

Emergence of Networks and Market Institutions in a Large Virtual Economy*

Curtis Kephart[†] Daniel Friedman[‡] Matt Baumer[§]

University of California Santa Cruz

November 8, 2015

Abstract

A complete set of transactions, more than 40 million within a 1.8 year span, allows us to track the evolution of the trader network and the goods network in an on-line trading community. The computer platform was designed to make barter exchange as attractive as possible; money was not part of the design and all players were created equal. Yet, within weeks, several specific goods began to emerge as media of exchange, and not long after that various sorts of specialized traders began to appear. We track their progress using network-theoretic metrics such as node strength, assortativity, betweenness and closeness. By the end of our sample, virtually all trade was money-mediated and market makers played a major role.

JEL-Classification: B41, C45, D49, E42

Keywords: money as medium of exchange, market makers, virtual economy, market efficiency, network analysis.

*We are grateful to the people at Valve Corporation for financial support, advice, and data. For helpful comments we thank Elle Tian, the moderators at SteamRep.com, and audiences at UCSC, UCI, Berlin Behavioral Economics Seminar, ITAM, and Facebook.

[†]Economics Department, University of California Santa Cruz, CA 95064, Curtis.Kephart@googlemail.com

[‡]Economics Department, University of California Santa Cruz, CA 95064, dan@ucsc.edu

[§]Economics Department, University of California Santa Cruz, CA 95064, mbaumer@ucsc.edu

1 Introduction

How and when do new institutions emerge to facilitate trade, and how can we measure their impact? Such questions are classic but have new urgency in the early 21st century, as markets more tightly bind together economic activity across the planet, and mobile communications enable new ways to transact.

This paper makes a small empirical contribution pertaining to those large questions. In September 2010, Valve Corporation launched a high-performance pure barter trading platform for the user community of one of their more popular games, Team Fortress 2. We have access to *every* transaction from that platform over a 661.4 day interval, involving thousands of distinct types of goods and over 200 thousand different traders.

We analyze those data with several classic questions in mind. Given its best conceivable shot, how stable is barter? Were Adam Smith (1776, Book 1 Chapter 4), William S Jevons (1885) and Karl Menger (1892), among others, correct in predicting that commodity money will emerge to solve logistical problems inherent in barter? Do we see a unique medium of exchange? What roles do divisibility and durability play? Do trade specialists emerge as the market grows, as would seem to follow from the opening argument in Smith (1776)? If so, what kind — dealers (who carry inventory)? escrow brokers (who don't)? lenders? In general, do we see institutions emerge that lower transactions costs? Do we see equilibrium outcomes like the Law of One Price?

Despite recent theoretical elaborations such as Kiyotaki and Wright, 1989 or Ostroy and Starr, 1990), these questions have provoked remarkably little empirical work, then or now. Perhaps the best known is Radford (1945), who showed that cigarettes emerged as medium of exchange in a WW2 POW camp. Much like Radford's prison camp, we see trade in a broad spectrum of goods tends to become intermediated by a small subset of goods which act as media of exchange.

Our empirical investigation is also motivated by network-theoretic questions. What network architectures characterize barter versus monetary exchange? Or direct trade versus intermediated trade? Which network metrics can best demonstrate how goods networks evolve over time? Or how trader networks evolve?

Within the large and heterogeneous recent literature on economic networks (see Jackson, 2010 and Easley and Kleinberg, 2012 for recent overviews), our approach is distinctive in that we deal with (a) large, non-uniform, changing networks rather than static networks that are uniform (e.g., cellular automata) or small, (b) two distinct but interlinked networks (for goods and for traders) defined endogenously by actual transactions data, rather than a single exogenous network or a single endogenous network assumed to have achieved Nash equilibrium.

We are aware of two related empirical lines of research. A series of articles including Kirman (1997) and Kirman and Vriend (2001) study the market for fresh fish in Marseilles. For this differentiated perishable commodity, the authors focus on the stability of trading relationships between a few dozen sellers and several dozen buyers, using a sparse sample of periodic data. We need different techniques to study “continuous” trade (on average, nearly a trade every second) in our much larger networks for exchange of homogeneous durable goods.

Bech and Atalay (2010) use federal funds (overnight) loans to construct a trader network among US banks. They look mainly at a directed unweighted network averaged over time (although sometimes they consider weighted edges or time trends), and confirm stylized facts on node degree distributions. Our network metrics overlap with theirs, but our focus is on undirected weighted networks (although sometimes we consider directed or unweighted networks) and how they change over time. Some of our findings, including those on node degree distributions, stand in contrast to stylized facts established for other sorts of economic or social networks.

Section 2 opens by reviewing relevant aspects of modern network theory, including metrics such as node degree and strength, network assortativity, and betweenness and closeness centrality. It also shows how a set of barter transactions can be used to define a goods network as well as a trader network. The following section briefly describes the data; see Baumer and Kephart (2015) for more details. Section 4 presents results, beginning with an overview of trading volume, whose US dollar value averages well over 2 million per week over the second half of the sample. The analysis of trader networks discloses economically interesting violations of scale-free distributions for node strength, and the emergence of several different sorts of trade specialists. The analysis of goods networks discloses the emergence of

commodity monies — not one, but several. Price dispersion decreases over time but remains substantial. A concluding discussion is followed by supplementary appendices.

2 Theory and Notation

We first review definitions of network metrics, focusing on those appropriate for weighted undirected networks. Then we show how to construct an empirical network from a set of transactions, and sketch the differences we would see in classic monetary versus direct trade, and in direct versus intermediated exchange.

2.1 Network Metrics

Given a network described by an $I \times I$ adjacency matrix $Y = ((y_{ij}))$, the entry $y_{ij} \geq 0$ is called the *strength* of the directed edge (or link) from node i to node j . The network is called *undirected* if Y is a symmetric matrix so that $y_{ij} = y_{ji}$ for all $i, j = 1, \dots, I$. The network is called *unweighted* (or binary) if $y_{ij} \in \{0, 1\}$, i.e., an edge from one node to another either exists (1) or does not exist (0). In any case, by convention all diagonal elements $y_{ii} = 0$, i.e., nodes do not connect to themselves.

Node strength and degree. The *strength* of node i in an undirected network is the sum of its edge weights,

$$s_i = \sum_{j=1}^I y_{ij}. \quad (1)$$

The *degree* k_i of any node i is the number of edges of positive weight that include that node, a nonnegative integer given by

$$k_i = \sum_{j=1}^I \mathbf{1}_{[y_{ij}>0]}, \quad (2)$$

where the indicator function $\mathbf{1}_e = 1$ if event e occurs and is 0 otherwise. In an unweighted network, of course, node strength coincides with node degree.

In a directed network, the expression in equation (1) is called the out-strength of node i , and that the same equation with y_{ij} replaced by y_{ji} is called the in-strength. For reasons explained towards the end of section 2.3 below, when we want to symmetrize node strength

in a directed network we replace y_{ij} in equation (1) by $\max\{y_{ij}, y_{ji}\}$.

Occasionally it is helpful to distinguish nodes with lots of moderately weighty connections from nodes with just a few very weighty connections. For this purpose, following Barrat et al. (2004) and Opsahl et al. (2010), we consider Cobb-Douglas combinations

$$s_{\alpha,i} = s_i^{1-\alpha} k_i^\alpha, \quad (3)$$

focus on the case $\alpha = 0.5$, and compare it to the polar cases $\alpha = 0$ (so $s_{\alpha,i} = s_i$) and $\alpha = 1$ (so $s_{\alpha,i} = k_i$).

Assortativity. Do strong nodes tend to connect directly to other strong nodes, rather than to weaker nodes? A positive answer suggests that the network may be like our galaxy, with a weighty core and gossamer periphery. A sufficiently negative answer, on the other hand, may hint at some sort of specialization, e.g., internet service providers and customers.

An assortativity metric is, in essence, the correlation (across edges) of the strengths of each edge's two nodes. Conceptually, the expression is straightforward:

$$A(Y) = \rho_{[ij]} = \frac{Cov[ij]}{Var[i]} \equiv \frac{E(s_i - Es)(s_j - Es)}{E(s_i - Es)^2} = \frac{E(s_i s_j) - (Es)^2}{E(s_i^2) - (Es)^2}, \quad (4)$$

where the expectation operator E is understood to average over all edges ij .

The concept is easiest to implement in binary (i.e., unweighted undirected) networks, once it is understood that non-existent edges (ij such that $y_{ij} = 0$) are ignored. Newman (2002) proposed replacing node degree by excess degree in such networks, netting out the edge in question, because including it biases assortativity upward. That bias seems at least as important for unweighted networks, so we define $s_{i \setminus j} = s_i - y_{ij}$ as the *excess strength* of node i for edge ij , with expected value $\bar{s} = \frac{1}{H} \sum_{i,j=1}^I y_{ij} s_{i \setminus j}$, where $H = \sum_{i,j=1}^I y_{ij}$. Then the assortativity of an undirected weighted network Y is

$$A(Y) = \frac{\frac{1}{H} \sum_{i,j=1}^I y_{ij} s_{i \setminus j} s_{j \setminus i} - \bar{s}^2}{\frac{1}{H} \sum_{i,j=1}^I y_{ij} s_{i \setminus j}^2 - \bar{s}^2} \quad (5)$$

We have not seen equivalent expressions in the literature (see Noldus and Van Mieghem, 2015 for a recent review). Leung and Chau (2007) uses edge-weighted averages and covari-

ances but not excess strength in defining assortativity for weighted networks. Many authors follow Newman in using excess degree, but only in unweighted networks. A caveat: directed weighted networks, not used below, require a separate definition of $s_{j \setminus i}$.

Centrality. A node is central if it is on lots of shortest paths. To formalize this intuition, define a *path* $p = (n_1, n_2, \dots, n_k)$ as a sequence of adjacent nodes, i.e., nodes satisfying $y_{n_i n_{i+1}} > 0$ for $i = 1, \dots, k - 1$. Let \mathbb{P}_{ij} be the set of all paths from node i to node j , i.e., paths satisfying $n_1 = i$ and $n_k = j$. Define the distance from i to j along path $p \in \mathbb{P}_{ij}$ to be the sum of the reciprocals of the edge weights, $L(p) = \sum_1^{k-1} 1/y_{n_i n_{i+1}}$, and define a *shortest path* from i to j to be any $p^*(ij) \in \operatorname{argmin}\{L(p) : p \in \mathbb{P}_{ij}\}$. A shortest path is generically unique, and the *distance* from i to j is always uniquely defined by $d(i, j) = L(p^*(ij))$. The distance is always positive for $i \neq j$, and is smaller when the shortest path has fewer and weightier links (edges). By convention, the distance is $+\infty$ if $\mathbb{P}_{ij} = \emptyset$, i.e., if i and j belong to different connected components of the network.

Following Brandes (2001) generalization of Freeman (1979), define the *betweenness centrality* of node n as

$$B(n) = \sum_{i, j \neq n} \mathbf{1}_{[n \in p^*(ij)]} / \sum_{i, j \neq n} 1 \in [0, 1], \quad (6)$$

i.e., the fraction of all node pairs ij that have a shortest path that goes through n . The extreme case $B(n) = 0$ occurs for a node with only one edge, and other extreme $B(n^*) = 1$ occurs for the center node n^* of a star-shaped network.

An alternative intuition is that a node is central if on average it has a short distance to other nodes. Freeman (1979) and Newman (2001), define the *closeness centrality* of node n as

$$\tilde{C}(n) = 1 / \sum_{n' \neq n} d(n, n'). \quad (7)$$

A problem for our purposes is that if even one node n' is very weakly connected to other nodes (or is disconnected) then $\tilde{C}(n)$ will be pushed towards (or will equal) zero for all n . This creates problems in our empirical work, so we prefer to use the less standard Opsahl et al. (2010) definition

$$C(n) = \sum_{n' \neq n} \frac{1}{d(n, n')}. \quad (8)$$

As the sum of reciprocal distances instead of the reciprocal of summed distances, (8) is much

less sensitive to the weight of the lightest edge. One can see that $C(n)$ is the sum of harmonic mean weights (divided by number of edges) along shortest paths from n to all other nodes. Thus $C(n)$ will increase as distances shorten, as is desirable, but also as the number $I - 1$ of other nodes increase. So we will normalize it by dividing by $I - 1$.

2.2 Transactions

Our networks are not given exogenously, but rather are constructed from transaction data. We take as given a finite set of *active traders* $A = \{1, 2, \dots, M\}$ and a finite set of *tradable goods* indexed $n = 1, \dots, N$, and consider *bilateral barter transactions* observed over some finite time interval $[0, 1]$. Such a transaction is specified by naming the initiating trader $i \in A$, the counterparty $j \in A$, and the net trade vector $x \in R^N$.

Suppose that trader i initiates net trade x with counterparty j at time t . The convention is that i 's post-transaction holdings $\omega(i, t) \in R_+^N$ are related to her pre-transaction holdings $\omega(i, t-) = \lim_{\varepsilon \downarrow 0} \omega(t - \varepsilon) \in R_+^N$ via

$$\omega(i, t) = \omega(i, t-) + x. \quad (9)$$

Of course, j 's holdings satisfy

$$\omega(j, t) = \omega(j, t-) - x. \quad (10)$$

Using the notation $x_n^+ = \max\{0, x_n\} \geq 0$ and $x_n^- = -\min\{0, x_n\} \geq 0$, the convention can be restated by saying that i trades bundle x^- to j and acquires bundle x^+ in exchange, so $x = x^+ - x^-$ is the net trade vector.

Without loss of generality (just drop exceptions from the lists), we can assume that all M traders transact and that all N goods are traded at least once. Since self-trades and null trades are meaningless, we can assume without loss of generality that $i \neq j$ and $x \neq 0$. Netting out transactions in which a trader both acquires and relinquishes positive amounts of the same good n , we can say without loss of economic content that $x^+ \cdot x^- = 0$. For convenience and with only slight loss of generality, we assume that each price $p_n > 0$, so the *price vector* is a point in the strictly positive orthant, $p \in R_{++}^N$.

Given a price vector $p \in R_{++}^N$, the *value* of the bundle i acquires is $v^+ = p \cdot x^+$ and

the value of the bundle j acquires is $v^- = p \cdot x^-$. The transaction is *budget-balanced* at p if $v^+ = v^-$ or, equivalently, if $0 = p \cdot x \equiv \sum_{k=1}^N p_k x_k$.

2.3 Trader Network and Goods Network

Suppose that transactions $\langle i(t), j(t), x(t) \rangle \in A \times A \times R^N$ are observed at times $t = t_1, t_2, \dots, t_K$, where $0 \leq t_1 \leq t_2 \leq \dots \leq t_K \leq 1$, so the transactions are indexed by k . Given (constant) price vector p , the observed *trader network* is a weighted directed network with node set A . The directed edge weight from node ℓ to node m is the value of the bundles that ℓ acquires from m . Thus the *trader network* is defined by the $M \times M$ adjacency matrix $W = ((w_{\ell m}))$ with entries

$$w_{\ell m} = \sum_{k=1}^K v^+(t_k) \mathbf{1}_{[i(t_k)=\ell] \& [j(t_k)=m]} + v^-(t_k) \mathbf{1}_{[j(t_k)=\ell] \& [i(t_k)=m]}, \quad (11)$$

where $\mathbf{1}_e = 1$ if event e occurs and is 0 otherwise. Thus equation (11) ignores all transactions except those in which ℓ , as initiator or counterparty, acquires goods from m . Of course, adjacency matrix entries are nonnegative and, by convention, diagonal entries are zero. If all trades are budget-balanced, then (11) tells us that the adjacency matrix is symmetric so the trader network is undirected.

The same set of transactions also defines a *goods network*. The nodes of this network are $n = 1, \dots, N$, and the edge weights reflect the value of transactions in which one good is part of the exchange for another. We want the $N \times N$ adjacency matrix $Z = ((z_{nn'}))$ to be symmetric because edges represent *mutual exchange* values of goods — if the value flows from good n to n' for the initiator then it flows the opposite direction for her counterparty, and there is no reason here to privilege one party over the other.

Specifying the edge weights $z_{nn'}$ takes some thought when trades involving n and n' also include other goods. For example, suppose that the value $v_{n'}^- = p_{n'}^- x_{n'}^-$ of good n' constitutes half the value $v^- = p \cdot x^-$ of the vector x^- of goods sent by the initiating trader. Then it seems reasonable to assign half the value v_n^+ of good n acquired by the initiating trader to the edge nn' , and the other half of v_n^+ to other edges nn'' connecting n to the other goods n'' sent by the initiating trader. More generally, given the v 's associated with a trade vector

x , we could assign the weight $v_n^+(\frac{v_{n'}^-}{v^-})$ to the edge nn' when good n is a positive component of x^+ and good n' is a positive component of x^- . Treating the initiator and counterparty symmetrically, we would add the term $v_n^-(\frac{v_{n'}^+}{v^+})$ to account for the case where n' is a positive component of x^+ and good n is a positive component of x^- . Of course, both expressions are 0 when the two goods do not appear on opposite sides of the transaction.

If a transaction is not budget balanced, then the denominators differ in the two expressions and symmetry is lost. To recover symmetry, we adopt the convention that both denominators are

$$v = \max\{v^+, v^-\}, \quad (12)$$

and define the contribution $v_{nn'}$ to that edge weight of a trade x at time t by the equation

$$v_{nn'}(t) = v_n^+ \left(\frac{v_{n'}^-}{v} \right) + v_n^- \left(\frac{v_{n'}^+}{v} \right) \quad (13)$$

Using that expression the goods network matrix entries are

$$z_{nn'} = \sum_{k=1}^K v_{nn'}(t_k). \quad (14)$$

Since a good can't appear with opposite sign from itself, the diagonal entries are zero.

Note that there are several reasons why an observed transaction might not be budget balanced. Convention (12) makes good sense if the main reason is that the record keeping system missed an element of the transaction such as an explicit or implicit promise to repay. If instead the main reason for the imbalance is random noise in goods valuations (perhaps due to trader idiosyncrasies, imprecise pricing, or item indivisibility) then a better convention would be $v = (v^+ + v^-)/2$. We also apply convention (12) to trader networks when we want them to be undirected; the underlying ‘‘phantom value’’ logic is the same.

2.4 Classical Networks

Pure barter can be idealized as transactions in which each good is equally likely to be traded for any other, i.e., v_n^+ and v_n^- are both on average proportional to their value share r_n in the overall economy. Then (apart from sampling error) we would have $z_{nn'} = ar_n r_{n'}$ for some

$a > 0$ and all $n \neq n'$ in equation (14), so the goods network would be completely connected with edge weights proportional to the strength of each node. Assortativity would be near zero, and betweenness and closeness would have roughly uniform distributions assuming that the value shares are roughly uniformly distributed.

At the other extreme, money mediation can be idealized as a single good, $n = 0$ say, such that transactions take either the form $(-m, x^+)$, i.e., the initiator *buys* the bundle x^+ for $m > 0$ units of good 0, or the form $(m, -x^-)$ in which the initiator *sells* the bundle x^- . The corresponding $(N + 1) \times (N + 1)$ symmetric adjacency matrix has $2N$ positive entries, all in row 0 and column 0; the other N^2 matrix entries are all 0. This matrix defines a “star-shaped” goods network around node 0 — a minimal number of edges in the goods network, but each as weighty as its node strength permits. Assortativity would be quite negative, betweenness would be 1.0 for the money good and zero for everything else, and closeness would be high for all goods, especially the money good.

Trader networks can reveal institutions that facilitate transactions. At one extreme, there could be uniform bilateral trading relationships, whereby any trader is equally likely (perhaps after correcting for differing transaction capacities among traders) to trade with any other. As before, this would give us a fully connected graph described by a symmetric matrix $((w_{\ell m}))$ with strictly positive off-diagonal entries, and metrics similar to those just described for barter.

At the other extreme, there could be a universal store, designated as trader 0. If the other M ordinary traders sell their bundles only to the store and buy bundles only from the store (or even conduct barter transactions but only with the store), then the adjacency matrix has exactly the same bordered-zero form just discussed, again representing a star-shaped network. Less extreme forms of intermediation will leave traces in the network metrics, but supplementary analysis will be necessary to distinguish among market makers, speculators and brokers.

3 The Data

As explained in more detail in Baumer and Kephart (2015), our data come from the video game Team Fortress 2 (TF2), sponsored by Valve Corporation. Launched in October, 2007




Figure 1: Item Barter Trading Platform. Here the initiating trader (You) offers a game license and one unit of refined metal and the counterparty (Test4321) offers ten different hats. The transaction occurs when both traders check the central blue box.

using a standard computer game business model (revenues mainly from players purchasing permanent rights to access the game platform), Valve made TF2 “free-to-play” — i.e., zero price for game access rights — in July 2011, gaining revenue from a company store (dubbed the “Mann Co. Store”) where some popular items could be purchased for US dollars or other national currencies. It should be noted that the traded items in TF2 with appreciable value, unlike those in many other games, offer little or no direct advantage in playing the game; they are mainly for show, or to establish identity or to make a style statement.

TF2 took a new turn in September 2011 with the advent of the barter (“Steam Trade”) platform shown in Figure 1; beta versions had been seen a few months earlier. Our data consist of *all* barter transactions beginning in August 2011, when the current accounting system was first introduced, through May 2013 — over 40 million bilateral barter transactions involving nearly 2 million unique trader identities and over 1000 distinct types of items (or goods), gathered in a relational database of more than 0.5 terabytes. Table 1 shows a tiny truncated extract.

Table 1: Example data snippet.

TradeID	PartyA	PartyB	Time	AppID	AssetID	NewID	Origin	EconAssetClass
1	1203	1876	234	440	3881	4120	1	100
2	4256	172	245	440	3942	4136	0	1949
2	4256	172	245	440	4135	4137	1	1585
3	993	8384	250	440	4133	4138	0	207

Party A is the initiating trader i , Party B is the counterparty j for the trade at Time (only last three digits shown) t_k , where the TradeID k appears in the first column. AppID 440 refers to TF2 transactions, AssetID and NewID are tracking numbers for particular units of the item (or good) specified in EconAssetClass. Origin is the indicator variable that the good is a positive element of x^- . Note that a transaction as defined in Section 2 may correspond to several lines in the database, e.g., trade $k = 2$ here consists of the second and third rows.

Thus the data correspond closely to the theoretical structure introduced in Section 2, but with some notable exceptions. Writing $x \in R^N$ suggests that goods are perfectly divisible, but in TF2 each good has indivisible units (each with its own AssetID). Section 2 discussed pure exchange of perfectly durable units in fixed total supply, but in TF2, new units of goods can appear randomly, and some goods can be purchased in the company store. A few goods can be produced by consuming others. In particular, any user can convert 2 junked weapons into 1 unit of scrap metal, convert 3 units of scrap metal to one unit of reclaimed metal or the reverse, and convert 3 units of reclaimed metal into one unit of refined metal or the

reverse. Also, a treasure chest (or crate) can be opened via a purchasable key to produce its (heretofore hidden) contents, and the key and crate are thereby irreversibly consumed. A parallel to these crates are scratch-off lottery tickets: pay the cost (one key and one crate) and receive an immediate update about the prize that you won, with the dollar value of this prize ranging from less than one dollar for the most common up to an estimated value of over \$10,000 for the most sought after items.

The barter trading platform allows trade in items other than those used in TF2, such as the Left4Dead2 game license seen in Figure 1. Of the 70 million transactions we observe, 44 million involve at least one TF2 item, and nearly 41 million involve only TF2 items. Transactions involving non-TF2 items were more common near the end of the period covered by our data, and are excluded from our analysis.

4 Results

Before analyzing how network metrics evolve over time, we take a look at the overall growth of the exchange economy. Figure 2 shows that the number of trades K rose from below 100,000 in the first week to more than 500,000 per week in less than a year, and remained above that level for the rest of the sample. Similarly, the number of unique trader identifiers active each week was approximately 25,000 in the first few weeks and increased to 200,000 within a year, leveling off thereafter. (For the Steam Trading platform as a whole, growth trends continued unabated but, as noted earlier, mainly for games other than TF2.)

Figure 3 shows roughly similar trends for the weekly value of trade $V = \sum_{k=1}^K v(t_k)$. Values are determined using the daily price vector whose construction is described in Baumer and Kephart (2015). The unit of account is a key, which over the entire sample period could be purchased at Valve’s store for US\$2.49 or the equivalent in other national currencies. A substantial fraction of transactions are not budget balanced; indeed 41% are *one way* transactions, with either $v^+ = 0$ or $v^- = 0$. As noted earlier, we assign transaction value $v = \max\{v^+, v^-\}$ when we need to form undirected networks.

Weekly trade value bounces around in the 1 - 1.5M key-equivalent range during the last year, or about US\$2.5 - 3.75 million. Trade value is highly correlated with trade count ($\rho = 0.966$) and with unique trader count ($\rho = 0.938$), but the trade values are less sensitive




Figure 2: Weekly transaction counts and unique trader account over the sample period. Vertical shaded bars indicate promotional events



Figure 3: Weekly values over the sample period. Values are measured in key-equivalents; they can be multiplied by 2.5 to get approximate US Dollar equivalents.

than trade counts to special promotions. (It seems that most promotional items have low prices, and those with high prices have low trade volume. There also seems to be an inflow of new market participants trading common, relatively low value items.)

4.1 Trader Network

Node Degree Distribution. Figure 4 shows the distribution of node degree (k_ℓ = number of counterparties of trader ℓ) in trader networks obtained from all transactions in a single week. Each panel shows the degree histogram in log-log scale. For all 6 weeks shown (and all other weeks as well), most of the histogram declines roughly linearly (in logs), suggesting

that the Pareto distribution (known by physicists as a power law or by network theorists as a scale-free network) dominates here as it does in so many technology and social networks (Barabási and Albert, 1999, Pastor-Satorras and Vespignani, 2001, Liljeros et al., 2001). The slope seems perhaps less steep in later weeks, suggesting that the Pareto exponent γ may decrease over time.

Beginning in Panel C, we see something different and more economically interesting. Starting in early October 2011 a group of traders emerge who trade with an order of magnitude larger set of counterparties. This subpopulation tends to become larger and more disconnected from the main mass over time. Thus we see the spontaneous emergence of large traders, who have thousands of counterparties every week. This mass of large traders is not consistent with the Pareto distribution, so we see that our network at first appears to behave in a fashion to what would one might naïvely expect based on common patterns in social networks but exhibits a departure from this behavior soon after the first day of our dataset.

Assortativity. As a first step in elucidating the economic impact of the large traders, we check trends in assortativity. Figure 5 traces the standard (excess degree) measure of assortativity $A(W^b)$ in the unweighted (“binary”) trader network W^b computed weekly from transaction data. In the first few weeks, assortativity is surprisingly positive, indicating that at first traders with many counterparties tended to trade with others of high degree, and traders with few counterparties tended to trade with each other. The level is comparable to the top benchmarks in Newman (2002), including movie actors. (In the network whose edges indicate whether the actors have appeared together in the same film, big name actors tend to work with other big names, hence the positive assortativity.) The zero assortativity benchmark (using excess degree strength!) is a random graph.

Over the next six months, trader network assortativity plummets towards Newman’s lowest benchmarks, world wide web links (-0.065) and internet wiring (-0.21). (The latter is negatively assortative since individual homes and businesses mostly connect to internet backbones.) The downward trend ends with a 2011 Winter holiday event that brought an influx of new traders. After February 2012, assortativity mostly bounces around in a negative range bounded by the internet wiring and hyperlink benchmarks. Most upward jumps are short lived and coincide with special promotions that attract new participants to the trading




Figure 4: Degree distribution in trader network (log scales).



Figure 5: Weekly assortativity $A(W^b)$ in the unweighted trader network. Horizontal dashed lines show benchmarks from Newman (2002), and vertical shading indicates promotional events that attract new traders.

platform. These sharp increases during the promotions is likely due to the introduction of a variety of promotional goods which, along with the accompanying influx of players due to the events, induce small traders to be willing to search and match to other small traders more often relative to non-promotional event time periods.

What drives these trends? Figure 6 takes a finer grained look, applying equation (5) to two different subsets of the transaction data. The solid red line shows trader assortativity $A(W^{hi})$ for the undirected weighted trader network W^{hi} constructed from transactions more valuable than the median for that week, and the dashed blue line does the same for the network W^{lo} constructed from lower-than-median- v transactions.

In the first few weeks assortativity is between 0 and 0.15 for both subsets. Thereafter, the red line bounces around a modestly upward trend, suggesting that big trades tend to occur mainly between big traders. At the same time, the blue line quickly trends down and eventually settles down in modestly negative territory. Thus it appears that, after the market matures, small traders who want to exchange goods of relatively low value turn to large traders, perhaps specialists, who are willing to accommodate them.

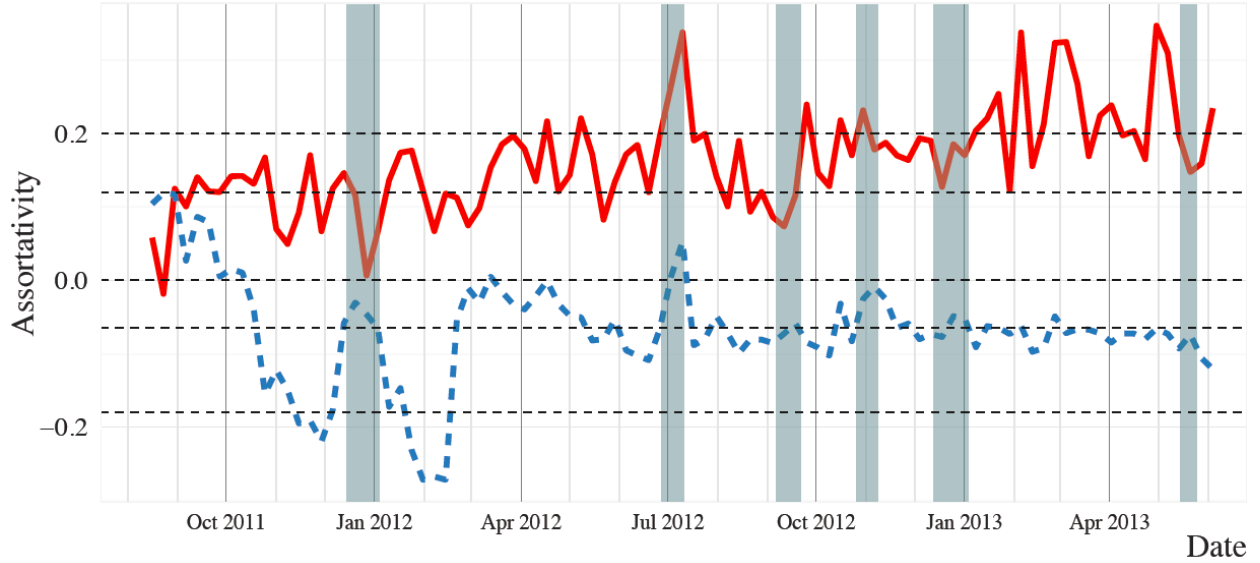


Figure 6: Weekly assortativity $A(W^s)$ for weighted trader networks built from high value ($s = \text{hi}$, solid red line) and from low value ($s = \text{lo}$, dashed blue line) transactions.

Market Makers. To better understand the large traders, we sorted them by characteristics such as weekly transaction count and value, frequency of one-way trades, and profitability. The group that emerged in October 2011 consisted of 88 unique trader id’s, each of whom trade “inhuman” quantities, working 24 hours a day 7 days a week. We call them the Clump because they all move closely together in exploratory animations of trader activity. All evidence (see Appendix A for more details) indicates that the Clump consists of 88 automatons controlled by a single economic entity. The clump initiated over 17.5% of all TF2 transactions in our sample and had an overall gross profit margin (value received minus value delivered divided by the sum of value received and delivered) of roughly 2.1%, with a slight declining trend. Of their first trades with the Clump, 96.9% were one-way inward, delivering value to the Clump. Subsequent trades were commonly one-way and about half were outward, and only about 18% of these were for a good previously delivered to the Clump. We infer that the Clump provides some warehousing services but is primarily an inventory-carrying market maker for a broad range of goods, and that it grants trade credit secured by customers’ deposits.

A second sort of large trader emerged in October 2012 that specialized in 2-way trades. On closer examination, these traders predominantly accepted piles of junked weapons in exchange for metal at or near the conversion rate (2 weapons: 1 scrap metal) available

to ordinary traders. Apparently these specialists are not really exchange intermediaries, but rather are simply offering a convenience on the production side, analogous to CoinStar machines at grocery stores that give dollar bills in exchange for piles of coins. (Baumer and Kephart (2015) note that a sustained depreciation of metals relative to keys began in October 2012. Was that a coincidence? Available evidence is inconclusive.)

Closely associated with this CoinStar entity is a single account that specializes in trading metals for keys and the reverse, beginning in December 2012. Essentially all of its 1,500 counterparties are among the 10,000 accounts that utilized the CoinStar service. We refer to this account as the MoneyChanger for reasons that will be apparent in the next section. We note here that the weightiest edge in the goods network is keys-metals, and that ever since it first appeared MoneyChanger has accounted for 10% of this edge weight. MoneyChanger's weekly average spread between buying and selling prices is about 2%, and it initiated all but one of its 37,000+ trades in our sample. We surmise that MoneyChanger is the primary market maker in the TF2 economy's thickest market.

Another sort of large trader emerged in late December 2012, eventually controlling 6 trader IDs. Of the nearly 400,000 trades involving these IDs, only a few hundred were one way and the vast majority were (within measurement error) budget balanced; gross margin was less than half of one percent on average gross trade value of 1.42 keys. These traders had over 30,000 unique counterparties, who on average transacted more than a dozen times, though the distribution was quite skewed and the median was only three times. Inspection of individual trades suggests a familiar business model: buy and sell goods at a narrow price spread for a range of standard goods, avoiding large inventories. That is, the six IDs are employed by a market maker who uses spot transactions, and does not offer trade credit or take deposits. This spot market maker's share of transactions tended up relative to the Clump, but remained less than half as large, and covered a somewhat narrower range of goods (in particular, they avoided exchange in high value items).

Speculators. A (long) speculator builds up inventory of a particular item whose price he expects to rise, and later sells it off, making a profit if his expectation is correct. (Short speculation seems infeasible in TF2.) To detect speculative behavior, we applied the Wald-Wolfowitz runs test (Bradley, 1968) to a given trader's sequence of transactions in a given good. A run ends, and a new run begins, each time the trader breaks a sequence of con-

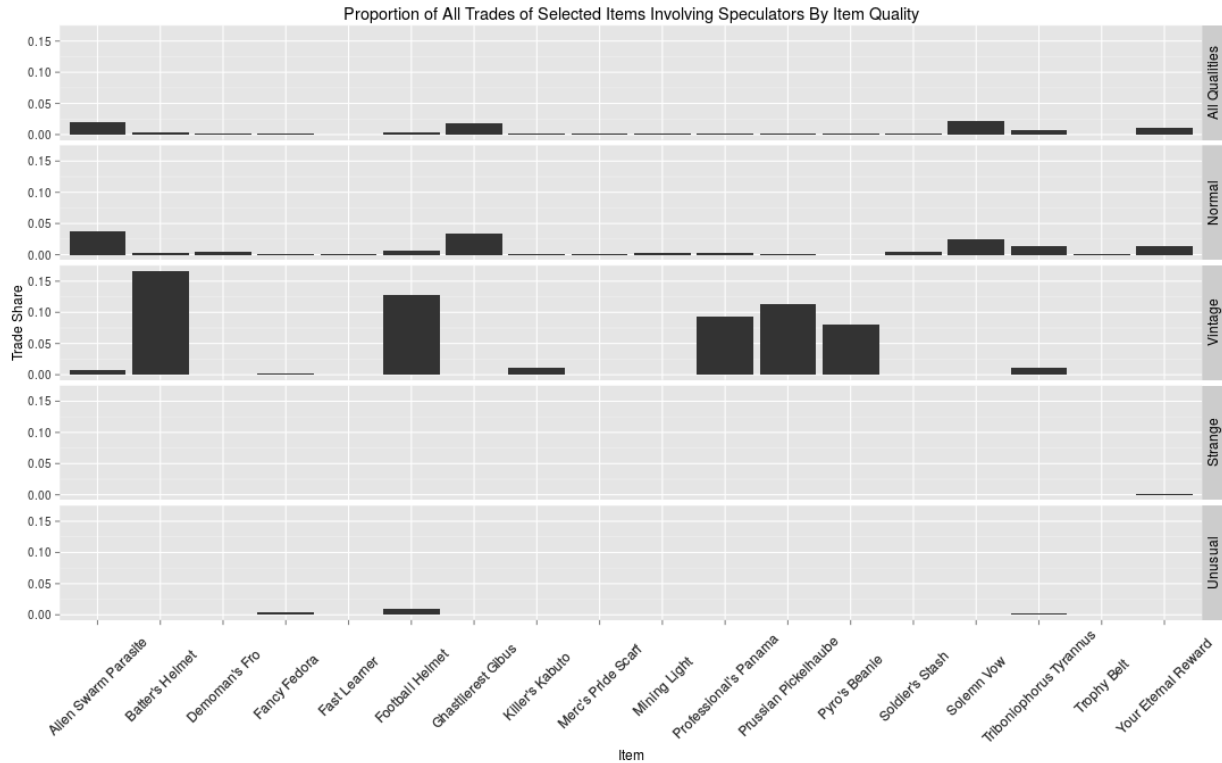


Figure 7: Suspected speculator trade volume as a percentage of total trade volume for selected items. Rows 2-5 report respectively normal, vintage, strange and unusual quality items; top row shows percentage over all four qualities.

secutive buys with a sell, or breaks a sequence of consecutive sells with a buy. We classify the trader as a suspected speculator if the Wald-Wolfowitz runs test Z-score falls in the $p = 0.001$ lower tail under the null hypothesis of exchangeability (essentially, serial independence). That is, suspected speculators tend to have relatively few (hence relatively long) runs, suggestive of inventory accumulation and liquidation.

We examined the 72 items (18 goods each with 4 qualities) shown in Figure 7, chosen because price trends seemed conducive to speculation. The share of trades by suspected speculators is mostly well under 5% of overall trading volume. The only exceptions are five of the vintage quality goods where that share reached 5 to 17%; all five exceptions are attributable to a single account that traded over 1000 of each of those hats but never had a maximum position of greater than 100. So this suspect was not really a speculator, it seems, but perhaps instead was more a hobbyist or an erratic middleman.

So far we haven't found any traders making large profits via speculation; all our evidence

suggests that speculation is not a major TF2 economic activity. This may help explain the puzzle of how it was possible to sustain for so many months a steady depreciation of metals against keys.

Brokers. Another sort of trade specialist facilitates trade of valuable items between two parties who trust the specialist but not each other. For example, trader A may agree to send US\$100 to trader C in exchange for a very special hat. Trader B (a broker or escrow agent) might agree, for a modest fee, to hold the Paypal transfer until C sends him the hat, and then to send C the money and send A the hat. We have no data on Paypal transfers, but can observe B engaging in two one-way trades in rapid succession for the same good; the signature is a short holding time and specialization in particularly high value goods.

We screen for brokers by analyzing the complete transaction history of goods available in unusual quality with a particular effect that tends to command prices of at least 20 keys. We then search for individual accounts which displayed at least twenty episodes of receiving a high-value good in a one-way trade which is then delivered in a second one-way trade within 48 hours. In the first three such accounts we detected, we very conservatively estimate the value of brokered exchange at 5000 keys or US\$12,500; the median holding time was 7 minutes.

Our sampling of unusual grade goods so far has detected tightly-defined brokerage in 9 to 15% of all of the trades involving these high value goods trades. We conclude that brokerage plays a substantial but not dominant role in TF2's markets for high-end goods.

4.2 Goods Network

For goods networks, the main questions are whether traders eventually abandon barter in favor of monetary exchange, and whether transactions costs decline substantially as the market matures.

Strength. As a first clue on barter versus money mediated exchange, we consider item strengths as defined in equation (1). Figure 8 graphs the strengths of every item in the TF2 goods network, normalized so that the strengths of all items always sum to 1.0. Putting aside for the moment the top line, we see that Keys account over 20% of the trade value by the end of the sample. Both Earbuds and Refined Metal are about as important as Keys in early




Figure 8: Weekly normalized item strength s_i in weighted goods networks.

2012, but by the end of the sample each is around 10%. (Earbuds traded for about 40 units of refined metal at the start of the sample and over 140 units by the end, so refined metal’s high strength also indicates very large trade quantities.) These three items are strongest, but another three also show consistent strength: Bill’s Hat and Reclaimed and Scrap Metal. The other items that show up on the graph are Max’s Severed Head (actually a hat), the Hat of Undeniable Wealth and Respect (HOUWAR), as well as blips for special event keys and crates.

Given the fact that metals are convertible in both directions at a fixed rate of three lesser-quality metal to one of next higher value type, it makes sense to combine the three grades of metal into a single composite. We tentatively define money as the combination of metals, keys, earbuds and Bill’s hats. Although relative prices can vary among these four components, we will see later that they all serve as media of exchange.

Simply summing the four (or six, counting the metal grades separately) components’ strengths, we approach the 50% benchmark for idealized monetary exchange. Does this mean that barter has disappeared? Not necessarily; some of the sum comes from “currency market” trade, i.e., from edges within the set of six nodes. A better indication of the extent of monetization comes from collapsing these six items into a single node and calculating its




Figure 9: Weekly assortativity $A(Z)$ in goods networks. Connected orange dots and green triangles indicate respectively the unweighted (binary) full network and reduced (single money node) network values, while dashed blue squares and dashed purple indicate the corresponding weighted network values.

strength in the resulting goods network, as shown in the top red line of Figure 8. We see that the six money items jointly account for about 30 to 34% of value flows in the reduced network — quite a lot but substantially below the benchmark.

For an alternative perspective, see Figure 16 in Appendix A1. It shows the $\alpha = 0.5$ Opsahl et al. (2010) strength for all items in the economy, combining node degree and node strength. As one might expect, the lower denomination currencies look more important with this metric, and there is an overall upward trend due to increasing number of distinct traded items over time.

Assortativity. Figure 9 shows that unweighted goods networks have assortativity $A(Z^b)$ that remains very close to zero, but that assortativity $A(Z)$ in the weighted networks, computed using equation (5), is surprisingly negative, even compared to the internet wiring benchmark. We just saw that the six tentative money items have weighty edges with each other, so we again collapse them into a single composite good and obtain extremely negative assortativity, around -0.5. This is a very strong hint that other goods tend to trade with this composite good, so it may indeed be the main medium of exchange.

Betweenness. The most definitive evidence on money come from the betweenness

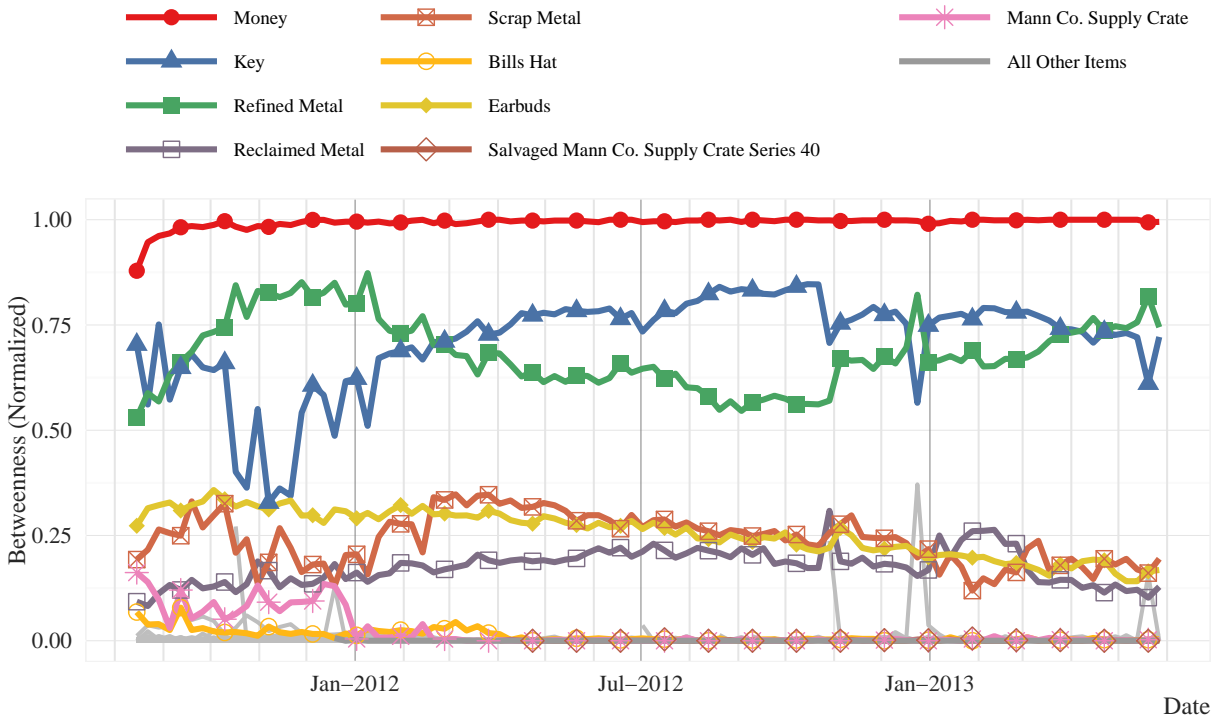


Figure 10: Weekly betweenness centrality $B(n)$ in weighted goods networks. The top line pertains to the network with six nodes collapsed to one (“Money”) node; other lines pertain to the full goods network.

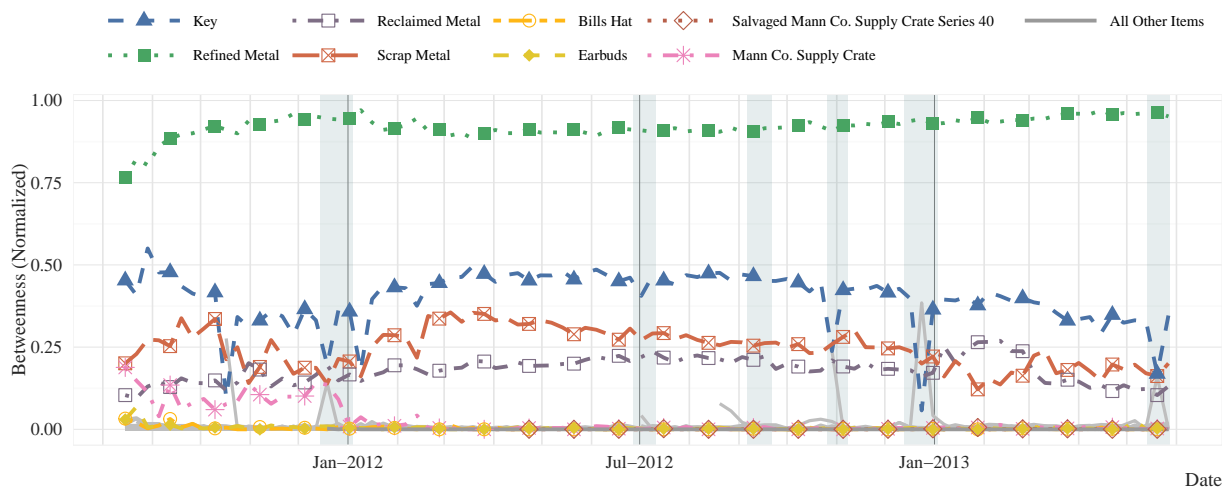


Figure 11: Weekly betweenness centrality $B(n)$ in weighted goods networks constructed for transactions involving only items valued below 0.95 keys.

metric $B(n)$. Figure 10 shows that $B(n)$ is essentially zero for the vast majority of goods. All long-lived exceptions are among the tentative money goods, especially keys and refined metal, each with betweenness usually in the 60-80% range. (Short-lived exceptions are mostly event keys and crates, which seem to be substitutes for ordinary keys as media of exchange — their up spikes in Figure 10 coincide with down spikes for keys, especially around the 2012 holidays and the end of the sample.)

Once again, we reconstructed the goods network with the same composite money good as before. The graph shows that even in the first week over 85% of trade by value went through composite money, and within a few weeks, virtually all trade did so for the rest of the data sample. We interpret this as conclusive evidence that the TF trading platform completed its transition to monetary exchange by October 2011.

Metal, especially refined metal, is the main sort of money used in the low value transactions. Figure 11 shows that keys also play a role here, and that special event keys are close substitutes when they appear. For mid-tier items, Figure 12 shows that keys have dominated since late 2011; refined metal plays a supporting role that diminishes over time. For top-tier items, Figure 13 shows that earbuds and Bill’s hats come into their own, though keys are almost as important early on and eventually attain the highest betweenness centrality even in this segment.

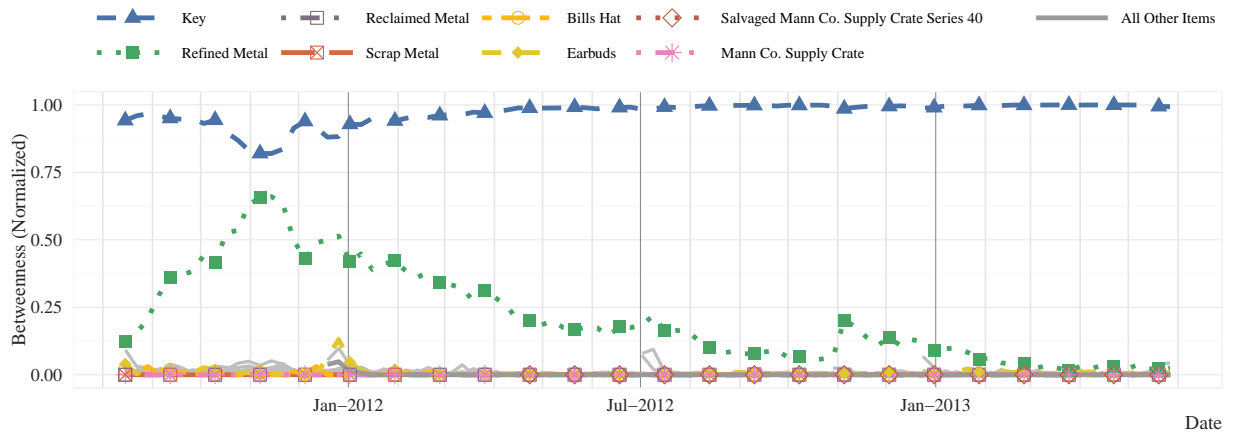


Figure 12: Weekly betweenness centrality $B(n)$ in weighted goods networks constructed for transactions involving only items valued valued between 0.95 and 5 keys.

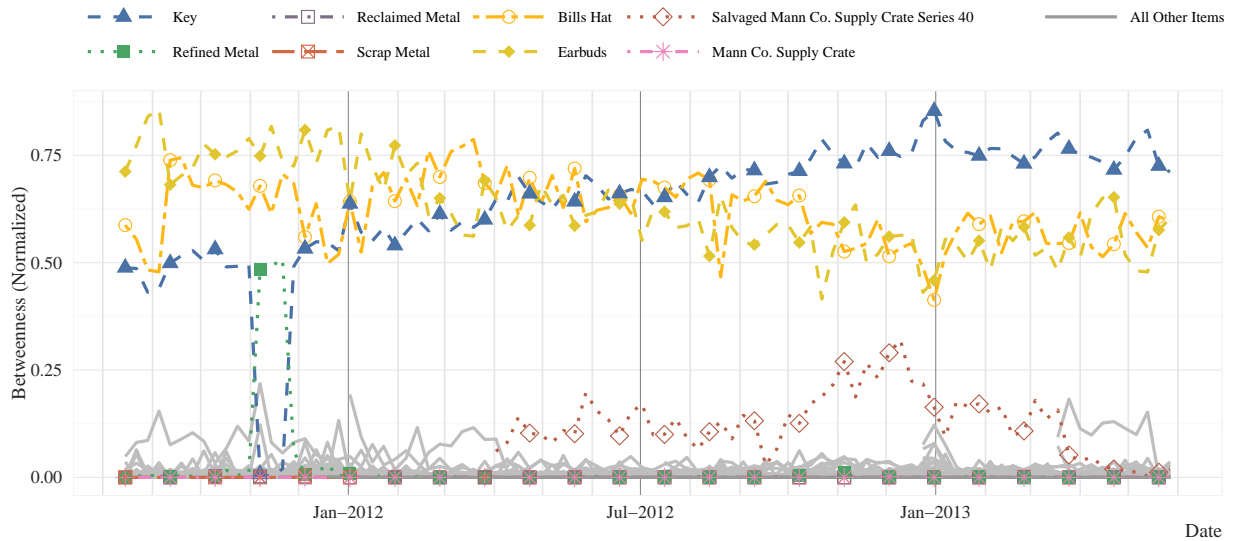


Figure 13: Weekly betweenness centrality $B(n)$ in weighted goods networks constructed for transactions involving only items valued above 5 keys.

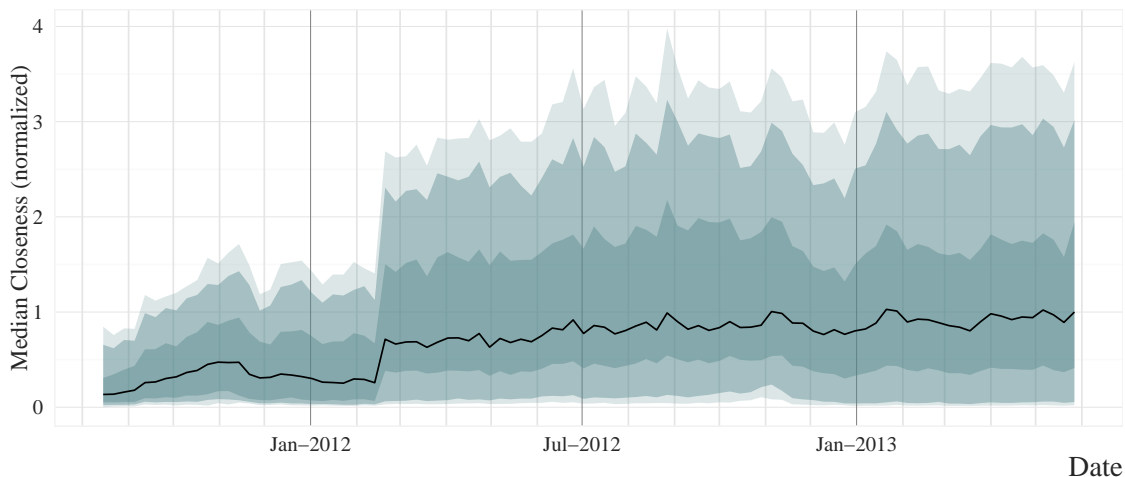


Figure 14: Weekly normalized closeness distribution for the weighted goods network, computed via equation (7) divided by the weekly number of goods, and scaled so that median is 1.0 in the final week. Line is median and shaded ribbons emanating from median span 25th - 75th, 10th - 90th, and 5th - 95th percentiles.

The question now remains, why are there six money goods rather than one? Some relevant evidence appears in Figures 11 - 13, which break down the goods networks by the maximum price of the non-money goods (other than the six money goods) involved in the transaction.

Closeness. Figure 14 shows the distribution of normalized closeness $C(n)$ on a weekly basis. The distributions are usually unimodal and don't have especially long or fat tails. (By contrast, the betweenness centrality distributions are very skewed as $B(n)$ is quite large for a handful of goods and very close to zero for everything else.) Between the first few weeks and the last median closeness increases seven fold while the shape seems to change little. The big jump in mid February 2012 seems to be due to mainly to allowing trade in previously untradeable items and improvements in the user interface; the number of trades and traders nearly doubled at this time. Overall, the Figure shows that shortest paths between goods became weightier and shorter over time, i.e., it became easier and easier to trade one arbitrary good for another.

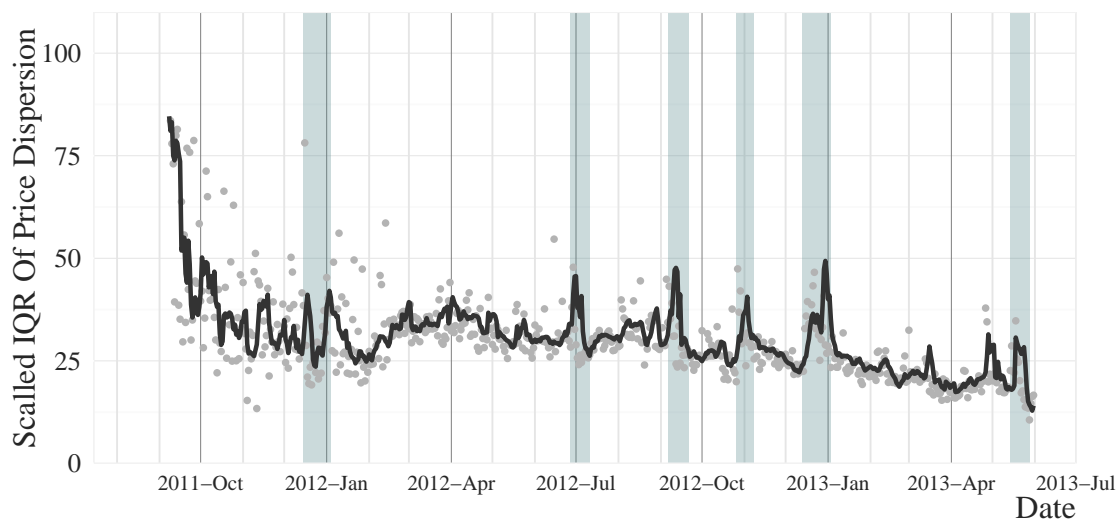


Figure 15: Overall price dispersion as a percentage of mean price. Gray dots show daily value-weighted averages of normalized interquartile range, and black lines show trailing value-weighted 7-day averages, over all items with at least 100 price observations in previous 30 days. Vertical shading indicates promotion events.

4.3 Price Dispersion

Did the emergence of money and specialist traders improve efficiency? More specifically, did transactions costs decline and prices become more unified? In a monetary economy, the most direct measure of transactions cost is the spread between bid and offer prices, the current cost of a round-trip transaction. That is, the direct measure is the net loss (as a percentage of the mean price) when you sell an item at the highest bid price and immediately repurchase it at the lowest ask price. Since bid and ask prices are not part of our data set (nor do they generally exist in TF2), we need an observable proxy for transactions costs.

As explained in the Appendix, we believe that SIQR, the interquartile price range scaled by the median price, is a good proxy for transaction cost, as well as a robust direct measure of price dispersion. It also aggregates well across goods, so we take its value-weighted average as in Figure 15.

The Figure shows that overall SIQR is quite high in the Summer of 2011, but by Winter 2011 it declines to under 50 percent, and is mostly in the 25-35 range in 2012. We infer that a typical (in terms of value) round trip trade would return only about $100 - 75 = 25$ percent of its original value in early months, but would return 65-75% in 2012. The overall SIQR

spikes briefly during promotion events, probably due to the influx of new traders and new goods of unclear value. In 2013 SIQR declines modestly and mostly remains below 25, and in the last few days of our sample it falls to about 16, suggesting that round-trip costs that are about one sixth of original value.

The emergence of the Clump in October 2011, and its sudden disappearance for ten days at the end of July 2012, had no discernible effect on SIQR. As noted earlier, the thickest bilateral market is keys-for-metal, and there SIQR drops below 10 percent in October 2011. From January 2012 until the end of our sample, the SIQR for that money conversion rate mostly bounces around in the 3-8 percent range, while that for earbuds is mostly around 5% and that for Bill's hat is mostly around 12%.

5 Discussion

What conclusions can we draw from our results? Although Valve engineers created a trading platform that was entirely egalitarian, we found that several sorts of specialists soon emerged to facilitate trade. The node strength distribution in the trader network grew a longer and fatter upper tail, and then calved off sets of very large traders with different specialities and different business plans.

One set of nodes evidently was controlled by a single entity (we call it the Clump) that traded actively in a wide range of commodities and maintained modest inventories. We see the Clump as a classic intermediary. It maintains a modestly profitable spread between buy and sell prices and offers trade credit secured by deposits. An apparent competitor emerged later, an intermediary that also made markets for low- to mid-value goods via spot transactions. Another specialist (MoneyChanger, as we call it) made the market for key-metal exchanges, the weightiest edge in the entire TF2 economy. Evidence so far suggests a substantial role for brokers (or escrow agents) in the market for high-end goods, but little role for speculators. The dramatic drop in trader network assortativity in the first year suggests that, taken together, the specialist traders indeed facilitate trade, reducing small traders' search costs and frictions.

The results on goods networks are equally enlightening. Although Valve designers created perhaps the most efficient barter platform that the world has ever seen, the evidence

indicates that nevertheless indirect, monetary exchange soon evolved. Betweenness metrics show that the composite money good was already quite prevalent by the time our data sample begins, and became essentially universal in exchange within a few more months, consistent with classical economists' writings on money.

On the other hand, the classical economists predicted that a single medium of exchange would prevail, but we found that the money composite in TF2 consisted of 6 distinct goods (or 4, if one lumps together the three different grades of metal). In our interpretation, the multiplicity of commodity monies is due to the indivisibility of TF2 goods. Indivisibility limits the competition between low- and high-denomination commodity moneys; it is awkward to trade dozens or hundreds of units of a low denomination money for a valuable good, or to make change when paying for a cheap good using a unit of a high denomination money. Thus a high denomination money good may be a complement rather than a substitute for low denomination money.

A future avenue we hope to explore involves the collapse of earbuds which started around the beginning of 2015, dropping from a value of approximately 30 keys down to their current estimated value of 5 keys. We propose that Valve's introduction of a centralized dollar denominated posted-price marketplace replaced the use of earbuds as the preferred medium of exchange for high-value items and thus the value-in-exchange of earbuds became equal to their value-in-use as a cosmetic item which was much less than their exchange value when they were used as a primary currency.

Indeed, although only one good, the dollar, is money in the US, it is also true that coins and bills are indivisible. Four popular denominations (quarters, 1-dollar bills, 5's, 20's) span two orders of magnitude in value. Likewise, in TF2 trading, proto-money goods may compete within each denomination range — Max's Severed Head, HOUWAR, earbuds and Bill's hats seem to have competed with each other in the high range with only the latter two surviving as media of exchange, but none of these items seemed to compete with metals in the low range.

The fairly steady increase over time in the closeness metric suggests easier trading as TF2's market institutions matured. More direct evidence comes from our measure of price dispersion, SIQR, which also serves as a proxy for transactions costs. The overall value of SIQR dropped sharply during the first year from around 75% to well under 50%, and

by the end of our sample was below 25%. Even in the thickest markets, SIQR remains substantially above the 0% level implied by a strict Law of One Price. Our interpretation is that indivisibilities remain important for low to medium value goods, and that markets are thin for high-end goods.

Each economy, including the modern global economy that we all inhabit, has its own peculiarities, and one must be cautious in generalizing. Our paper contributes some new evidence on how economies can self-organize, a new data point to combine with those already available. This new data point may be especially useful because it comes with unprecedented detail on transactions, and is relatively independent of those already known by historians. We cannot conclude too much from the TF2 economy, but it does add new support (and new caveats) to classical perspectives on money, and to the view that institutions emerge spontaneously to reduce transactions costs and facilitate trade.

Much work remains. Towards the end of our sample, the Steam Trading platform supported considerable trade for virtual goods for games other than TF2, and some trades crossed the boundary between different games. We conjecture that these data may provide a new perspective on international finance questions, especially those concerning what happens when previously separate economies begin to interact with each other. Valve and its user community both continue to innovate, so the story continues.

Our main technical contributions are to propose new ways to construct two different networks from barter transactions data, and to adapt existing network metrics to describe how these networks evolve. We hope that our readers are inspired to further refine and extend these metrics.

6 Appendix: Supplementary Analyses

Simplifying networks. Given a general directed weighted network described by an $I \times I$ adjacency matrix $Y = ((y_{ij}))$, we can construct an undirected weighted network by replacing y_{ij} by $y_{ij}^s = \max\{y_{ij}, y_{ji}\}$ for all $i, j = 1, \dots, I$, whose adjacency matrix Y^s is symmetric. Likewise, we can construct a directed unweighted network $Y^b = ((y_{ij}^b))$ from Y by replacing all positive entries (or alternatively, all entries exceeding some positive threshold value) by the entry $y_{ij}^b = 1$.

Happily, the operation $Y \mapsto Y^s$ commutes with the operation $Y \mapsto Y^b$. That is, we arrive at the same undirected unweighted network Y^{sb} whether we first symmetrize using $y_{ij}^s = \max\{y_{ij}, y_{ji}\}$ and then binarize using the indicator function, or first binarize then symmetrize. This is not true if we symmetrize using $y_{ij}^m = [y_{ij} + y_{ji}]/2$. Hence, besides the empirical advantages mentioned in the text, the max procedure also has a theoretical advantage, which can be helpful given the ubiquity of metrics (such as node degree) based on simple undirected unweighted networks.

Node Degree and Strength. Looking at strength without considering degree potentially could be deceiving. For example, in goods networks, some non-money items may exhibit relatively high weights – possibly because they are very popular or important to economic life. There may also be goods that have high degree because they appear on the other side of trade for many other goods, and yet not really be money, akin to pocket lint that often accompanies coins. Therefore it seems worthwhile to combining degree and strength, as in Opsahl et al. (2010) modified degree with tuning parameter α set to 0.5.

Figure 16 shows Opsahl et al. (2010) strength for all items in the economy. Composite Money (in the reduced goods network) is strongest, but all six of its constituents also exhibit considerable strength. (The weakest of them, scrap metal, tracks the strongest other item, Max’s Severed Head, fairly closely.) Keys emerge as the strongest constituent, with refined metals and earbuds vying for second place.

The clump. Of the 88 members of the Clump, less than 0.00003 of their transactions are initiated by non-clump traders or are between two clump traders, and these exceptions are confined to a couple of days and have low value. Over our sample the clump transacted

Opsahl et al (2010) Strength

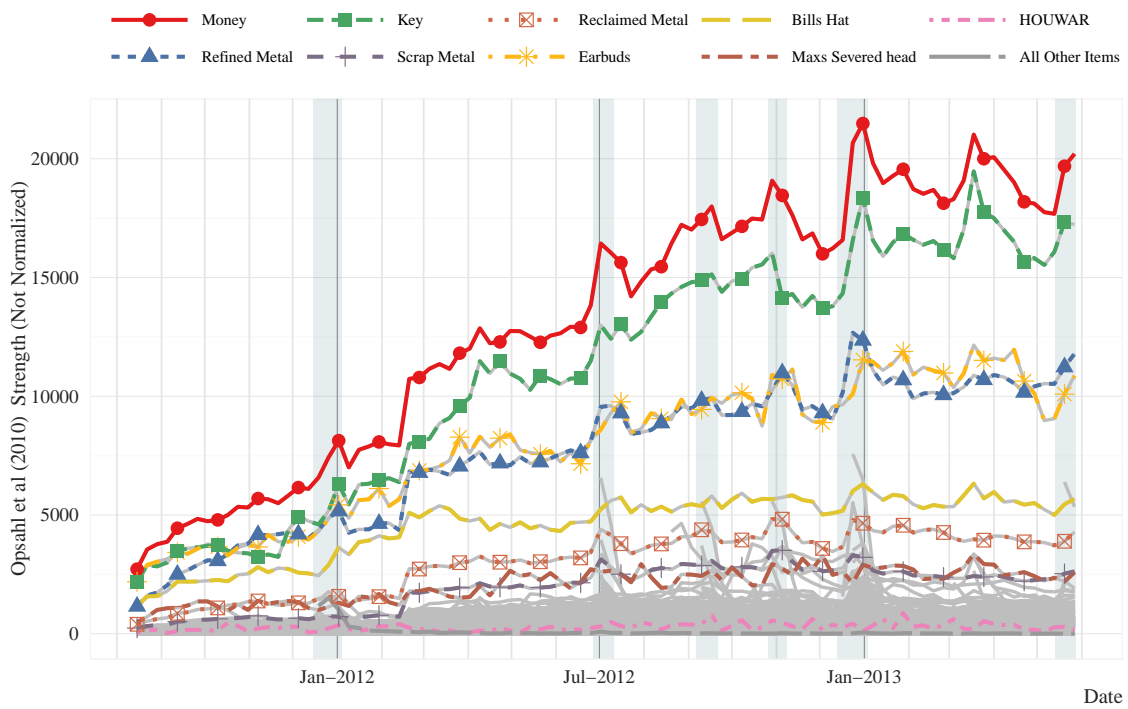


Figure 16: Weekly strength $s_{\alpha,i} = s_i^{1-\alpha} k_i^\alpha$ for $\alpha = 0.5$. Money strength is computed for the reduced network (collapsing Money's six constituent nodes); all other strengths are for the full goods network.

7 million times with over 135 thousand unique counterparties, accounting for over 17.5% of all TF2 trading. Overall gross profit margin (value received minus value delivered divided by the sum of value received and delivered) was roughly 2.1%, with a slight declining trend. Most of its transactions, 94.7%, are one-way, and 91.7% of counterparties trade more than once. Of their first trades with the Clump, 96.9% were one-way inward, delivering value to the Clump. Subsequent one-way trades are increasingly likely to be outward; by the second trade 55% withdraw value. As the number of trades with the Clump increase the percent of trades that withdraw value approaches two-thirds, but only about 18% of these were for a good previously delivered to the Clump by that trader.

As an inventory-carrying market maker for a broad range of goods, what value does the Clump provide to customers? In a post on TF2 forum, an apparent customer explained that it “is fast and straightforward. -Prices set up upfront. -Don’t have to delve into a forum/server looking for someone having what i want. -Don’t have to chase a user i want to trade with. -No unnecessary [sic] haggling/price changing/offer changing/trade requests during the trade. If buying [at relatively] high [price] is what i have to pay for the convenience of automated trading, so be it. I’m not on it for the benefit, but for the hats. I consider it a price for the service offered.” (Steam User Forum, October 31 2012)

Brokers. Our value estimate is conservative because unusual quality items are not uniform — there are several different forms of unusual. Our single estimated price for unusual items understates the price of the most desirable sorts, which are more likely to be brokered. For example, burning flames is one sort of unusual effect, and recently a burning flames hotties hoodie was valued at over 250 keys while we priced it at 40 keys, the median across all unusual hotties hoodies. We thus believe that our median price estimates are in fact lower bounds for the valuations of these thinly traded unusuables since lower value items tend to be traded more frequently. Because the services of a third party broker are more likely to be requested for relatively valuable items, we believe that the true total value flow that has been mediated by brokers may actually be much larger than the 5000 key estimate which we present.

Denominations of money. Exchange rates differ from day to day, but over the second half of our sample, Bill’s hats traded for about 8 keys and earbuds usually for 21- 27 keys. This roughly 3:1 ratio is less than the 4:1 or 5:1 ratio for popular coins and bills, but there is

a possible historical precedent for a compressed ratio. According to Wikipedia, “In the Great Recoinage of 1816, the guinea was replaced as the major unit of currency by the pound and in coinage with a sovereign. Even after the coin ceased to circulate, the name guinea was long used to indicate the amount of 21 shillings (£1.05 in decimalised currency). The guinea had an aristocratic overtone; professional fees and payment for land, horses, art, bespoke tailoring, furniture and other luxury items were often quoted in guineas until a couple of years after decimalisation in 1971.”

Price dispersion as a proxy for transaction costs. Once most transactions go through a money good, it becomes much easier to detect and arbitrage price discrepancies. Indeed, if anyone posts (or even hints) that they are willing to buy at a price higher than the price at which someone is willing to sell, then anyone aware of the two prices could accept both offers and pocket the difference.

This standard argument implies that every accepted ask price observed over a short period of time is above all accepted bid prices. Assuming equal numbers of the two sorts of transactions, we conclude that all prices above the median are accepted asks, whose median thus is at the 75th percentile, while the median of accepted bids is at the 25th percentile. Hence their difference, the interquartile range, is a proxy for round-trip transaction cost. To maintain comparability across goods, it makes sense to express the interquartile range as a percentage of median price, so for item n at time t , we define

$$\frac{SIQR_{nt} = 100 * [p_{75,nt} - p_{25,nt}]}{p_{50,nt}}, \quad (15)$$

where $p_{z,nt}$ is the z^{th} percentile of the imputed prices associated with good n and time interval t . We concede that $SIQR_{nt}$ may somewhat overstate the round trip cost, but there is no reason to think that the degree of overstatement changes systematically over time or across goods. Of course, by definition, $SIQR_{nt}$ is a robust and direct measure of price dispersion.

To aggregate SIQR across goods, we take the value-weighted mean adjusted for sample size,

$$SIQR_t = \frac{\sum_n \eta_{nt} SIQR_{nt}}{\sum_n \eta_{nt}}, \quad (16)$$

where the weight $\eta_{nt} = p_{50,nt} k_{nt}^{1.5}$ is the square root of the number of transactions k_{nt} (to capture the sample precision) times a robust estimate of the relevant transaction value

$p_{50,nt}k_{nt}$. The sum is taken over all goods n for which the time interval $[t - 30, t]$ includes at least $k_{\min} = 100$ transactions involving good n . To aggregate $SIQR_t$ across time, we simply take a simple trailing average.

Our results are robust to a variety of other choices of k_{\min} ; of course, lower choices of k_{\min} generally result in choppier time series. Also, we find similar trends in SIQR when we replace $k_{nt}^{1.5}$ in the definition of the weight by $k_{nt}^{0.5}$, as would be appropriate if, instead asking how much should you expect to lose on a round trip for an *item of typical value*, you asked the question for a *typical trade of whatever size*.

Tensors and Networks. Suppose that each trader i is characterized by a vector of traits, $v^i \in R^I$. Here it may be convenient to treat the traded bundles separately, replacing notations $x^+, x^- \in R_+^N$ by $b^i, b^j \in R_+^N$ to denote the bundles acquired by the two traders via barter swap. This enables us to give a more complete description of a transaction as an element $T_{j\ell}^{ik} \in R^I \otimes R^I \otimes R_+^N \otimes R_+^N$. That is, transactions are tensors of order (or mode) 4, comprising descriptions of the initiating trader, the bundle she acquires, the counterpart trader, and the bundle he acquires.

This space of tensors has dimension L^2N^2 , as opposed to the N dimensional transaction space $A^2 \times R^N$ used in the text. (The space A as written is finite and discrete, hence dimension 0, but it is possible to embed it in an interval, in which case A^2 would add two more dimensions.) For example, with $I = 7$ relevant characteristics and $N = 1,000$ goods, the tensor space would have dimension of roughly 50 million, while the cartesian product space would have dimension of only about 1 thousand.

The extra dimensions in the space of tensors allow us to capture arbitrary interactions among the characteristics of both traders and elements of both traded bundles. That sort of richness can be important in many applications but it is not yet clear to us how important it would be for our application even if we knew traders' demographic info etc. In fact, we currently only know their trading history, so it seems infeasible to take this approach.

If using tensors turns out to be useful at some later stage, we can take advantage of existing techniques to greatly reduce the dimensionality: cluster analysis, and principal component (and singular value) decompositions generalized to tensors.

7 Bibliography

References

- Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. *science*, 286(5439):509–512, 1999.
- Alain Barrat, Marc Barthélemy, Romualdo Pastor-Satorras, and Alessandro Vespignani. The architecture of complex weighted networks. *Proceedings of the National Academy of Sciences of the United States of America*, 101(11):3747–3752, 2004.
- Matt Baumer and Curtis Kephart. Aggregate dynamics in a large virtual economy: Prices and real activity in team fortress. *University of California Santa Cruz Working Paper Series*, 2015.
- Morten L Bech and Enghin Atalay. The topology of the federal funds market. *Physica A: Statistical Mechanics and its Applications*, 389(22):5223–5246, 2010.
- James V Bradley. Distribution-free statistical tests. 1968.
- Ulrik Brandes. A faster algorithm for betweenness centrality. *Journal of Mathematical Sociology*, 25(2):163–177, 2001.
- David Easley and Jon Kleinberg. Networks, crowds, and markets, 2012.
- Linton C Freeman. Centrality in social networks conceptual clarification. *Social networks*, 1(3):215–239, 1979.
- Matthew O Jackson. An overview of social networks and economic applications. *The handbook of social economics*, 1:511–85, 2010.
- William Stanley Jevons. *Money and the Mechanism of Exchange*, volume 17. 1885.
- Alan Kirman. The economy as an evolving network. *Journal of evolutionary economics*, 7(4):339–353, 1997.
- Alan P Kirman and Nicolaas J Vriend. Evolving market structure: An ace model of price dispersion and loyalty. *Journal of Economic Dynamics and Control*, 25(3):459–502, 2001.

- Nobuhiro Kiyotaki and Randall Wright. On money as a medium of exchange. *The Journal of Political Economy*, pages 927–954, 1989.
- CC Leung and HF Chau. Weighted assortative and disassortative networks model. *Physica A: Statistical Mechanics and its Applications*, 378(2):591–602, 2007.
- Fredrik Liljeros, Christofer R Edling, Luis A Nunes Amaral, H Eugene Stanley, and Yvonne Åberg. The web of human sexual contacts. *Nature*, 411(6840):907–908, 2001.
- Karl Menger. On the origin of money. *The Economic Journal*, 2(6):239–255, 1892.
- Mark EJ Newman. Scientific collaboration networks. ii. shortest paths, weighted networks, and centrality. *Physical review E*, 64(1):016132, 2001.
- Mark EJ Newman. Assortative mixing in networks. *Physical review letters*, 89(20):208701, 2002.
- Rogier Noldus and Piet Van Mieghem. Assortativity in complex networks. *Journal of Complex Networks*, page cnv005, 2015.
- Tore Opsahl, Filip Agneessens, and John Skvoretz. Node centrality in weighted networks: Generalizing degree and shortest paths. *Social Networks*, 32(3):245–251, 2010.
- Joseph M. Ostroy and Ross M. Starr. The transactions role of money. In B. M. Friedman and F. H. Hahn, editors, *Handbook of Monetary Economics, Ed 1*. North-Holland Publishing Company and American Elsevier Publishing Company.
- Romualdo Pastor-Satorras and Alessandro Vespignani. Epidemic spreading in scale-free networks. *Physical review letters*, 86(14):3200, 2001.
- Adam Smith. An inquiry into the nature and causes of the wealth of nations: Volume one. 1776.