realize profit of $r\tau$, while losing the τ trunits.

---

[8] This discussion may give the impression that honesty is only advantageous if the seller plans a sizeable sequence of honest transactions. As will be clarified below, the incentive for honesty exists even for horizons of a single transaction.

2. Sell the trunits, and realize a profit of $b\tau$, while relinquishing the τ trunits.

3. Act honestly, earning $(1-c)r\tau$, keeping the τ trunits and gaining an additional $p\tau$ trunits, for a total of $(1+p)\tau$. These trunits can then be used in one of three ways: cheat on future transaction(s), sell the trunits, or engage in future honest transaction(s). To compare the profitability of the original three actions, however, we need only consider one of these options: since trunits can be sold at price b, the seller can earn at least this much with the remaining trunits. The $(1+p)\tau$ trunits will fetch $b(1+p)\tau$ on the market. The total profit in this case, then, is $(1-c)r\tau + b(1+p)\tau$.

Since $b > r$, $b\tau > r\tau$; selling the trunits is more profitable than cheating. Since $c \leqslant 1$ and p is positive, $(1-c)r\tau + b(1+p)\tau \geqslant b(1+p)\tau \geqslant b\tau > r\tau$, so engaging in an honest sale is also more profitable than cheating. Cheating, then, is the profit minimizing choice; there is an economic incentive to be honest even with the introduction of trunit sales. For any single transaction then, a rational seller will choose to be honest, since this is the profit maximizing strategy. Note, too, that unless c and p have their maximum and minimum values (1 and 0) respectively, $(1-c)r\tau + b(1+p)\tau > b\tau$, meaning that engaging in honest sales is more profitable than simply selling trunits—agents are encouraged to engage in honest sales. This is equally true for a seller who has built trunits over a long history of honest sales, and one who has simply purchased trunits for immediate use.

Honesty continues to be the most profitable policy for sellers in the more general case, where sellers' strategies can consist of schedules of multiple transactions, b is allowed to change over time (as suggested in the next section), etc. An outline of this proof is provided as an appendix.

One critical point must be noted. If the seller has no further goods to sell, engaging in honest trades is not possible. Without the ability to sell trunits, his only (profitable) option was to cheat with the remaining trunits. Now, there is another choice. Moreover, it is more profitable to sell the trunits than to use them to cheat—the exit and surplus trust problems have disappeared.

Commodity Trunits provides an incentive for honesty—rational sellers choose to behave honestly because it benefits them. Beyond the loss of trunits, Commodity Trunits does not directly punish dishonest sellers in other ways (e.g., banning them from the marketplace, modifying/censoring their product offers, etc.). Thus, if an irrational seller (with sufficient trunits) wishes to cheat, he is not directly prevented from doing so. Such irrational acts will result in financial losses for the the seller, in the form of lost future sales and the lost revenue from selling the trunits. Repeated bad behaviour results in repeated losses of trunits, i.e., mounting financial losses. Further, sustaining a pattern of consistently bad behaviour will require the seller to repeatedly acquire new trunits, which in turn will require the continual input of money by the seller. It follows that finite financial resources, and alternate uses for the money, should curtail dishonesty even by irrational agents. Beyond this, if those administering the marketplace are concerned about large numbers of irrational sellers, additional protections might be added: e.g., the use of an escrow service for sellers who are repeatedly dishonest, increasing trunit penalties for seller with large numbers of defaults, banning bad sellers, etc.

4.3 Comparison to Bonds

Trunits may initially appear to be equivalent to a system of using cash bonds to secure transactions: under bonding, money is used to secure a sale, and the seller loses that money if he does not fulfill his obligations. In fact, Trunits does share some characteristics of such a system. It is worth noting that similarity to bonds is not a flaw; bonds are quite powerful in inducing honest behaviour. Rather, we acknowledge some key practical limitations of a bonding system, while noting some of the important features of Commodity Trunits.

First, a bond is money. A trunit is a commodity, one that may be bought and sold. In effect, a trunit is a license that allows the seller to engage in sales of up to r dollars at a time. Trunits can be freely created or destroyed by the market operator, while cash cannot (or should not!) be.

Second, in the case of cash bonds, something must be done with the cash in the event the seller cheats. The problem here is that whoever receives this cash, whether market operator or buyer, has an incentive to treat the seller unfairly. Fairness would require arbitration of alleged cheating incidents, which is not practical for large electronic marketplaces. In contrast, since trunits that are 'lost' by the seller are simply destroyed by the market operator (and have no value to him), there is no such problem under the Trunits model.

Third, honest transactions allow the seller to grow his trunit balance substantially via the reward for honesty, which allows for increasing sales levels (desirable for both seller, and market operator). This is not possible under bonding without the input of additional funds. Significant capital outlay is required only if the seller wishes to ‘jump-start’ his sales volume, rather than actually building trust.

This last point highlights an interesting property of the mechanism. While long-term traders can build trust, for traders who wish to conduct small numbers of sales, the system is very similar to bonding. This is quite positive, in fact. Most trust/reputation systems make it difficult for a short-term seller to effectively sell goods (or get full value for them) because his reputation is undeveloped. Using the best known example, under the eBay feedback system (eBay, Inc., 2007), sellers with established histories (obvious from their high feedback scores) are strongly preferred over new sellers. In (Zacharia et al., 1999), new sellers are intentionally treated as equivalent to disreputable ones, in order to prevent agents from shedding bad reputations by creating new user accounts. Moreover, if a system does allow such new participants to sell effectively, it is very difficult to ensure they will be honest, given their short-term focus. The method presented here overcomes these obstacles as well as a cash-bond system, but accomplishes this seamlessly within the context of a system that allows traders to build trust/reputation over the long term.

5 Validating the market for trust

The safety of Commodity Trunits hinges on a key requirement: that the price of trunits b exceeds the required ratio r. In this section we establish that, while non-trivial, this requirement can be maintained.

The proposed mechanism is dependent on sellers being able to buy and sell trunits whenever they want; otherwise a seller with too few trunits may lose out on sales, or worse, a seller with too many trunits may cheat with them. If agents are restricted to trading trunits only with the market operator, agents may wish to sell more trunits than the operator can afford to purchase, because of the growth in trunit balances due to the premium for honesty. Instead, we look to the other obvious choice—allowing traders to sell trunits to one another as well. Note that the direct sale of trunits from one agent to another is numerically equivalent (and equally safe) to the sale from the first agent to the operator, followed by a purchase by the second agent from the operator at the same price.

Although it is safe for agents to buy and sell trunits where $b > r$, unfortunately this condition is not guaranteed when we allow free buying and selling of trunits among agents. The problem is essentially one of economics.

Trunits are created when honest sales occur, and destroyed by cheating. In a properly functioning marketplace, growth through trunit creation will exceed the loss from cheating. This suggests a growing trunit supply.

Trunits are required to cover each sale made. Over time, as the volume of sales increases (via entry of new traders, or increased activity per trader), the number of trunits required to support this activity will rise. This suggests increasing demand.

It is a fundamental principle of microeconomics that the balance between supply and demand will determine the price of a good in a competitive market. Consider a marketplace with a current equilibrium trunit price of b, and where $b > r$ (i.e., at present, the required safety property holds). If the increase in demand outpaces the increase in supply, then b will rise. If supply growth outpaces that of demand, however, b will drop. Over time, b may fall below r—cheating may become more profitable than honesty.

To prevent this, the operator might simply try to enforce a minimum price. Unfortunately, enforcing a price above the natural equilibrium results in a surplus, where more agents wish to sell trunits than wish to buy them (and thus, encourages cheating). Instead, we employ two techniques. First, our operator buys up excess trunits in the marketplace, to the degree possible given the revenue taken in to that point. Second, the fixed required ratio r is no longer enforced. These two techniques are detailed in the following sections.

5.1 A free market for trust

As price b floats, it is natural to consider allowing the ratio r to be determined by the market as well. Sellers would be free to offer any price/trunit combination they desired, and buyers could decide how many trunits would be required for them to consider a transaction safe.

Consider again the scenario where supply growth outstripped growth in demand, to the point where $b < r$. Here, we would expect to see a surge in cheating in the marketplace. This is likely to have two consequences. First, buyers are likely to increase the trunits demanded to enter into a transaction, so that they feel safer making purchases (*trunit inflation*). Since $\tau = \frac{V}{r}$ (and V is determined by the good sold), increasing τ lowers r, pushing it towards safety. At the same time, increased trunit requirements mean that sellers need more trunits—as demand increases, so does b. These two effects work to restore $b > r$.

Ideally, however, we would prefer that agents be able to adapt to changing market conditions without actually experiencing (or committing) acts of cheating. The market operator (who maintains trunit balances for traders, and hence has access to details on trunit trades) can help to ensure this by providing information to agents—in particular, a *recommended safe ratio*. This allows the market operator to influence r and b significantly, and also relieves agents of the burden of acquiring and processing the data themselves.

5.2 Simulation

The use of simulation to investigate economic markets has a long history (e.g., Naylor, 1971). To investigate the feasibility of maintaining the required safety property in a marketplace, we ran extensive simulations under a well-defined scenario. Our simulation method and findings are detailed in this section.

As outlined above, the trunit price b and required ratio r were allowed to float. An active market operator bought and sold trunits (as described below), and recommended a safe ratio to buyers. For simplicity, buyers faithfully adhere to the recommended ratio. Under these circumstances, evaluation of the method is simplified. If the market operator can set any required ratio r, he can always ensure that $r < b$, as long as $b > 0$; if b reaches zero, then cheating cannot be less profitable than honestly selling trunits, so the market has failed.

Since trunits are only used by sellers, buyers influence the price of trunits in two important ways: *i*) through the purchases that they seek to make, which require trunits to cover them; *ii*) through the quantity of trunits that they demand in order to cover a sale. To study trunit pricing, then, sales of goods to buyers can be modeled using simple frequencies/rates of sales. Since the operator sets r, we need not consider *ii* as part of a buyer model.

Each selling agent is modeled using a variety of parameters, including the following noteworthy values:

- The price for which it sells its goods (V). This is a fixed, advertised price, which is randomly chosen from a uniform distribution with range [1, 100].

- The cost incurred by the agent in selling its goods (c), as a fraction of selling price. This is drawn from a normal distribution, with mean of 0.5.

- The initial rate of sale, expressed in sales per day. This is drawn from a uniform distribution, with range [1, 10].

- The rate of change in the agents' sales volume. This is expressed as an annual percentage change; it is drawn from a normal distribution, with the mean specified by an input parameter to the simulation.

The rate of change may be positive or negative.

- The probability that the agent will decide to exit the market. This is a fixed probability for all agents, specified as an input parameter and expressed as the fraction of sellers who will leave the market per year.
- The agent's conservatism, drawn from a uniform distribution with range [0.5, 1]. This reflects differing temperaments that agents might have, in terms of the degree of certainty they require before taking a 'risky' action. For example, an agent with a conservatism of 0.7 will sell surplus trunits today unless there is at least a 70% chance that the price will be higher on a future day within his horizon.
- The agent's horizon, the number of days he looks into the future when deciding whether or not to buy/sell trunits today. Horizons are drawn from a uniform distribution with range [1,10].

Selling agents forecasted future trunit prices using regression analysis. Uncertainty was incorporated using prediction intervals based on each agent's level of conservatism. Our agents' behaviour, then, consists of the four following rules:

1. **Buy for current need:** If trunits are needed for the current day to execute sales, buy these trunits if it is profitable to do so, considering both the acquisition cost and forecast selling price of the trunits at completion of the transaction.
2. **Buy for future need:** Beyond any trunits bought for the current day, if the agent expects to need additional trunits within its horizon, purchase those trunits today unless the agent expects that the price of trunits will be cheaper within its horizon.
3. **Sell surplus:** If an agent has surplus trunits beyond its expected needs, sell those trunits unless the agent expects that the price will be higher within its horizon.
4. **Speculative sale:** If an agent has trunits that it does not need today, but might need in future, sell these trunits if it is expected that they can be repurchased for less within the horizon.

We believe that this model is sufficient to establish the viability of a trunit marketplace—agents prefer to hold trunits as price increases, and not to hold them as price decreases.

5.2.1 Simulation execution

Here, we briefly outline the simulation execution. The duration of a transaction (i.e., the time between the initial sale, and the trunits being returned to the seller in the case of an honest sale) was 14 days. Sellers joined and left the market each day, and new sales for each active seller were randomly determined. Based on the number of new sales and its forecast of future parameters, each agent determined its trading strategy (i.e., the prices and quantities at which it would buy/sell). Based on this, the equilibrium price that balances supply and demand for trunits was determined, and trunit trades were executed at these prices. (This method is a simplification, based on microeconomic theory, used instead of developing a full bid-ask marketplace.) Subsequently, sales of goods to buyers are executed.

Agents require trunits to conduct sales. Scarcity of trunits would constrain sales, which is likely undesirable. Additionally, serious trunit shortages may result in spikes in trunit prices, trunit speculation, and so on. While trunits are a commodity, we note that: a) they can be freely created by the operator; b) that instability in the (trunit) marketplace is undesirable, since it adds complexity and risk for traders. Thus, the market operator creates and sells trunits whenever insufficient trunits are available to prevent jumps in price. This includes the situation of market start-up—initial quantities are available from the operator.

Trunit growth may cause trunit prices to fall. Prices that drop too quickly are problematic, for several reasons. First, if prices are dropping dramatically, it might not be profitable for agents to engage in sales at all: the gross profit from the sale might be less than the loss incurred by buying trunits and then selling them at a dramatically lower price. Moreover, the velocity of a quickly dropping price increases the danger that prices may crash entirely, reaching zero.

Ensuring that prices do not drop at an inappropriate rate is partly a matter of setting market parameters properly. If trunit growth does not greatly exceed actual sales growth in the marketplace, prices should not drop precipitously. (The interplay of market parameters is examined in our simulation results, below.)

Beyond this, the market operator can also moderate price decreases by purchasing trunits, ‘mopping up’ excess supply. Given that the operator can take in revenue from trunit sales, such policies may be necessary to inspire the confidence of market participants: the perception that the operator is profiting from trunit sales might undermine his credibility in administering the system. For this reason, we would recommend that the market operator’s role be performed with complete transparency to participants. To ensure that the operator

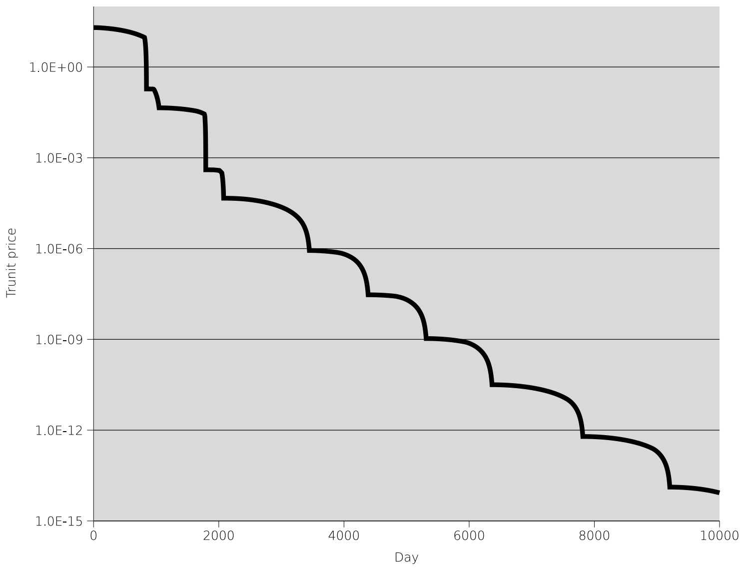

Figure 1: A single simulation run

does not run a deficit, however, the ability to purchase trunits is limited by the amount of revenue taken in previously by creating and selling trunits.

5.3 Simulation results

Initially, we sought to answer two questions: a) Is it possible to sustain a long-term market for trunits? b) How will prices change over time in such a marketplace? To answer these questions, we experimented with parameter values in ranges that we intuitively considered to be reasonable. The results of one such representative run are depicted in Figure 1, to illustrate market operation. (Aggregate results over numerous runs are presented later in this section.) In this simulation, the market was in operation for 10,000 days (more than 27 years). The rate of growth in the number of agents was set at 20% per year, the average growth in agents' sales volumes was 10% per year, the rate at which agents exit the market was 10% per year, and p was set to 0.2.

The market price of trunits b never reached zero over the span of 10,000 days, and b never fell below the required ratio r—cheating was never more profitable than honesty. This confirms the existence of sustainable trunit marketplaces.[9]

---

[9] Over the long term, trunit prices fall to levels at which the values can be cumbersome for participants. For convenience, trunits can be 'merged' when prices are low, in the same manner that stocks undergo 'reverse splits' when prices are low.

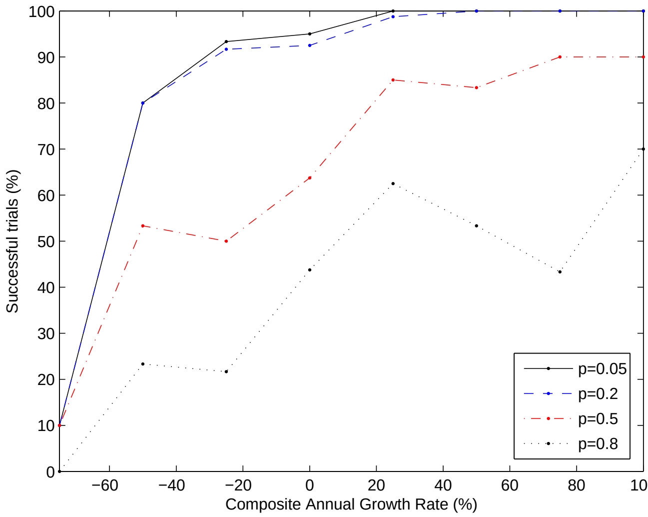

Figure 2: Market sustainability

Note that price drops in a stable, predictable manner. This is desirable, since it makes it easier for agents to predict future prices. The ability to make accurate predictions makes the market safer, and hence more attractive, to sellers. Further, as this market would continue to run beyond 10,000 days, the pattern suggests that prices would continue to drop but remain above zero. Finally, prices appear to drop very quickly, to very low levels. This might initially seem to be troubling: agents may not want to invest in trunits, since quickly dropping prices mean lost value for those holding the trunits. It is worth reiterating, however, that this chart depicts more than 27 years; price decreases are not as fast as they may appear. More importantly, these dropping prices are factored into the agents' profit calculations and decision making—they did not consider the rate of decrease in prices a deterrent, and continued to buy trunits and execute sales. This means that they still found it profitable to operate in the market.[10]

5.3.1 Investigating the impact of market parameters

The results above verify that scenarios exist under which it is possible to manage a sustainable marketplace for trunits, one that preserves the incentive for honesty. From here, we investigated the range of market

---

[10] This may seem trivial, that a growth rate of 20% in trunits (p) is offset by a net positive growth in the marketplace activity. Note, however, that while the market growth parameters are *annual* rates, trunits can grow by 20% every 14 days (i.e., the duration of a transaction.)

parameters for which we could expect a market to succeed (i.e., maintain $b > 0$). We focused our study on parameters most likely to effect supply and demand, and hence price. We tested all combinations of the following parameter values: the premium for honesty p (0.05, 0.2, 0.5, 0.8), the rates at which sellers join (0%, 25%, 50%) and leave (0%, 25%, 50%, 75%) the market, and the average rate of change in existing agents' sales (-25%, 0%, 25%, 50%). For each combination of parameter values, 10 simulations were run, to determine the fraction of trials in which the marketplace was successfully sustained. For presentation, we have consolidated the parameters impacting sales volumes into an approximate composite growth rate. For example, rates of agent growth of 25%, change in sales of 50%, and agents exiting of 25% are reflected in the chart as an approximate composite growth rate of $25\% + 50\% - 25\% = 50\%$. All trials with the same composite growth rate (regardless of differing specific parameter values) have been aggregated in the results depicted in Figure 2. Simulations began with 100 sellers, and ran for 10 years (3650 days); the figure depicts the percentage of trials for which the market was successfully sustained.

We expected that market failures might result from high values of p (since trunit creation enhances supply) and for low rates of market growth (since lower sales volumes result in lower demand). Our expectations were confirmed to be correct: relatively high values of p (0.5, 0.8) resulted in relatively low success rates, particularly at negative rates of market growth. More reasonable premium rates resulted in markets that were sustainable with very high probability: at $p = 0.05$, markets were sustainable in 100% of cases for even the lowest positive growth rate tested. Even in conditions of market stagnation or decline, p values of 0.05 and 0.2 were sustainable in over 90% of trials. It is expected that lower values of p would yield even higher success rates. These results match the intuition that trust should grow slowly, as each agent demonstrates its trustworthiness.

We note two important points here. First, an operator may wish to provide significant rewards for honesty, yet be concerned about possible future conditions of slow/negative market growth. The operator, then, can adjust p if necessary depending on observed market conditions, to assure market sustainability. Second, rates of success decline for highly negative growth rates. We suggest that if a market is in a state of rapid decline, issues beyond the trust model demand the attention of the operator.

In short, for a range of realistic market conditions, trunit marketplaces are sustainable.

6 Discussion

6.1 Addressing vulnerabilities

Several models for addressing trust in electronic marketplaces can be seen as predictive: the future (un)trustworthy behaviour of partners is learned in order to avoid or discount less desirable business partners (e.g., Jøsang and Ismail, 2002; Teacy et al., 2006; Tran and Cohen, 2004; Yu and Singh, 2002). Kerr and Cohen (2007a,b) discuss the vulnerabilities faced by these models, where agents are dishonest in ways that are strategic, but difficult to predict. In contrast, Commodity Trunits offers important protections against vulnerabilities, when agents act rationally.

Commodity Trunits provides the same protection against reputation lag, value imbalance, and re-entry as Basic Trunits. Moreover, Commodity Trunits offers a solution to the exit problem—a seller earns more money by honestly selling trunits than by using them to cheat. The closely related *surplus trust* problem also disappears: if a seller accumulates more trust than required to cover his predictable periodic transactions, they can likewise be sold.

Trunits does not regulate the behaviour of buyers. As such, it does not yet (on its own) provide complete protection against some collusive attacks involving both buyers and sellers. One such attack, that of bad-mouthing, involves buyers giving false negative reviews of a seller, in order to damage that seller's reputation. A typical goal of such an attack is to improve the chances of a competitor (i.e., part of the coalition) of winning sales, at the expense of the target. Trunits, alone, provides no protection for sellers from bad-mouthing. Note, however, that bad-mouthing will not directly undermine the stability of the market for trunits—bad-mouthing artificially destroys trunits, reducing available supply, which will exert upward pressure on trunit prices.

A valuable property of Commodity Trunits, however, is that it provides protection against an important form of collusion: *ballot-stuffing*, where fabricated positive ratings are used to artificially inflate reputations, in order to increase sales. Bhattacharjee and Goel (2005) suggest that ballot-stuffing can be discouraged if transaction costs (e.g., commissions) are larger than the expected gain from the activity. Let a selling agent's cost ratio $c = c_{oe} + c_{se}$, where c_{oe} includes all operating costs in furnishing the good for sale, and c_{se} is the selling expense (in particular, commission.) A ballot stuffing transaction requiring τ trunits would return a premium of $p\tau$ trunits, but incurs only the selling expense, $Vc_{se} = r\tau c_{se}$. To simply purchase the same quantity of trunits would cost $bp\tau$.

Ballot stuffing is unattractive, then, if $r\tau c_{se} > bp\tau$, or $rc_{se} > bp$. Since the operator can exert control over r (as explained above), he can maintain $r = kb$ for some constant k, where $0 < k < 1$. Thus, $kbc_{se} > bp$, or $kc_{se} > p$. Since the market operator controls both the premium and the commission rate, they can be set to make ballot-stuffing unprofitable, but this requires that $c_{se} > p$. (A ratio that only approaches this inequality may still discourage the activity.)

The result above shows that a rational seller, trying to earn profit within the market, will not ballot-stuff when market parameters are set appropriately: he makes more money by not doing so. Another possibility, however, is that an attacker might wish to cause the trunit market to fail. Ballot-stuffing, by artificially creating trunits, might increase supply, exerting downward pressure on prices. If enough such pressure were applied, the market might crash.

Note, however, that when market parameters are set as specified above, the market operator can effectively prevent this from succeeding. As noted above, the value of the trunits acquired in a ballot-stuffing transaction is less than the cost of the commission paid to the operator. This means that for every quantity of trunits created by ballot-stuffing, the operator gains sufficient revenue to buy it back when it is dumped on the market.[11] Thus, the attacker cannot create more trunits via ballot-stuffing than the operator can afford to purchase.

To verify this, we conducted another set of experiments where attackers aggressively attempted to crash the market by engaging in massive ballot-stuffing for a period of time to build trunits, then dumping all of their trunits abruptly on the market. Trails were run with the following parameters: quantity of ballot-stuffing equal to $\{50\%, 100\%\}$ of the total legitimate sales in the marketplace, and dumping occurring after $\{5\%, 50\%, 75\%\}$ of the duration of the simulation run. All combinations of these values were tested, against all combinations of the parameter values specified in Section 5.2 (except that p was set at 0.015 and the commission rate at 0.2, to satisfy the requirements noted above.) The operator cannot tell which sales are ballot-stuffing and which are legitimate, so simply bought up trunits as necessary to maintain market stability.

In every such run, the attackers were unsuccessful in causing the marketplace to fail—the operator was

---

[11] This revenue will be sufficient to purchase the trunits at all points in the future, as long as trunit prices do not increase. As noted in Section 5.2.1, the operator creates trunits as required to prevent price jumps.

always able to maintain stability. This confirms the robustness of the trunits marketplace in the face of ballot-stuffing.

It may appear that under this system, the market operator might not earn legitimate profits for running the marketplace: operational revenue (commissions) will be eaten up buying trunits resulting from ballot-stuffing. Note, however, that only the commissions from ballot-stuffing transactions (i.e., revenue that would not have existed in the absence of ballot-stuffing) are required to buy back ballot-stuffing trunits. Commissions earned by legitimate transactions can be realized as profit.

While this shows that Commodity Trunits, even on its own, can provide some security against ballot-stuffing, we hasten to note that much stronger protection might be realized. Ballot-stuffing relies on collusion between buyers and sellers; the integration of a system for ensuring buyer honesty will provide substantially improved protection.

6.2 Desirable marketplace properties

Under Commodity Trunits, buyers only trade with sellers who have sufficient trunit balances, with no further modeling required. This contrasts with trust and reputation modeling systems (e.g., the work of Teacy et al. (2006); Tran and Cohen (2004); Yu and Singh (2002)) that track sellers and develop methods to determine which ones are sufficiently reputable. One important contrast, therefore, is the amount of computational effort required for these systems.

At the same same time, Trunits is very egalitarian in its treatment of sellers. It should be noted that where the same risk-ratio r applies to all traders, a seller cannot use a large trunit balance to make its offering more attractive than that of other sellers; any seller with sufficient trunits to cover a transaction can compete. In fact, a new and an established seller competing for a sale will be treated equally if they have enough trunits to cover the sale, regardless of their total accumulated trunit balances. In contrast, many models (such as BRS (Jøsang and Ismail, 2002) and TRAVOS (Teacy et al., 2006)) treat buyers with longer histories differently than those with shorter histories, even if they have been honest with the same relative frequency.

Similarly, we consider it to be desirable for every buyer in the system to be equally protected. If new participants are more vulnerable than established buyers (as, for example, in models where buyers rely on direct experience to choose trustworthy partners), it can serve as a disincentive for new buyers to enter the market. Under Trunits, a seller's ability to engage in a sale is determined by that seller's trunit balance, regardless of the identity of the buyer—every buyer receives the same degree of protection.

Beyond the egalitarianism, there are other advantages for buyers as well. When different sellers have different degrees of trustworthiness, a buyer may feel compelled to chose the 'most trustworthy' seller, even if that seller's good does not match his preferences as well as one offered by a less trustworthy seller. Under Trunits, however, each agent with sufficient trunits can be considered equally trustworthy—the buyer is then free to choose products that best match his needs.

6.3 Contrast with other incentive-based approaches

Some research in trust and reputation modeling bears similarities to our work. The work of Braynov and Sandholm (2002) develops a mechanism to protect buyers by attempting to regulate seller behaviour. Under this incentive-compatible mechanism, the profit maximizing strategy for sellers is to be honest. While the soundness of the system is established by the authors, it targets a somewhat different marketplace than most proposals. In the envisioned scenario, each seller declares its degree of trustworthiness; based on this information, buyers choose sellers. Note that the incentive is designed so that profit is maximized by truthfully *declaring* trustworthiness, rather than necessarily acting in a trustworthy manner. Also, there is a fairly strong assumption that buyers know the sellers' cost functions.

Researchers in peer-to-peer networks, have shown interest in trust and reputation schemes, in order to address problems like freeloading (where users take advantage of resources without likewise contributing) and the insertion of useless content (disguised as desirable content) into the system. Some of the proposals (e.g., Gupta et al., 2003) are superficially similar to Trunits, controlling transactions through the use of transactional units; these units may even be referred to as 'trust'. Generally, under such systems units are gained through the provision of resources, and surrendered in the consumption of resources—in effect, this is an exchange process, where the units perform the role of currency more than trust. Such systems may be very effective in preventing problems such as freeloading, but they do not protect against the sort of vulnerabilities with which we are concerned.

It should be noted that predictive models, such as those mentioned in Section 6.1, can have an incentive effect—if the seller knows he is being modeled, he may act honestly in order to improve his prospects for future sales. In this case, however, the implicit incentive may be unclear or difficult to measure, so agent behaviour may be unpredictable, and safety from vulnerabilities may be difficult to ascertain.

7 Future Work

Trunits makes no attempt to regulate or predict the behaviour of buyers. Operation of the model depends upon the honesty of buyer feedback, however, assuming the use of a parallel mechanism for eliciting honest responses. Considering the trustworthiness of buyers is an important area of future investigation. The ideas of (Jurca and Faltings, 2003) may be especially valuable here, as they constitute methods to promote buyer honesty, to complement our current work promoting seller honesty.

As noted, Commodity Trunits provides some protection against ballot stuffing, but is not resistant to other forms of collusive attack. Collusion is a problem to which most (if not all) current trust and reputation systems are vulnerable, and is a central issue for future research. Incorporation of a system to ensure the honesty of buyers, as noted above, may provide protection against collusion; such integration/expansion of the system must be carefully conducted to ensure that honesty is clearly profit-maximizing. Additionally, exploration of variable rewards may be fruitful. For example, rewards might be reduced for reviews from new reviewers (as by Zacharia et al. (1999)), or for repeated reviews by the same reviewer (as by eBay, Inc. (2007); Zacharia et al. (1999)).

As presented, Trunits is a centralized model. There is nothing about the mathematics of the model that necessitates centralization, however. The possibility exists for a decentralized implementation in which trunit balances are stored locally for each seller, and updated as transactions occur. Such an implementation would require secure storage and update of trunit balances; a cryptographic scheme such as that used in Gupta et al. (2003) might permit this.

8 Conclusion

Trunits is a novel trust model for electronic marketplaces, which seeks to ensure rather than predict the trustworthiness of participants. The Basic Trunits system, under which trunits are assigned to a specific trader and are not transferable, has a number of attractive characteristics. It provides clear incentives for honesty, which can be measured and verified by both system designers/operators and market participants; this serves as a basis for rational (and predictable) behaviour on the part of agents. Basic Trunits provides solutions to a number of important vulnerabilities, including reputation lag, value imbalance, and re-entry. Basic Trunits is computationally simple; further, decision making is extremely simple for buyers (in determining whether to trust a seller) and sellers (in determining whether or not to cheat). Basic Trunits is egalitarian, in that all buyers receive equal protection regardless of their experience, and all sellers with sufficient trunits are considered equally trustworthy, regardless of the length of their history.

While Basic Trunits is an important step towards secure trust models, it has certain limitations. In particular, it provides no protection against collusion, and is subject to the surplus trust and exit problems. It also provides no single solution to the start-up problem, the question of how an agent acquires an initial quantity of trunits.

Commodity Trunits is an extension of Basic Trunits, one in which agents are permitted to buy and sell trunits. Commodity Trunits retains many of the characteristics of Basic Trunits, while providing important enhancements. Commodity Trunits remedies the start-up problem, by allowing agents to purchase trunits at start-up. Commodity Trunits alleviates concerns of artificially constrained trading volume—any time an agent has insufficient trunits, it can purchase the required quantity. Commodity Trunits provides a solution to the important exit problem, where an agent who is leaving the market can cheat freely. Similarly, Commodity Trunits provides a remedy to the surplus trust problem, in which an agent can use surplus trust beyond its routine needs to cheat. Commodity Trunits also provides protection from an important form of collusion, ballot stuffing. Commodity Trunits provides seamless handling of traders with both short-term and long-term market presences, ensuring honesty in both cases.

The security of Commodity Trunits depends on the ability to maintain a price of trunits higher than the value that can be gained by using them to cheat. Simulation verified that it was possible to maintain this price relationship; in fact, Trunit marketplaces were shown to be stable and sustainable over the long term, fostering a thriving and safe trading environment.

This research constitutes important progress towards providing secure trust and reputation models, a direction we feel is important to promote broad participation in electronic marketplaces.

9 Acknowledgments

The authors wish to thank the Natural Sciences and Engineering Research Council of Canada (NSERC) for its generous support.

References

Rajat Bhattacharjee and Ashish Goel. Avoiding ballot stuffing in ebay-like reputation systems. In *P2PECON ’05: Proceeding of the 2005 ACM SIGCOMM workshop on Economics of peer-to-peer systems*, pages 133–137, New York, NY, USA, 2005. ACM Press. ISBN 1-59593-026-4. doi: http://doi.acm.org/10.1145/1080192.1080203.

Sviatoslav Braynov and Tuomas Sandholm. Incentive compatible mechanism for trust revelation. In *AAMAS ’02: Proceedings of the first international joint conference on Autonomous agents and multiagent systems*, pages 310–311, New York, NY, USA, 2002. ACM Press. ISBN 1-58113-480-0. doi: http://doi.acm.org/10.1145/544741.544814.

Chrysanthos Dellarocas. The digitization of word of mouth: Promise and challenges of online feedback mechanisms. *Management Science*, 49(10):1407–1424, 2003a. ISSN 0025-1909. doi: http://dx.doi.org/10.1287/mnsc.49.10.1407.17308.

Chrysanthos Dellarocas. Goodwill hunting: An economically efficient online feedback mechanism for environments with variable product quality. In *AAMAS ’02: Revised Papers from the Workshop on Agent Mediated Electronic Commerce on Agent-Mediated Electronic Commerce IV, Designing Mechanisms and Systems*, pages 238–252, London, UK, 2002. Springer-Verlag. ISBN 3-540-00327-4.

Chrysanthos Dellarocas. Efficiency and robustness of binary feedback mechanisms in trading environments with moral hazard. Working papers 4297-03, Massachusetts Institute of Technology (MIT), Sloan School of Management, April 2003b. available at http://ideas.repec.org/p/mit/sloanp/1852.html.

eBay, Inc. Web site, 2007. URL `http://www.ebay.com/`. http://www.ebay.com (accessed October 2007).

Eric Friedman and Paul Resnick. The social cost of cheap pseudonyms. *Journal of Economics and Management Strategy*, 10(2):173–199, 2001.

Diego Gambetta, editor. *Trust: Making and Breaking Cooperative Relations*. Blackwell, 1990.

Minaxi Gupta, Paul Judge, and Mostafa Ammar. A reputation system for peer-to-peer networks. In *NOSSDAV '03: Proceedings of the 13th international workshop on Network and operating systems support for digital audio and video*, pages 144–152, New York, NY, USA, 2003. ACM Press. ISBN 1-58113-694-3. doi: http://doi.acm.org/10.1145/776322.776346.

Audun Jøsang and Roslan Ismail. The beta reputation system. In *Proceedings of the 15th Bled Electronic Commerce Conference e-Reality: Constructing the e-Economy*, June 2002.

R. Jurca and B. Faltings. An incentive compatible reputation mechanism. In *Proceedings of the IEEE Conference on E-Commerce*, Newport Beach, CA, USA, 2003. URL `citeseer.ist.psu.edu/jurca03incentive.html`.

Reid Kerr and Robin Cohen. Trunits: A monetary approach to modeling trust in electronic marketplaces. In *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'06) Workshop on Trust in Agent Societies*, Hakodate, Japan, 2006.

Reid Kerr and Robin Cohen. Towards provably secure trust and reputation systems in e-marketplaces. In *Proceedings of the Sixth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'07)*, Honolulu, Hawaii, USA, 2007a.

Reid Kerr and Robin Cohen. Guaranteed security in trust and reputation systems. In *Proceedings of the Sixth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'07) Workshop on Trust in Agent Societies*, Honolulu, Hawaii, USA, 2007b.

Niklas Luhmann. *Trust and Power*. Wiley, 1979.

Roger C. Mayer, James H. Davis, and F. David Schoorman. An integrative model of organizational trust. *The Academy of Management Review*, 20(3):709–734, 1995.

Thomas H. Naylor. *Computer Simulation Experiments with Models of Economic Systems*. John Wiley & Sons, 1971.

W. T. Newlyn. *Theory of Money*. Oxford University Press, 1971.

J. B. Rotter. Generalized expectancies for interpersonal trust. *American Psychologist*, 26(5):443, 1971.

Denise M. Rousseau, Sim B. Sitkin, Ronald S. Burt, and Colin Camerer. Not so different after all: A cross-discipline view of trust. *The Academy of Management Review*, 23(3):393–404, 1998.

W. T. 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. ISSN 1387-2532. doi: http://dx.doi.org/10.1007/s10458-006-5952-x.

Thomas Tran and Robin Cohen. A learning algorithm for buying and selling agents in electronic marketplaces. In *AI '02: Proceedings of the 15th Conference of the Canadian Society for Computational Studies of Intelligence on Advances in Artificial Intelligence*, pages 31–43, London, UK, 2002. Springer-Verlag. ISBN 3-540-43724-X.

Thomas Tran and Robin Cohen. Improving user satisfaction in agent-based electronic marketplaces by reputation modelling and adjustable product quality. In *AAMAS '04: Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 828–835, Washington, DC, USA, 2004. IEEE Computer Society. ISBN 1-58113-864-4. doi: http://dx.doi.org/10.1109/AAMAS.2004.145.

Oliver E. Williamson. Calculativeness, trust, and economic organization. *Journal of Law and Economics*, 36(1):453–486, 1993.

Bin Yu and Munindar P. Singh. Distributed reputation management for electronic commerce. *Computational Intelligence*, 18(4):535–549, 2002. doi: 10.1111/1467-8640.00202. URL `http://www.blackwell-synergy.com/doi/abs/10.1111/1467-8640.00202`.

Giorgos Zacharia, Alexandros Moukas, and Pattie Maes. Collaborative reputation mechanisms in electronic marketplaces. In *HICSS '99: Proceedings of the Thirty-second Annual Hawaii International Conference* *on System Sciences-Volume 8*, page 8026, Washington, DC, USA, 1999. IEEE Computer Society. ISBN 0-7695-0001-3.

List of Figures

Appendix A Safety in the general case

Consider the general case, where a seller can freely engage in multiple transactions, with arbitrary timing. As described in Section 5, b and r can vary, and sellers can hold and trade trunits freely. We outline an inductive proof of safety on n, the number of transactions engaged in by a seller (a *schedule*).

Base case: Any single transaction, n = 1. Any seller faced with a single transaction (requiring τ trunits) has three choices: execute the sale honestly, cheat, or forsake the sale and instead honestly sell his trunits. We consider the latter two options here, addressing the first afterwards.

A sale may be covered with existing trunits, and/or purchased ones. Let τ_h be the number of existing trunits used for the transaction, and τ_p the number that are purchased, so $\tau = \tau_h + \tau_p$. If the agent cheats, he gains (at most) $r\tau = r(\tau_p + \tau_h)$, and incurs a cost of $b\tau_p$ to purchase the trunits. If he simply sells his trunits instead of engaging in the transaction, he gains $b\tau_h$. Honestly selling the trunits is more profitable if $r(\tau_p + \tau_h) - b\tau_p < b\tau_h$, simplifying to $r < b$. Since this is true by assumption, then for any single transaction, forgoing the sale and selling the trunits honestly is more profitable than cheating.

While cheating minimizes profit, will the agents want to engage in actual sales when b can vary? Consider the case where an agent must purchase τ trunits to engage in a sale. He can purchase the trunits at price b_{start}, and any remaining trunits can be sold at price b_{end}. His profit for a single transaction, then is $(1-c)r\tau + b_{end}(1+p)\tau - b_{start}\tau$. Since a seller would likely only engage in a sale if his anticipated profit is greater than 0, we set the inequality $0 < (1-c)r\tau + b_{end}(1+p)\tau - b_{start}\tau$. Ideally, all traders should wish to engage in sales, so we set $c = 1$, the maximum cost ratio. Simplifying, this yields $\frac{b_{start}}{b_{end}} < (1+p)$, that is, the rate of decrease in trunit prices (per transaction duration) cannot exceed the premium for honesty. As discussed, in our simulation, this condition was never violated: agents always found it profitable to engage in honest sales.

Induction step: Assume that honesty maximizes profit for any schedule of n transactions. Consider an arbitrary schedule of size $n+1$. We choose one transaction arbitrarily (i). In the absence of i, the remaining set of transactions constitutes a schedule of size n. The addition of i to the schedule does not prevent any other transaction from being executed (or vice versa), since trunits can be purchased. Thus, the n transactions can be executed independently of the execution of i, and by the induction hypothesis, yields maximum profit when executed honestly.

Consider now the introduction of i (with starting time t_1) into the schedule. If existing trunits are used for i, it might deprive a transaction in n (starting at time t_2) of some quantity of trunits τ. This saves the purchase of τ trunits at price b_{t_1}, but requires a purchase at price b_{t_2} to execute the later transaction. The net impact on profit is $\tau(b_{t_1} - b_{t_2})$. In the absence of i, however, we could simply sell the τ trunits at time t_1 and repurchase the same quantity at t_2: the net impact on profit is $\tau(b_{t_1} - b_{t_2})$, identical to the case where we use the trunits for transaction i. By similar argument, the holding and disposal of trunits for i has no impact on the execution or profitability of the n transactions, and vice versa. i, then, may be treated simply as a single, isolated transaction: an instance of our base case. As shown above, executing the single transaction honestly maximizes profits. Thus, for any schedule of $n+1$ transactions, executing them honestly maximizes profit. By induction, profit for any schedule of size greater than 0 transactions is maximized by honestly executing every transaction. It should be noted that whether an agent cheats or not depends on only two parameters: b and r.