Inter-Market Competition of Reputation Infrastructure: An Agent Based Approach

James "Augie" Hill

April 25, 2011

Abstract

E-commerce markets are prone to cheating between participants because of the anonymity of Internet interactions and the lack of governmental jurisdictional authority to prosecute cheaters. For this reason, there has been much research and development into mechanisms that prevent cheating. Sets of mechanisms that seek to reduce or eliminate cheating in online markets are known as *reputation systems* or *reputation infrastructure*. Previous work on reputation systems has considered markets as existing alone. In reality, however, markets exist alongside other markets, so this paper builds on previous work to examine the dynamics of a world in which there are multiple markets differentiated by reputation infrastructure. In the model presented in this paper, agents with varying strategies are matched to play Prisoner's Dilemma, and then can switch to a different market if they choose. Experiments are performed to examine general behavioral dynamics and the effects of varying the model parameters.

1 Introduction

The Internet is a medium upon which people can interact anonymously. One of the types of interactions that occur on the Internet is the exchange goods and services, which is termed *e-commerce*. These exchanges typically occur in digital marketplaces where sellers list goods and services and buyers can browse the current offerings and initiate an exchange.

There are some interesting aspects to e-commerce marketplaces that make them different than real-world marketplaces. The most important is that of identity. In an online community, it is simple to change one's pseudonym, or the identification handle under which one trades, by simply signing up for a new account, and thus shed all of one's history and obligations. In the real world, identity change, also known as whitewashing, can be accomplished by producing false documentation, but one would also need to change one's appearance before rejoining the same marketplace, so this is not a practical attack. In an e-commerce marketplace, it is also a simple process to assume multiple identities and control them all nearly simultaneously, for which there is no practical analog in the real world. A possible solution to the problems related to pseudonymous identity in e-commerce markets is to verify the real-world identity of every participant, but the significant overhead cost of this solution makes finding an alternative online-based solution extremely desirable. Another important aspect that distinguishes e-commerce markets from real-world markets is policing. When an incident occurs, the e-commerce market's computer servers, the cheating participant, and the cheated participant may all be in different jurisdictional areas, so the assignment of jurisdiction between governmental authorities is still unclear on the Internet, and usually results in no police intervention. Because of this, there is a vacuum of governmental authority, and cheating participants can commit their transgressions without real-world punishment.

Because of the ease of shedding one's historical abuses and a lack of governmental policing, e-commerce marketplaces are inherently more prone to cheating behavior than are real-world marketplaces [1]. For this reason, there has been much effort towards designing mechanisms that both incentivize good behavior and provide information on the likelihood of market participants to swindle fellow participants. The set of mechanisms that e-commerce marketplaces employ to inform agents of the historical actions of other players is termed a *reputation system* or a *reputation infrastructure*.

Previous research has focused on the relative effectiveness of varying reputation infrastructures in a single-market environment. By considering each market in its own separate environment, a possibly important aspect of the real world is abstracted away: the consumer choice between markets. In the real world, players can choose the market in which they will participate. This paper investigates the dynamics of a world where there are multiple markets from which players can choose to conduct their business. The markets are differentiated by their reputation infrastructures, and thus vary in the amount and quality of the information available to market participants.

Some of the most interesting previous work in the area of reputation systems has focused on decentralized methods, particularly with peer-to-peer markets of digital goods in mind, where there may be only limited centralized information and control of the network. In a decentrazlied system, participants differ in their access to information, as opposed to a centralized system, where the same information is available to all participants. Because of the interesting aspects unique to decentralized reputation systems, this paper will focus largely those types of reputation infrastructures.

The model described in this paper functions most similarly to the model in [1] with a few exceptions. There are some number of marketplaces to which agents can choose to exist after some amount of time. Agents have options to change marketplaces at any time after the minimum number of games in a market have been reached, but this requires starting a brand new pseudonymous identity that only exists within a single marketplace. Each agent can keep their pseudonyms within all marketplaces in which they have interacted and come back to them at a later time. In order to focus the efforts of this paper on the effects of cooexisting markets, the effects of multiple possible pseudonyms per agent have been abstracted, which was also abstracted in the related work [1, 2, 3].

The markets are designed with the various reputation infrastructures described in [1], and an additional market without any reputation infrastructure is included to act a control. All of the agents begin in the control market. The primary metric is the number of agents of each type in each market over time. Because the utility of an agent depends upon both the strategy mix and the reputation infrastructure, I believe this model may produce an unstable system with interesting dynamics.

My hypothesis for this model is that there will be an interesting dynamic that exists where agents that have found a good market may not take the short-term risk associated with a new unknown marketplace to find out if that market will eventually improve their utilty per time step, and thus they will be stuck in a local optimum when there exists a better global optimum. I also think that there will be a dynamic interplay between populations of agents with different strategies. Particularly, I think that the cheating agents will "chase" altruistic agents from market to market.

This paper takes the following form. The related research is first described, followed by a description of the model, and then the experimental results.

2 Related Work

There has been much research in the past fifteen years towards the design of better online-based systems for helping market participants assess the risk of interacting with other market participants [1, 2, 3]. Mechanisms have been researched with the goal of disincentivizing whitewashing, and systems of signals that provide information on the amount of risk in dealing with a particular market participant have been modeled. These signals of risk are also referred to as the amount of *trust* or *reputation* an agent carries.

The previous research examined modeled markets as existing alone in the world [1, 2, 3]. However, this is not an accurate representation of reality, where markets cooexist and compete for agent participation on several aspects - including features, niche, fees, recommendations, and reputation systems. The primary motivation for this paper was to examine the conclusions of the work in [1] when there are multiple markets from which agents can choose to participate.

The models in the referenced papers address the problems of designing reputation systems for decentralized markets. In a centralized market, all agents are connected to all other agents through a mediator, and the mediator provides the same reputation information to all participants. In a decentralized market, each agent is connected to a local subset of the other agents in the market, and reputation information is collected only from those local agents. Because of this limitation, all agents have a different and limited view of the markets. The challenge, therefore, is to create information sharing mechanisms that provide the most useful information using the fewest resources. Decentralized markets present interesting challenges while also having important applications, and thus the work presented in this paper also considers decentralized markets and corresponding reputation systems.

This paper builds primarily upon the work presented by Mui et al [1], which uses evolutionary game theory as a basis of a simulation framework devised with the purpose of understanding the various aspects of reputation in the domain of distributed artificial intelligence. The authors of [1] first describe context-specific reputation. Clearly, people will be better at some things than other, so an aggregation of all reputation data points from various contexts would create a muddy picture of the actual reputation of that person. The authors next describe the personalization of reputation. For example, there may be a certain global reputation score, but there may also be a reputation score based on what people an agent trusts think of the agent in question. It's the difference between creating the score based off of the feedback from everybody versus the people one trusts. Another described aspect to reputation is the groups an agent belongs to. Groups of agents have reputations that often become a signal about the members of that group, even to the persons without reputation data points within the group.

The authors of [1] also describe the differences between direct and indirect reputation. They note that it has been found in previous research that direct interaction means more to people than indirect interaction, but that indirect reputation is used as a bootstrapping method of assessing reputation.

The authors next introduce a framework for simulating reputation that is based upon the ideas of evolutionary game theory. They base their model on the idea from previous work that agents cooperate through reciprocal altruism (also known as the "tit-for-tat" strategy). The model is of a population of agents interacting with a random second agent playing Prisoner's Dilemma. After a generation, agents produce offspring proportional to their current fitness. The authors always ran their simulations with initially half the population being reputation tit-for-tat agents and the other half being all-defect strategies. The reputation tit-for-tat strategy is the same as the tit-for-tat strategy, except it bootstraps its interactions with an agent not encountered before by using the available reputation information to choose the action it believes is most likely to be chosen by the opponent. In particular, the authors' goal was to determine the conditions under which reputation-based tit-for-tat was evolutionarily dominant using several different notions of reputation.

While this paper is based primarily upon the work of [1], the model herein does not use an evolutionary game theory approach, in which the population of agent strategies adjusts proportionally to the fitness of that strategy. The first reason for not using this method is that of simplicity. If conclusions can be drawn using a simpler model, then there is no reason to creating any complications. The argument can be made either way as to which model is a more accurate representation of reality, and which one is more accurate depends upon the nature of human behavior in the real world. If it is true that people will change their strategies to whichever is more profitable, then the evolutionary strategy model is more accurate, but if it is true that a person's strategy is innate, then the fixed strategy model is more accurate. It is also likely that the real world is a mix of these two models. For the sake of simplicity, this paper will use a fixed strategy model.

The primary hypothesis of the authors in [1] was that the addition of reputation to the tit-for-tat strategy for use in the initial encounter with a particular opponent would reduce the number of encounters per generation necessary for agents playing the reputation tit-for-tat strategy to dominate the all-defect strategy. Their results show that the number of generations to evolutionary dominance decreases depending upon the reputation infrastructure of the market in the following order: none, observed, grouped, propagated with a depth of 1, and propagated with depth of 3, with the final infrastructure type requiring less than one-twelfth the encounters per generation as in the control market for reputation tit-for-tat strategy to dominate the all-defect strategy. The authors do not provide enough information about the specific parameters of each of the types of reputation infrastructure that produced the given ranking of effectiveness. Without specifying parameters, they also could not present conclusions on the robustness in ranking to variations in the parameters that describe the reputation infrastructures. A primary investigation of this paper is thus to determine if different market parameter values could produce a different order of effectiveness than those presented in [1] and the robustness of those rankings.

The purpose of work in [2] is to introduce a framework for distributed reputation that produces a "safe, fair, and profitable" agent-mediated marketplace. The framework attempts to combine the understanding of complex interactions from the domains of sociology and economics to provide a "comprehensve" model of how trust and recommendation affect decisions in a market.

In their agent-based model, the marketplace is one for the buying and selling of information goods and services. There are two types of agents: producers and consumers. The producer agents seek to maximize their profits, which for each producer is a function of its cost function and its pricing strategy. To simulate unreliability as in real-world applications, producers are instantiated with a level of competence and a character type that will control how the agent behaves and whether it will succeeed in its tasks. The inner workings of the consumer agent are far more complex that the producer agent.

As part of their experiments, the authors of [2] model the effects of four different types of "hazards" on the marketplace to see how their consumer agent behavior will behave: consumer market entry, incompetent producers, malicious producers, and dishonest recommendations. In this simulation of dishonest recommendations, two thirds of the consumers produce distorted recommendations with the incentive of lowering the price of the producer from which they purchase goods. They find in this simulation that trust and reputation applied to the recommendations of other consumers improves the utility of the altruistic users.

The work in [2] provides very interesting insight into the likely realistic model of consumer behavior in a marketplace by bringing together all of the socio-economic research into a single comprehensive agent-based model. However, many of the details of this model are beyond the scope of this paper, which is focusing on the effects of multiple markets on agent dynamics. Depending on initial findings, it may be important to examine detailed decision-making processes in future work as has been done in [2].

In [3], the authors propose the use of a reputation management system as a tool for the management of risk in e-commerce interactions. They argue in particular in favor of a bottom-up approach to reputation management, where the reputation information comes from the users of the e-commerce system themselves. The primary contribution of [3] with concern to this paper is *market flexibility*, or the rate of inflow and outflow of agents to and from a marketplace, which can have a significant impact on which strategies are dominant within that marketplace depending on its reputation infrastructure. Market flexibility will certainly be an important dynamic of the model described in this paper.

The authors model a consumer-to-consumer online marketplace in which participants, when exchanging goods or services, are modeled as playing a round of Prisoner's Dilemma. They make a convincing argument that this is a simplification without any loss of detail.

In their agent-based model, each agent has strategies, goods to sell, goods to buy, the importance of reputation to the agent, and the amount of time before the agent forgets a transaction from the past. The strategies available to an agent are all-cooperate, all-defect, tit-for-tat, and random.

The authors make the distinction between negative and positive reputation systems. In a negative reputation system, the agents calculate the reputation of others based solely on the number of times the opponents have defected, while the positive reputation system focuses on the number of times the opponents have cooperated. It is not clear that the distinction is significant enough to warrant much consideration.

In their experiments, the authors found that when the market flexibility was low, the dominant strategy was all-cooperate, while when the market flexibility was high, the dominanent strategy was all-defect. This makes sense because the all-defect strategy is best for short-term utility and the all-cooperate strategy is best for long-term utility. They also found that using a positive reputation system resulted in cooperation becoming the dominant strategy even when the turnover rate was high.

3 Model

The model consists of a fixed set of markets and agents. Every tick, agents in each market are randomly matched to play a game of Prisoner's Dilemma, which is representative of a producer-consumer interaction in tangible goods markets. If there are an odd number of agents, one agent is not matched. When agents are matched, each choose an action according to their strategy and the opponent with whom they have been matched. The result of the game is then reported to each of the agents.

Agents are assigned a different unique pseudonym in each market. Opposing agents know only the pseudonym of the agent they are playing, so they cannot recognize when they are playing the same agent in a different market. Agents are differentiated by the strategies they employ. There are four, which are as follows.

Table 1: **Agent Strategies**

Always Cooperate (*C*)	Chooses to cooperate no matter the opponent.
Always Defect (*D*)	Chooses to defect no matter the opponent.
Tit-for-tat (*TFT*)	Chooses to cooperate upon first meeting an agent. On subsequent meetings, chooses whatever the opposing agent chose in their most recent meeting.
Reputation Tit-for-tat (*RTFT*)	Upon first meeting an agent, uses the reputation information available from the market to calculate how likely the other agent is to cooperate. If that probability is higher than the defection threshold, then the agent chooses to cooperate. Otherwise, it chooses to defect. On subsequent meetings, the agent chooses whatever the opposing agent chose in their most recent meeting.

There are four markets differentiated by reputation infrastructure. The reputation infrastructure defines the available historical information, which the reputation tit-for-tat agents use to compute the probability that an opposing agent will choose to defect. The four types of market reputation infrastructures are defined as follows.

Table 2: **Market Reputation Infrastructures**

None	Control. There is no information available about opposing agents in this market.
Observed	RTFT agents are assigned a set of other agents of al types in the market to observe. The game results of the observed agents are reported to the observing agents after each tick. RTFT agents calculate probability of cooperation by dividing the total number of observed cooperations by the total number of observations.
Grouped	Agents are grouped by their strategies, and each group is assigned a random name. A percentage of each group is randomly switched with the other groups. RTFT agents calculate probability of cooperation by dividing the number of cooperations by encountered agents of the same group by the number of encounters with agents of the same group.
Propagated	Upon meeting an agent for the first time, RTFT agents ask other agents in the market about the history of the opponent. Those agents in turn ask other agents, whom in turn ask other agents, etc. RTFT agents calculate probability of cooperating by dividing the total nuber of reported cooperations by the total number of reports.

All agents are initially placed in the None market. Agents must have played a certain number of games in a market before they can leave that market. They will leave the market if their average recent payoff is less than the payoff threshold. The minimum games parameter is also used as the window of games over which averages are calculated. If they choose to leave the market, they will choose a new market at random from the market to which they have not been yet. If they have been to all of the markets, they will choose a market at random.

Table 3: **Parameters**

Name	**Description**	**Default Value**
agentsC	Number of C agents per market.	20
agentsD	Number of D agents per market.	20
agentsTFT	Number of TFT agents per market.	20
agentsRTFT	Number of RTFT agents per market.	20
minGames	The minimum number of games that an agent must play before it can leave a market. Also used to set a window for averaging several metrics.	10
payoffThreshold	The average payoff below which agents will choose to leave a market.	
defectionThreshold	Upon first meeting an opponent, the calculated probability of cooperation below which an RTFT agent will choose to defect against that opponent.	
gdrMixing	The percent of each group that is randomly switched with other gorups in the Grouped market.	
observeCount	The number of agents that an RTFT agent is assigned to observe in the Observed market.	
prFanout	In the propagated reputation market, the number of agents that an RTFT agent asks about the history of an opponent. Also the number of agents that the asked agents ask about the history of the opponent, etc.	
prRecursion	In the propagated reputation market, the depth of recursion of information requestions. For example, a prRecursion of 2 means that the RTFT agent will ask an agent, who will also ask another agent.	

The Prisoner's Dilemma game is set with the following payoffs, which are designed to represent the payoffs of an exchange in a market of goods.

Table 4: **Prisoner's Dilemma**

	Cooperate (C)	Defect (D)
C	1, 1	-1, 2
D	2, -1	0, 0

The case where both agents cooperate is analogous to the successful trading of goods. The good that each had before the exchange was worth less to them (1) than the good that each has after the exchange (2), and thus both agents end up with positive payoffs ($-1+2=1$). The case where both agents defect is analogous to nothing occuring after an exchange has been agreed upon. The state of each agent is unchanged from before the game. The case where one agent cooperates and one agent defects is analogous to the defecting agent saying they will send a good in exchange for the cooperating agent's good, but never doing so. The cooperating agent sends their good to the defecting agent, losing the value that they held for that item (-1) and not getting anything in return. The defecting agent receives the good from the cooperating agent, gaining the value that they hold for that item (2) and not losing anything in the exchange.

4 Experiments

Experiments performed on this model will provide general insight into the dynamics of an e-commerce model in which there are multiple markets with varying reputation infrastructures.

The experiments will also determine which of the reputation infrastructures is best under various conditions. They will determine the conditions under which the results of [1] are replicable, and answer the outstanding question of how robust the concluded rankings of reputation infrastructures in [1] are to parameter variations. In this section, a market has achieved *dominance* when it has at least 80% of the RTFT agents in it after the *minGames* bootstrapping period. Once a market has achieved a critical mass of RTFT agents, it will remain the dominant market.

4.1 Metrics

There are various metrics that will be used to determine the relative effectiveness of reputation infrastructures. None are direct measurements, but are rather indications of effectiveness. The likelihood for dominance of a market depends both on its reputation infrastructure as well as the other agents that are in the market. For example, a market could give an RTFT agent perfect foresight about their opponent, yet if the agent only plays D agents, the RTFT agent will achieve a total payoff of 0, after which it would leave the market. Therefore, while effectiveness may be measured indirectly, no one measure can predict which market will become dominant.

4.1.1 Market Population

When a market has a large population of RTFT agents, it is indicative of an effective reputation system in that market. This has to do with the *stickiness* of the market, which I define as the average time agents stay in the market. Recall that agents stay in a market for the minimum required number of ticks plus as long as their average payoff is greater than the payoff threshold. When all markets are evenly mixed, agents with a relatively low payoff threshold will tend to switch markets often, while agents with a relatively high payoff threshold will tend to switch makets seldomly, but if the payoff threshold is set just so, certain markets may repel agents while other markets hold on to them. When this occurs, the market to which agents *stick* tend to become the dominant market.

4.1.2 Average Payoffs

A straightforward metric for gauging the effectiveness of a market is average payoff. There are two ways of calculating average payoff: one is from the perspective of the agent, and the other is from the perspective of the market.

From the perspective of the agent, average payoff is calculated by averaging the payoffs from the last *minGames*. If the average payoff is lower than the *payoffThreshold*, then the agent will choose to leave the market.

From the perspective of the market, average is calculated by dividing the sum of all payoffs of agents of a certain type in the last tick by the total number of games played by agents of the same type in the last round. This average payoff does not tell the whole story of the market, however. If the variance of payoff between agents of that type is high, then there will likely be more agents of that type that will choose to leave the market. A few bad games could push an agent's individual average payoff below the *payoffThreshold*. Thus, effective markets will also be consistent in the payoffs for RTFT agents.

4.1.3 Average Wrong First Encounters by RTFT Agents

When an RTFT agent first encounters an opponent, they must make an educated guess as to what action the opponent will choose. Guessing incorrectly when playing a D agent means that the RTFT agent will get a payoff of -1 instead of 1, a significant setback. This metric is calculated by dividing the sum of the first encounters where the agent guessed incorrectly by the total number of encounters. A market with a better reputation infrastructure will have a lower value because the average RTFT agent will get higher payoffs on first encounters with D agents.

4.2 Population Dynamics

One of the purposes of this reseach is to examine the dynamics of an agent population in a model with multiple markets. This section descibes an interesting scenario where the parameter values for the markets are chosen to be similar to those in [1]. The parameters for this scenario are as follows.

payoffThreshold	0.5
defectionThreshold	0.5
gdrMixing	0.05
observeCount	5
prFanout	5
prRecursion	3

Figure 1 shows a snapshot of the population of the markets after ticks 0, 10, and 20. Each market is represented by a histogram showing the number of agents of each type currently in that market. The red bar represents the C population, the green bar represents the D population, the blue bar represents the TFT population, and the yellow bar represents the RTFT population.

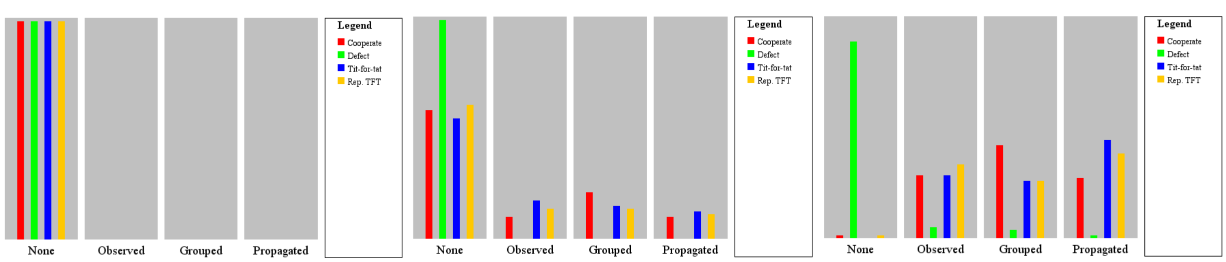

Figure 1: (Left) Tick 0: all agents begin in the None market. (Middle) Tick 10: first chance for agents to leave the None market. (Right) Tick 20: the exodus of non-D agents from the None market continues until there are none left.

All of the agents begin in the None market. The non-D (C, TFT, and TFT) agents start to leave when the minimum number of games have been reached. As shown by the histograms for tick 10, the non-D agents begin to leave the None market, indicating that some but not all of the non-D agents are doing poorly. At this point, all of the D agents remain in the None market, indicating they are all doing very well in this market against this mix of agents. After some of the non-D agents leave the None market, the ratio of D agents to non-D agents increases, and thus the average payoffs of the non-D agents decreases further, which pushes even more of them to leave the market. As demonstrated by the histograms for tick 20, this snowball effect accelerates until all of the non-D agents have left the market. Recall that once agents leave a market the first time, they will visit all of the other markets before returning to the same market.

Figure 2 shows the population of the markets for the same simulation after ticks 25, 90, and 140.

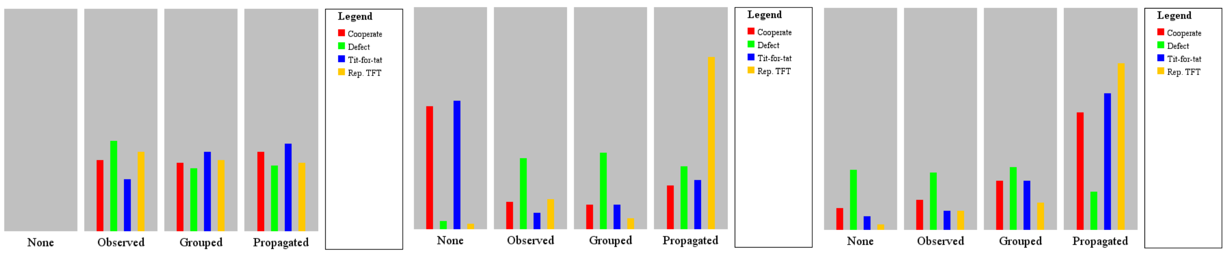

Figure 2: (Left) Tick 25: all of the agents have left the None market. (Middle) Tick 90: many of the C and TFT agents have transitioned back to the None market, the D agents are evenly dispersed amongst the other markets, and the RTFT agents are congregating in the Propagated market. (Right) Tick 140: many of the D agents have transitioned back to the None market, which pushed the C and TFT agents to leave, eventually congregating with the RTFT agents in the Propagated market.

The histograms from Figure 2 show that all of the agents have left the None market by tick 25. After the non-D agents began to leave the market, the average payoffs of D agents began to decrease, which caused some of the D agents to leave the market. When there were no non-D agents left in the market, the maximum payoff became 0, so all of the D agents ended up leaving. While the D agents remain in the None market, the non-D agents have stayed put in the markets to which they have randomly switched. Notice that the populations of non-D agents in the other markets are roughly equal. Because there are no D agents to play against them, their payoffs are guaranteed to be positive, and thus there is never any reason to change markets again.

In tick 90 of the simulation, the D agents have made their way about the other markets, causing many of the C and TFT agents to flee. These markets are safer for RTFT agents because they can make use of the information made available by the reputation infrastructure. The market these agents end up fleeing to is the None market, which the D agents have not yet transitioned back to. In the None market, the C and TFT agents are again guaranteed positive payoffs because there are no D agents present. As the C and TFT agents have transitioned to the None market, the Propagated market has become *sticky* for the RTFT agents. Instead of transitioning back to the None market, the RTFT agents have *stuck* to the Propagated market because the reputation information is good and there are enough non-D agents with which they can interact.

By tick 140, the C, TFT, and RTFT agents have congregated in the Propagated market, while the D agents end up distributed roughly equally between the markets. Between ticks 90 and 140, the D agents transitioned back to the None market, where they found a large number of C and TFT agents to take advantage of, thus causing the C and TFT against to flee once again. This time, however, they found a safe haven in the Propagated market because of the large number of RTFT against with which they could safely interact. It was also the conclusion of the Mui et al paper [1] that the propagated market was the most effective of the markets considered under these conditions.

The conclusion can be drawn from the graphs of Figure 3 that there is a negative correlation between the populations of the non-D agents and the D agents. It is intuitive that D agents cannot exist alone in a market as long as the payoff threshold is above 0 because the mutual payoff of two D actions is 0. Thus, D agents require non-D agents in order to reach their desired level of payoff will "seek" non-D agents. Conversely, and by similar reasoning, non-D agents are repelled by D agents and attracted to other non-D agents. Thus, the hypothesis that D agents will "chase" the non-D agents is correct only before there is a dominant market. When a dominant market arises, the D agents spread out between all of the markets.

The agent population of each market at each time step are illustrated in the following set of graphs. As shown in the None market graph, there is an almost instantaneous drop in the population of non-D agents to 0, followed shortly by an almost instantaneous drop in the population of the D agents to 0. The Propagated market depicts a steady increase in the number of RTFT agents, followed later by a sharp increase in the number of C and TFT agents.

Upon examining the Observed market more closely, one can see that it had a larger RTFT population than the Grouped and Propagated markets following the exodus from the None market. However, the increasing number of D agents corresponded to a decrease in the number of RTFT agents. If the reputation infrastructure of this market had been better, or *stickier*, the RTFT agents would have stayed in that market for longer, which would have attracted more TFT and C agents, and could have turned this market into the dominant market. This was not the case, so we can conclude that the reputation infrastructure of the Observed market is not as good as the Propagated market, which eventually became the dominant market, under these conditions.

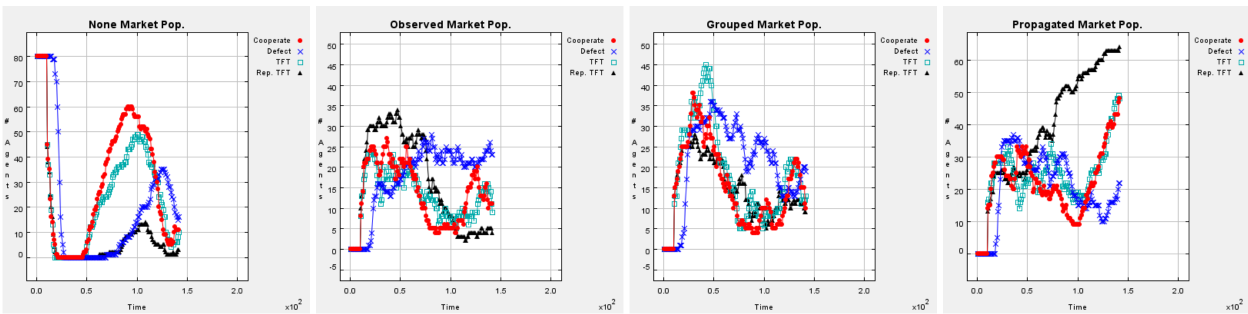

Figure 3: Market populations by agent strategy: (left) None, (middle left) Observed, (middle right) Grouped, (right) Propagated.

The average payoff for each of the agent populations at each time step is illustrated in the following set of graphs. Note that an average payoff of 0 may also mean that there are no agents in the market. In the RTFT agent payoff graph, it is clear that the average for the agents are almost always positive. The only instance of negative payoffs occurs in the None market early in the simulation. RTFT agents appear to get a consistently high payoff in the Propagated market, which improves as the simulation progresses. While the RTFT agents have the best average payoff in the Observed market for about 25 ticks, it does not become the dominant market. The explanation for this is that the D agents were not in the market at the time in any significant numbers, so once they entered the market, the C and TFT agents began to leave, and the average RTFT agent payoffs decreased.

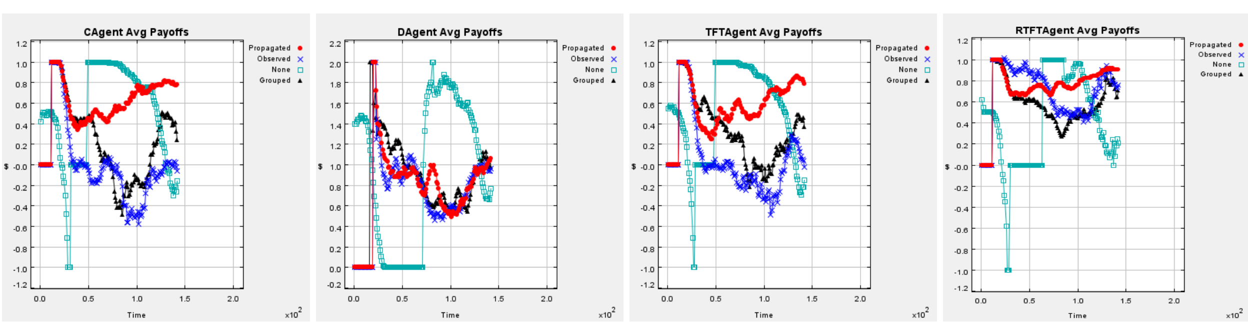

Figure 4: Average payoffs of agent strategies by market: (left) C, (middle left) D, (middle right) TFT, (right) RTFT.

The following is a graph that depicts the average number of wrong guesses made by RTFT agents in each market. Notice that the Propagated market consistently produces the best outcome for this metric. The briefly higher value of this metric in the Propagated market corresponds to tick 25, when the D agents began to move into the Observed, Grouped, and Propagated markets. Once the D agents have been in a market for some amount of time, the RTFT agents have enough information to learn which agents are the D agents, and the metric improves.

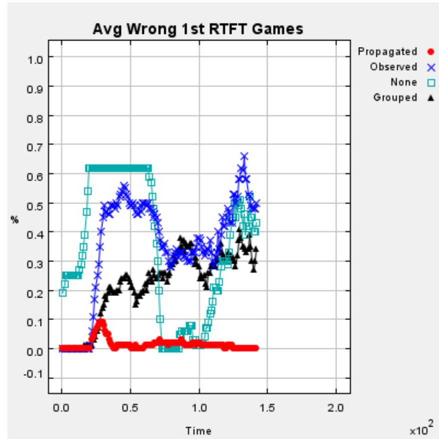

Figure 5: Average wrong first guesses by RTFT agents by market.

Given that a random guess would be incorrect 50% of the time, the Observed market does a surprisingly bad job of informing the RTFT agents. As can be seen in the graph, the average for this metric in the Observed market is around 65% near the end of the simulation. It is especially surprising that this spike occurs near the end of the simulation, where one would expect the RTFT agents to have learned the identities of the D agents.

4.3 Market Dominance

For this section, a sweep of the parameters was run to better understand the relationship between the parameter settings and the outcome of the simulation. The outcome that is being examined is the dominance of a particular market. The following table describes the parameter value ranges that were swept.

payoffThreshold	$\{-0.25, 0.00, 0.25, 0.50, 0.75, 1.00, 1.25\}$
defectionThreshold	$\{0.00, 0.25, 0.50, 0.75, 1.00\}$
gdrMixing	$\{0.00, 0.02, 0.04, 0.06, 0.08, 0.10\}$
observeCount	$\{1, 5, 9, 13, 17, 21\}$
prFanout	$\{1, 3, 5\}$
prRecursion	$\{1, 3, 5\}$

4.4 Payoff Threshold

The payoff threshold decides how "picky" the agents are in regards to their average payoff. What we would expect to see is that when the payoff threshold is relatively low, agents will tend not to change markets, and when the payoff threshold is relatively high, agents will tend to change markets often. Given the market dominance data, we would expect to see that the None market would be dominant when the payoff threshold is low, then non-None market dominance when the payoff threshold reaches some medium point, and finally very few dominant markets at all when the payoff threshold is high. The following is a table of the percent of simulations that resulted in dominant markets as *payoffThreshold* increases.

These results indicate that the hypothesis stated above was correct because we see the expected variations in the data. We can more specifically conclude that when *payoffThreshold* is between about 0.25 and 0.5, a dominant market will emerge from the simulation. This conclusion was only true about 80% of the time, however. Whether a dominant market emerges therefore depends on one or more other parameters about 20% of the time.

	-0.25	0.00	0.25	0.50	0.75	1.00	1.25
No Dominance	72.5%	99%	22.6%	17.7%	99.1%	100%	100%
None	27.5%	1%	16.4%	18.4%	0%	0%	0%
Observed	0%	0%	1.1%	0.8%	0%	0%	0%
Grouped	0%	0%	26.8%	28.8%	0.1%	0%	0%
Propagated	0%	0%	43.1%	34.3%	0.8%	0%	0%

Table 5: Percent of simulations that resulted in dominant markets as *payoffThreshold* is increased.

4.5 Defection Threshold

The defection threshold controls how trusting the RTFT agents are of other agents. When *defectionThreshold*$= 0$, the RTFT agent behaves exactly the same as the TFT agent because it always chooses to cooperate on a first meeting. When *defectionThreshold*$= 1$, the RTFT agent will always choose to defect on the first meeting. The following table shows the percent of simulations that resulted in dominant markets as *defectionThreshold* increases between these two extremes.

	0.00	0.25	0.50	0.75	1.00
No Dominance	70.6%	73.1%	73.5%	69.5%	70.5%
None	25.5%	4.3%	4%	4.6%	8.5%
Observed	1.4%	0%	0%	0%	0%
Grouped	1.1%	7.9%	8.8%	10.2%	15.6%
Propagated	1.4%	14.7%	13.7%	15.7%	5.4%

Table 6: Percent of simulations that resulted in dominant markets as *defectionThreshold* is increased.

As shown in the *defectionThreshold* data, a dominating market rarely emerges when the RTFT agent behaves as a TFT would. It appears that *defectionThreshold* has a very small affect on whether a dominant market emerges compared to *payoffThreshold*, as is shown by the almost constant percent of simulations that did not result in a dominant market as *defectionThreshold* is increased.

4.6 Observed Market

The effectiveness of the Observed market is controlled by the *observeCount* parameter, which is the number of other agents in the market from which the RTFT agents receive game reports. For the *observeCount* parameter, we would expect to see an increase in the percent of simulations in which the Observed market becomes dominant as the value of the parameter increases.

	1	**5**	**9**	**13**	**17**	**21**
No Dominance	71.6%	71.9%	71.3%	71.2%	71.6%	71.6%
None	9.6%	9.1%	9.5%	10.0%	9.4%	9.7%
Observed	0.3%	0.4%	0.4%	0.3%	0.2%	0.2%
Grouped	8.5%	8.7%	8.5%	8.3%	8.5%	8.4%
Propagated	10.1%	10.0%	10.4%	10.3%	10.2%	10.0%

Table 7: Percent of simulations that resulted in dominant markets as *observeCount* is increased.

As is evident from the table, there is no correlation between the *observeCount* parameter and the incidence of dominance of the Observed market for the range of values considered. It appears that the maximum value considered was lower than the minimum value required to make the Observed market reliably dominant.

4.7 Grouped Market

The effectiveness of the Grouped market is controlled by the *gdrMixing* parameter, which determines how randomized the agent groups are. We would expect that as the groups are increasingly randomized, the RTFT agents will have a more difficult time guessing correctly upon first encountering another agent, and that this would result in a decrease in the incidence of the Grouped market becoming dominant.

	0.00	**0.02**	**0.04**	**0.06**	**0.08**	**0.1**
No Dominance	73.3%	72.0%	71.1%	71.2%	70.5%	71.0%
None	8.2%	8.6%	9.9%	9.8%	10.4%	10.4%
Observed	0.2%	0.3%	0.3%	0.3%	0.2%	0.3%
Grouped	16.5%	12.6%	8.7%	6.3%	4.2%	2.6%
Propagated	1.8%	6.5%	10.0%	12.4%	14.7%	15.7%

Table 8: Percent of simulations that resulted in dominant markets as *gdrMixing* is increased.

As is evident by the table shown, the expectations above were correct. As *gdrMixing* increases, the incidence of the Grouped market becoming dominant declines. It is somewhat surprising that only 10% mixing can have such a significant impact on the effectiveness of the Grouped market.

4.8 Propagated Market

The effectiveness of the Propagated market depends on both the *prFanout* and *prRecusion* parameters, which control the breadth and depth of information request propagation. We would expect that as either increases, there would be a corresponding increase in the incidence of domination of the Propagated market.

	1	**3**	**5**
No Dominance	70.1%	72.3%	72.2%
None	8.9%	9.5%	10.2%
Observed	0.3%	0.3%	0.3%
Grouped	11.7%	7.2%	6.4%
Propagated	9.0%	10.7%	10.8%

Table 9: Percent of simulations that resulted in dominant markets as *prFanout* is increased.

	1	**3**	**5**
No Dominance	70.7%	72.4%	71.5%
None	8.9%	9.4%	10.3%
Observed	0.3%	0.3%	0.3%
Grouped	11.8%	7.1%	6.5%
Propagated	8.3%	10.8%	11.3%

Table 10: Percent of simulations that resulted in dominant markets as *prRecursion* is increased.

As is evident from the tables shown, both *prFanout* and *prRecursion* are positively correlated with the incidence of Propagated market domination. It appears, however, that *prRecursion* has a more significant effect than *prFanout*.

5 Conclusions

The purpose of the work presented in this paper was primarily to examine the effects of modeling e-commerce markets differentiated by reputation infrastructure when they exist in the same world simultaneously. A model based on the work in [1] was designed, developed, and experimented upon with a range of parameter values. The first experiment detailed a scenario that replicated the results from [1] for a certain set of parameters, and the following experiments analyzed the effects of varying the parameter values on the outcome of the simulation.

In the course of this work, several key insights have been noted. This work demonstrates the importance of the *network effect* in building a successful e-commerce marketplace. Without the proper community with which an agent can interact, an effective reputation system will not be useful. This work has also shown that the innate parameters of market agents (*payoffThreshold* and *defectionThreshold*), which cannot be known in the real world, make a large difference to the dynamics of the system, so much care should be taken to gauge these parameters in the real world.

The source code for this project is available on the Center for the Study of Complex Systems lab servers under the directory */users/augie/term-project-cscs530/*.

References

[1] Mui L, Mohtashemi M, Halberstadt A. Notions of reputation in multi-agents systems. Proceedings of the first international joint conference on Autonomous agents and multiagent systems part 1 - AAMAS '02. 2002:280.

[2] Neville B, Pitt J. A Simulation Study of Social Agents in Agent Mediated E-Commerce. World.

[3] Yamamoto H. Managing Online Trade by Reputation Circulation: An Agent-Based Approach to the C2C Market. Simulation. 1997.