

The Social Dynamics of Economic Activity in a Virtual World

Eytan Bakshy, Matthew P. Simmons, David A. Huffaker, Chun-Yuen Teng, Lada A. Adamic

School of Information
University of Michigan
Ann Arbor, MI 48103

{ebakshy,mpsimmon,davidhuf,chunyuen,ladamic}@umich.edu

Abstract

This paper examines social structures underlying economic activity in Second Life (SL), a massively multiplayer virtual world that allows users to create and trade virtual objects and commodities. We find that users conduct many of their transactions both within their social networks and within groups. Using frequency of chat as a proxy of tie strength, we observe that free items are more likely to be exchanged as the strength of the tie increases. Social ties particularly play a significant role in paid transactions for sellers with a moderately sized customer base. We further find that sellers enjoying repeat business are likely to be selling to niche markets, because their customers tend to be contained in a smaller number of groups. But while social structure and interaction can help explain a seller's revenues and repeat business, they provide little information in the forecasting a seller's future performance. Our quantitative analysis is complemented by a novel method of visualizing the transaction activity of a seller, including revenue, customer base growth, and repeat business.

1. Introduction

The marriage of social and economic systems has long been of interest to social scientists and has become increasingly important in the monetization of social media systems (Castronova et al. 2009). The economies of virtual worlds have received scholarly attention early on (Castronova 2001) and Second Life (SL) has often been highlighted in the popular media due to the sheer amount of trade activity. In a single month, residents of SL participate in over 40 million transactions, the majority of which represent the exchange of user-created content. In addition to trade activity, we observe a host of other social activities common to many other online communities: frequent chat-based communication; the formation of friendship ties; and the ability to create and join groups based on a variety of common interests. Together, SL provides a rich environment to study the emergence of exchange-based economies in the context of social behavior.

While virtual world economies have been booming, we know little about how such economic activity emerges

through social practices. In this paper, we analyze the patterns in trade of goods and services, both free and paid, within the SL community. We specifically examine the value of operating within groups, chatting with customers, aligning with fellow sellers, and other structural and social mechanisms.

This work is motivated by several research questions: (1) How much economic activity occurs in a virtual world? (2) to what extent is economic activity supported by social interactions? (3) What is the role of groups in creating and sustaining markets? and (4) What social interaction variables best explain seller success?

2. Background

Second Life is a 3D immersive, massively multiplayer virtual world made up almost entirely of user-created content. Linden Lab, the developers of Second Life, support users in the creation of content in several ways. First, they provide a programmable environment and modeling tools to assemble physical landscapes and create objects. Second, they provide a trade system for users to charge for services, and buy, sell, and exchange objects. Finally, they sell real estate – essentially, simulator cycles, where users can interact and deploy objects. In these respects, Second Life is akin to the world wide web: users are provided with a platform where they can develop and deploy content, potentially for profit. And like many Internet communities, users can develop relationships with one another, chat, and create groups.

This has given rise to a social and economic system that resembles the brick-and-mortar world. Individuals can specialize in the production of basic building components, and work together to assemble increasingly complex products and businesses. Entrepreneurs and speculators can build marketplaces and other infrastructure (e.g., shopping malls) to support sellers. Sellers can distinguish themselves by diversifying products, using advertising outlets, and providing immersive shopping experiences (Ondrejka 2008). Groups, clubs and events can serve as additional avenues to support business activity, and are used to sell goods or services and feature new products.

A second important feature of SL's economy is that many items are transferred for free. Sometimes giving away free items can be a purely social or altruistic act, such as sharing content with a friend or a new user. Other times, mer-

chants give away freebies in order to promote other items or events. Freebies are also given away at events or dance clubs as a way to draw in crowds. We expect that social behaviors surrounding the exchange of free goods would differ from that of paid goods. Studies show that giving away code or other virtual “gifts” creates new social relations that are fundamentally different than those that emerge from cost-based transactions (Bergquist and Ljungberg 2001), and that gift-giving often results in reciprocation (Taylor and Harper 2003). There has been previous empirical work on the study of economics in virtual economies. Recently, Castronova, et al. (2009) showed that commodities in massive multi-player online games (MMOGs) resemble that of the real-world economy, including food, clothing, weapons, furnishings and base materials. In MMOGs items are created by the game designers (not the players) and are often required for advancement in the game. Even in absence of “designed” goals, the founders of SL have noted that clothing and accessories, used to customize characters and enjoy the in-world experience, are highly popular in SL (Ondrejka 2004).

Others have focused on variables that increase economic activity in virtual worlds. For example, Malaby (2006) argues that users can increase sales through a combination of connections with others (i.e., social capital) and their credentials or reputation. Similarly, Balasubramanian & Mahajan (2001) argue that economic opportunities arise when virtual groups demonstrate high social interaction and a well-defined focus. Others have also demonstrated the importance of social ties in product adoption (Leskovec, Adamic, and Huberman 2007; Bakshy, Karrer, and Adamic 2009; Aral, Muchnika, and Sundararajana 2009).

3. Description of Data

Linden Lab, the creators of Second Life, provided a dataset in anonymized form spanning different time periods during 2008 and 2009: weekly snapshots of the social and group affiliation graphs, daily chat volume between users, and transactions between users including the time, transaction type, and amount of money transferred. Weekly user-to-user friend lists, user-to-group affiliations, and transaction data cover the period between September 2008 and June 2009. Data on chat volume was available for portions of March-May 2009.

The social network data consists of “friend” lists created by users. SL allows users to specify permissions for each relationship: whether the other user can see when they are online, their in-world location, or even whether they can modify the users’ objects. The network consists of 4.2 million unique users and 43 million unique relationships between users. We only retain those ties in which the “friends” can both see each others’ online status, following the assumption that friends would want to know when they are both in-world. We note that basic statistical features of this thresholded graph, such as the clustering coefficient, resemble that of real-world social networks ($C = .34$) compared to the friendship graph ($C=0.01$) (Newman and Park 2003).

Groups offer a potential means to categorize an aspect of an individual’s interests in Second Life, with different groups corresponding to different interests. There are

520,321 populated groups and close to 23 million group affiliations that were active at some point during the period covered by our data set. Users can be a part of up to 25 different groups at a time in SL. As in many other instances of groups in online networks (Backstrom et al. 2006), groups in SL have a highly skewed distribution in size. For example, while a few dance and club oriented groups had tens of thousands of members, the median size is just 7, and over 14 thousand “groups” had only a single member.

4. Overview of Economic Activity

4.1 Summary

We begin with an overview of the economic activity and basic social behavior of SL users in April 2009. For the purposes of the paper, we focus on transactions types that were either gifts or object transfers, leaving out membership dues, data upload fees, classified ads, etc.*

As shown in Table 1, monthly revenue averages 103,000 Linden Dollars (L\$) (roughly \$413 US Dollars at an exchange rate of 260 L\$ / US Dollar) for all users that receive payments from other users. We will refer to users making money and spending money as sellers and buyers, respectively. Unsurprisingly, there is a long tail in seller activity: while a few sellers have tens of thousands of customers, 67% of active sellers have fewer than 5 customers over the one month period.

In terms of social behavior, sellers on average have both long friend lists and chat with many other users. They belong to an average of 22 groups just shy of the limit of 25 groups.

| | Mean | 50% | 90% | 99% |
|-------------------|--------|-----|--------|---------|
| Free recipients | 13 | 3 | 21 | 85 |
| Free transactions | 49 | 7 | 91 | 363 |
| Pay customers | 17 | 2 | 22 | 253 |
| Pay transactions | 136 | 10 | 183 | 1,670 |
| Revenue (L\$) | 33,191 | 350 | 16,458 | 337,079 |

Table 1: Mean and quantile ranges for transaction behavior of users in April 2009.

4.2 Examining Individual Seller Activity

To provide an intuition for economic behavior on SL, we next examine the activity of a particular seller focused on fashion. This seller has a storefront in SL and specializes in avatar customization, designing a collection of female bodies across several features (e.g., hair style, hair color, eye color, etc.) and ethnicities. She also creates a variety of female clothing including ball gowns, casual clothes, high fashion and lingerie. Here we report on her behavior in April 2009. During this period, she had 773 for-pay transactions with 327 unique customers, meaning that almost half

*For this reason our data represent only a fraction of monthly economic activity published on the web by Linden Lab: <http://secondlife.com/statistics/economy-data.php>, the sample summarized here only includes 29,215,958 for-free and 36,081,807 for-pay transactions in the month of April 2009.




Figure 1: Example of four transaction histories of sellers that started businesses after September 2008. Each row of points represents the amount of money spent in each day by a single customer. Customers are sorted along the y-axis according to first purchase order. Point color intensity corresponds to the logarithm for the total transaction amount. Sellers in (a) and (b) exhibit slowing growth rates and a high number of repeat customers. Sellers shown in (c) and (d) have fewer repeat transactions, but with linear and super-linear growth in customer base.

her sales were from repeat customers. Overall, she made 70 thousand L\$ (approximately \$270 USD). Her customers have on average been a part of the SL community for almost year. She also gave away quite a bit, with 356 zero-cost transactions to 172 users. She likely gives away items as a promotion.

She is quite social, sending 16,114 messages to 2,891 users, and prolific, with an average message length of 551.4 characters, suggesting that she might be broadcasting advertisements to her customers. Although we cannot immediately tell if chatting occurs before, during or after a sale, 107 of the users she chats with are also her customers. She has 238 friends in her buddy list, of which 15 are her customers, suggesting that “friending” users is not in her selling strategy. She had also been a member of 32 different groups, mostly representing boutique shops, fashion and shopping, as well as designer groups (i.e., creating better textures for

SL objects).

In order to track the rate of growth of such a seller’s customer base and identify repeat business, we developed a novel visualization of user transaction activity, shown in Figure 1. Each of the four figures represents the complete transaction history of a seller starting from the seller’s first customer. Customers are ordered according to the time of their first transaction with the seller. Points indicate the days on which customers transacted with the seller, and their intensities correspond to the total amount of money exchanged that day.

The transaction history of the fashion seller is shown in Figure 1(d). We can see that her customer base grew rapidly after January 2009, that these users spent more money per day, and that many of these customers were repeat buyers. A variety of other behaviors can be observed from the figures. For example, in Figure 1(a), we see a seller that hosts

weeklong biweekly events. The seller in Figure 1(b) specializes in club apparel and also exhibits a large returning customer base. This particular seller appears to have closed their business operations over the winter holiday.

5. The Impact of Social Ties on Economic Activity

We note that personal contact is not a prerequisite for the exchange of goods in Second Life. Users can simply select an item from their inventory, target another user or user profile and click “give item”. Or they can create an item and set it to be “For Sale”. Customers can buy such items without directly interacting with the seller. Similarly, users are not required to be “friends” or be on a buddy list in order to access zero-cost items or paid goods. Nevertheless, as we will show, social ties play an important role in a significant portion of economic activity in SL. In this section we examine the impact of social ties, gleaned from communication behavior and explicitly stated friendships, on transaction activity.

5.1 The Social Network

Users can acquire and retain customers through advertising, but they can also utilize the explicit social networking feature (i.e., the buddy graph) to either directly connect with customers or to let word-of-mouth about their goods spread along the social ties of their existing customers. In order to understand how much of a role direct social ties between seller and buyer play in the distribution of content, we measure the proportion of transactions occurring along the social network. As one might expect, free transactions tend to occur more often (39% of the time) on the social graph than do pay transactions, of which only 7% occur between friends. However, we do not find a large difference in the amount of reciprocity, 15.5% vs. 14.4% for free vs. paid transactions, suggesting that the role of free transactions often falls outside of social gift giving.

It would be impractical for a seller to befriend thousands of customers. Indeed, as Figure 2 shows, sellers with thousands of customers are friends with only a small portion of them. However, when the number of customers does not exceed 100, on average about 10% are linked to the seller through the social graph. This suggests that small sellers are more likely to have a personal relationship with their customers.

As mentioned above, a seller can benefit in another way from social networks, and that is via word-of-mouth advertising. We find that 41% of first purchases by a user were “social adoptions”, that is, they were preceded by one of the user’s friends buying from that same seller first. We note that friends frequenting the same seller does not necessarily mean that they communicated about the seller. Rather, it may be due to homophily effects wherein friends are drawn to the same seller independently due to shared interests (Anagnostopoulos, Kumar, and Mahdian 2008; Aral, Muchnika, and Sundararajana 2009).

Nevertheless, the high proportion of social adoptions leads us to believe that users do frequently share informa-




Figure 2: The mean proportion of customers who share a group or are friends with a seller, binned according to the size of sellers’ customer base. Error bars represent 95% confidence intervals.

tion about sellers. To further illustrate this point, we visualize subgraphs of the social graph induced by the set of a seller’s customers, shown in Figure 3. This figure provides an illustration of possible causal paths for seller adoption over the social network.

5.2 Personal Messaging

In the previous section we saw that social ties, expressed in friend lists, correspond to many of the trades taking place in SL. However, our results so far tell us nothing about the strength of those ties. Since friends are easily added, but not all of those ties are subsequently maintained, we next examine more dynamic interaction between users, in the form of chat activity. Chat activity also captures interactions between users who did not codify their connection via the friend graph.

In April 2009, 14 million pairs of users chatted with one another. In order to examine the relationship between communication and economic activity we identified the dyads that had both exchanged a personal message and transferred either a free or for-cost item. Almost half of dyads that exchange free items also chatted with one another (i.e., 7.7 and 3.1 million pairs, respectively), and we find a moderate Spearman rank correlation between the number of messages and zero-cost transactions $\rho = 0.42$.

The relationship between chatting behavior and paid transactions is weaker. Of the 4.5 million dyads that engaged in a paid transaction, a mere 12% chatted. Still, we find a smaller, but significant Spearman rank correlation, $\rho = 0.27$. Therefore while social interaction likely mediates both kinds of economic activity, it’s role is more tightly linked to the exchange of free goods.




Figure 3: Directed acyclic graph representation of possible paths of social contagion amongst customers of a seller. Each node corresponds to a customer. Edges are drawn from a source node to a destination node when the two customers are friends, and the source had purchased before the target. Contagion flows downward along the graph so that predecessor nodes are placed above nodes that they may have influenced.

6. The Impact of Groups on Economic Activity

Much activity, economic and otherwise, is coordinated through groups. Groups allow for one-to-many communication, shared land ownership, and fee collection. Examining the groups to which the top sellers belong to gives us a sense of some of the most profitable activities in SL. Comparing the frequency of words in the names of groups that the top 100 sellers are members of to word frequencies in all group names, we find a high occurrence of words associated with gaming, in particular Zyngo games. This marks gaming as an SL activity where significant amounts of money change hands.

We then probed further to see whether there is some stratification in transactions. A weighted PageRank on the transaction graph of sellers allowed us to rank users not just based on the money they directly take in, but on whether their customers include users who are themselves making money. Looking at the groups of sellers who had high PageRank relative to their revenue, we found high-end one-stop shopping, leasing, and financial services groups, concierge services (provided for large landowners by Linden Lab itself), and other services that experienced users often use to outfit themselves and their businesses. Thus trade in SL is not divided between buyers and sellers, but is rather a tiered ecosystem where users specialize and supply items and services to one another.

To further examine the dynamics between the different kinds of services and activities on offer in SL and their associated revenues, we extracted 8 different sets of groups by keyword: land (*land, real estate*), scripting (*script, build*), gaming (*win, game, zyngo*), concierge, financial (*finan, bank,*

investment), dance (*danc, DJ*), education (*educ*), fashion (*fashion, cloth, jewelry, skin*), and adult (*sex*). We then labeled each transaction with the categories of the groups of seller and buyer, and aggregated to obtain a picture, shown in Figure 4, of the flow of money between different kinds of groups. We observe that groups relating to core economic activity, i.e. land, concierge, and scripting having a high level of transaction activity among their members. Groups relating to non-essential, yet popular, activities, e.g. clubbing, adult, and education, tend to have less economic activity per member and are not as closely tied. The one exception is gaming, which includes games such as Zyngo that generate a great deal of revenue.




Figure 4: Flow of transactions between groups. Nodes represent sets of groups matching topic keywords, sized by av. amount received by members, colored blue to red according to av. amount spent. The edges are directed, with thickness representing proportion of spending by the first set of groups that flows to the others, and darkness corresponding to a normalized transaction amount. Self-loops represent trade occurring between members of the same group.

6.1 Co-grouped exchange

While groups can enable monetary transactions, they are also social structures, ones that are maintained through communication, social ties and reciprocity. To probe their role, we compare the frequency of monetary and free transactions that occur between sellers and buyers who share at least one group. We find that groups play a significant role in both paid transactions, with (48.2%) of them occurring between “co-grouped” sellers, and free transactions, with the proportion at 42.0%. Whether there are particular group structures that are more conducive to free as opposed to paid transactions is an interesting subject for future study.

We next focus on what groups might mean to an individual seller. A seller may choose to provide a completely generic product or service, such that any individual may become a customer. Alternatively, they may choose to specialize and occupy a niche market for goods. A seller with a generic product could potentially gain many customers, but

could have difficulty advertising the product without being able to target a specific audience, and may not be able to rely on word-of-mouth. This could be an especially large handicap in a virtual world setting where mass media play less of a role. The niche seller can advertise more cheaply by targeting the appropriate customers, and can rely on word-of-mouth to spread through group and dyadic communication. However, by definition, their potential customer base is more restricted.

To answer the above question of seller success and their targeting of groups, we measure whether the seller shares groups with his or her customers, and whether the customers belong to the same group or groups. The former captures shared interests between the seller and buyers, and the latter captures whether the customer base fits a particular niche.

Sharing groups with customers can indicate a shared interest, e.g. both the seller and buyer participate in the same role-playing game. In many other cases, for example in the case of the fashion seller described in section 4.2, a seller will create a group that customers join in order to receive notifications about new products. Still in others, sellers may advertise through different groups, e.g. a general fashion group, or a group that represents a niche interest compatible with the seller’s products.

As shown in Figure 2, sharing a group with customers is more scalable than adding dyadic social ties. Even sellers with thousands of customers on average share a group with more than 40% of them. The figure also suggests that groups play a bigger role for sellers with hundreds of customers, who on average share a group with more than half of them, than for smaller sellers who are less likely to share group ties with their handful of customers.

6.2 Group coverage

Even when a seller does not belong to the same groups as her customers, the interests which their goods and services cater to can often be described by the common interests embodied by a small number of groups. That is, when we take groups as means of subsetting users based on a particular shared interest, we can use the diversity of groups required to cover a seller’s transactions as a measure of the diversity of interests present in their customer base. By examining the minimum amount of disjoint groups required to cover 80% of any seller’s customers, we gain insight in to how many separate categories of interest any seller caters to. To determine the 80% coverage data for a particular seller we begin by removing any transactions that do not have an associated transfer of Linden Dollars (i.e., since free transactions may serve non-economic purposes). Given the set of paid transactions for a seller, we count all unique customers that purchase from that seller, and set a threshold (θ) of 80% of that count. We define \mathcal{G} to be the set of all groups that contain at least one customer of the seller’s customers. The goal is find the minimal subset $G \subseteq \mathcal{G}$ whose members exceed the threshold. We note that this is an NP complete set-covering problem, but can be well approximated by the following greedy algorithm:

$$B \leftarrow \mathcal{G}$$

```

 $G \leftarrow \emptyset$ 
while  $|B|/|G| \geq \theta$  do
  select a  $g$  that maximizes  $|g \cap B|$ 
   $B \leftarrow B - g$ 
   $G \leftarrow G \cup g$ 
end while
return  $G$ 

```

We obtain a baseline for comparison by computing group cover for randomly sampled customers. Our expectation is that a random group of customers belong to a more diverse set of groups than sets of customers affiliated with the same seller. To do this, we draw a random sample of customers from the population of all customers over the period of study. Each customer is drawn without replacement with a probability proportional to the total number of paid transactions they were part of. We then compute the same coverage measure for this random sample. This is done for sample sizes ranging from 2 to 4,000, which accounts for the number of customers for 99.8% of sellers. At each sample size, the sampling and coverage computation is repeated 20 times. This gave us a robust measure of the number of groups it would take to cover 80% of any set of customers of size n .




Figure 5: Number of groups required to cover 80% of a seller’s customers as a function of the number of customers the seller has. The null model shows the same but for randomly sampled sets of customers. Sellers are logarithmically binned by number of customers. Exponents represent a linear fit to the averaged data.

Comparing the number of groups required to cover 80% customers to a randomly sampled but equally large group of customers, shown in Figure 5, reveals that those interacting with the same seller are far more likely to belong to the same groups. Furthermore, the number of groups required grows more slowly with a growing customer base for the actual compared to randomized customer groups; to cover 80% of a seller’s n customers requires a number of groups proportional to $n^{0.57}$, while for randomly sampled users it grows as $n^{0.77}$. This shows that sellers are able to operate within niche markets that are represented by a small set of groups.

7. Explaining seller success

In this section we bring together the different aspects of a seller’s interactions mentioned above, including friendship, chat, and group membership as they relate to a seller’s business in one linear model.

Here we look primarily at social and group characteristics based on user activity in April 2009, and correlate them with the users’ business success. We measure a seller’s success in two different ways: whether they have high revenues, whether they enjoy repeat business from the same customers, or whether they are experiencing revenue growth. We limit ourselves to sellers who had 10 or more transactions. Incidentally, we find no correlation between the number of customers a seller has and the average value of their transactions, so it is not true that sellers with many customers are necessarily selling cheap items.

Table 2 shows that Second Life resembles real world associations. The number one predictor of how much a seller makes is how much their friends make. Whether this is a reflection of business partnerships or homophily based on income would be an interesting subject for future research. The size of one’s networks helped to explain revenue, with more active users earning more. The number of friends[†] and number of groups were correlated with revenue at $\rho = 0.337$ and $\rho = 0.200$ respectively. The higher a seller’s revenue, the lower the proportion of customers that they chatted with, but having a direct social tie to them did correlate with increased sales. Finally, higher revenue sellers were more likely serve a customer base that has more recently joined SL (lower age).

| | Estimate | Std. Error |
|----------------------------------|----------|------------|
| (Intercept) | 3.616 | 0.047 |
| av. made by friends [†] | 0.420 | 0.004 |
| % customers chatted with | -2.434 | 0.022 |
| % repeat business | 1.758 | 0.023 |
| mean customer age | -0.319 | 0.011 |
| # friends [†] | 0.166 | 0.006 |
| % customers who are friends | 0.304 | 0.030 |
| group customer cover z-score | -0.134 | 0.007 |
| # of chat partners [†] | 0.223 | 0.008 |
| # of groups | -0.004 | 0.001 |
| % co-grouped sales | -0.007 | 0.021 |

$R^2 = 0.382$

Table 2: Summary of covariates in a linear model predicting seller revenue in order of significance. All predictors are significant at the $p < 0.0001$ level, except for the *% co-grouped sales*, which is not significant.

Although correlations between revenue and social variables were mixed, repeat business was highly consistent with sellers interacting with their customers. We defined repeat business as the fraction of transactions occurring outside of the first 24 hour period following the first observed transaction between a pair of users in the month we stud-

[†] Variable is log-transformed to correct for skew.

ied. The most significant factors, shown in Table 3, were that the seller shared at least one group with their customer, that their customers had been in Second Life longer, and that they chatted with a larger portion of their customer base. All three point to more sustained relationships between sellers and customers. There is also a negative correlation between the number of groups required to explain a seller’s customer base (controlling for the number of customers) and the percentage of repeat business ($\rho = -0.247$). Although this variable was not significant for revenue, it does indicate that niche markets, represented by a small number of groups, enjoy more repeat business than non-niche ones.

| | Estimate | Std. Error |
|----------------------------------|----------|------------|
| (Intercept) | 0.242 | 0.007 |
| av. made by friends [†] | -0.005 | 0.001 |
| % customers who are friends | -0.086 | 0.004 |
| % co-grouped sales | 0.299 | 0.003 |
| # friends [†] | -0.023 | 0.001 |
| group customer cover z-score | -0.018 | 0.001 |
| % customers chatted with | 0.254 | 0.003 |
| # of chat partners [†] | 0.009 | 0.001 |

$R^2 = 0.367$

Table 3: Summary of covariates in a linear model predicting repeat business in order of significance. All predictors are significant at the $p < 0.0001$ level.

Finally, we attempted to predict revenue growth, and failed to do so. Month-to-month revenues were highly correlated (Spearman’s $\rho = 0.800$) but none of the variables capturing a seller’s social interactions were predictive of revenue growth. Indeed, we would not expect to be able to make such potentially profitable predictions using only a simple set of variables.

8. Conclusion

The role of social structures in economic activity has been difficult to study on a large scale, because such structures are normally hidden from view and previously had only painstakingly been observed in small settings. In this paper we present the first large-scale study relating fine-grained social structure to economic transactions. We capitalize on a rich dataset of micro-level economic activity in Second Life, a massive multiplayer virtual world that allows users to create their own objects and share them with others for a price or for free. After showing that the vibrant SL economy has many of the products and markets we see in “real-world” economies, including retail, real estate, entertainment, and financial sectors, we identify a set of characteristics that can be used to connect social and group activity to virtual markets.

In particular, we find that in the virtual world setting, social ties and group membership play a significant role in trade. Not only do groups tend to represent niches in the overall market, but sellers appear to capitalize on this by belonging to the same groups as their customers. Furthermore, we see a positive correlation between communication,

friendship ties and economic activity, and show that these forms of social capital can predict seller revenue and repeat business. We find that direct interaction, via personal messaging, turning friendship lists into customers, and becoming friends with wealthier sellers, all contribute to increased revenue, while personal messaging and finding customers within shared groups contribute to repeat business. Though we were able to identify variables that can explain current revenue, actual monthly growth remains elusive. In future work we plan to uncover the mechanisms that help sellers grow their customer base and encourage repeat business.

Many transactions in virtual worlds are free, with users sharing digital goods they acquire. We find that social ties play a larger role in the trade of free items than in for-fee items. As online businesses continue to debate the success of “free” versus paid content, services and products, our findings show that a large portion of trade will likely occur as consumers share digital goods with their contacts. However, social ties also play a role in generating revenue; becoming a paying customer is frequently due to social adoption; that is, having other friends who are customers of the same seller. In future work we will examine the differences in diffusion of paid and free content over the social network.

This paper establishes an empirical foundation for understanding the relationship between social behavior, including the social ties, group structures and communication networks that users engage in, and economic activity. Our work not only provides more evidence as to the importance of social structures within an economic system, but also shows how groups and social networks are key to the design and deployment of successful virtual communities.

Acknowledgements

We thank Linden Lab for sharing SL data and Erik Hofer and Alena Bacova for providing cluster computing resources for the project. This work was supported by NSF IIS-0746646, MURI award FA9550-08-1-0265 from the Air Force Office of Scientific Research and the Intelligence Community (IC) Postdoctoral Research Fellowship Program.

References

Anagnostopoulos, A.; Kumar, R.; and Mahdian, M. 2008. Influence and correlation in social networks. In *KDD '08: Proceeding of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, 7–15. New York, NY, USA: ACM.

Aral, S.; Muchnika, L.; and Sundararajana, A. 2009. Distinguishing influence-based contagion from homophily-driven diffusion in dynamic networks. *PNAS* 106:51.

Backstrom, L.; Huttenlocher, D.; Kleinberg, J.; and Lan, X. 2006. Group formation in large social networks: membership, growth, and evolution. In *KDD'06*, 54.

Bakshy, E.; Karrer, B.; and Adamic, L. 2009. Social influence and the diffusion of user-created content. In *Proceedings of the tenth ACM conference on Electronic commerce*, 325–334. ACM.

Balasubramanian, S., and Mahajan, V. 2001. The economic leverage of the virtual community. *International Journal of Electronic Commerce* 5(3):103–138.

Bergquist, M., and Ljungberg, J. 2001. The power of gifts: organizing social relationships in open source communities. *Information Systems Journal* 11(4):305–320.

Castronova, E.; Williams, D.; Shen, C.; Ratan, R.; Xiong, L.; Huang, Y.; and Keegan, B. 2009. As real as real? Macroeconomic behavior in a large-scale virtual world. *New Media & Society* 11(5):685.

Castronova, E. 2001. Virtual worlds: A first-hand account of market and society on the cyberian frontier. *The Gruter Institute Working Papers on Law, Economics, and Evolutionary Biology* 2(1).

Leskovec, J.; Adamic, L.; and Huberman, B. 2007. The dynamics of viral marketing. *ACM Transactions on the Web (TWEB)* 1(1):5.

Malaby, T. 2006. Parlaying value: Capital in and beyond virtual worlds. *Games and Culture* 1(2):141.

Newman, M., and Park, J. 2003. Why social networks are different from other types of networks. *Physical Review E* 68(3):36122.

Ondrejka, C. R. 2004. Aviators, moguls, fashionistas and barons: Economics and ownership in second life. *SSRN eLibrary*.

Ondrejka, C. 2008. *The Ecology of games: Connecting youth, games, and learning*. The John D. and Catherine T. MacArthur Foundation Series on Digital Media and Learning. Cambridge: MIT Press. chapter Education Unleashed: Participatory Culture, Education, and Innovation in Second Life, 229–251.

Taylor, A. S., and Harper, R. 2003. The gift of the gab?: A design oriented sociology of young people’s use of mobiles. *Computer Supported Cooperative Work (CSCW)* 12(3):267–296.