Shout!: Design and Analysis of an Online Marketplace for Retweets

by Ambika M. Krishnamachar
B.S., Massachusetts Institute of Technology (2015)
Submitted to the
Department of Electrical Engineering and Computer Science
in Partial Fulfillment of the Requirements for the Degree of

Master of Engineering in Electrical Engineering and Computer Science
at the
Massachusetts Institute of Technology
September 2016

Author:

_____

Ambika M. Krishnamachar
Department of Electrical Engineering and Computer Science
June 10, 2016

Certified by:

_____

Cesar A. Hidalgo, Associate Professor of Media Arts and Sciences
Thesis Supervisor
June 10, 2016

Accepted by:

_____

Dr. Christopher Terman, Chairman, Masters of Engineering Thesis Committee

Shout!: Design and Analysis of an Online Marketplace for Retweets

Ambika Krishnamachar

Submitted to the Department of Electrical Engineering and Computer Science in Partial Fulfillment of the Requirements for the Degree of Master of Engineering in Electrical Engineering and Computer Science at the Massachusetts Institute of Technology

Abstract

In a world where attention is limited, popularity is an asset that allows those endowed with it to command attention on demand. Popularity, which we can approximate as the number of contacts in a person's network, allows journalists to share their stories with wider audiences, musicians to promote their creations to more fans, and entrepreneurs to secure more crowdfunding. However, since most modern social networking platforms treat popularity as non-tradeable and private, people are unable to leverage the popularity of their peers in their marketing efforts. If users were able to access the social networks of their close friends, they could multiply their reach without expending the effort normally necessary to build an expansive network. Here I present Shout!, a platform that allows friends to act as a group to coordinate their social media presence. Shout! is an online marketplace for retweets we launched in 2016. With Shout!, users can set up micro-contracts with their friends to exchange future retweets. Shout! allows the user to trigger these retweets through their friends' accounts when they need them. Shout! provides value to its users by allowing them to trade their social capital. In this thesis, I examine the meaning of social capital and the link between social media interactions and more traditional forms of capital. I then describe the design of Shout! and the major decisions made when building it. Finally, I evaluate the use of Shout! by early adopters and use the data collected to explore open research questions, such as the relative prices of retweets between people with different levels of popularity. Based on preliminary analysis, we find that prices of retweets remain within a small range even as popularity levels have much more variation. We also find that existing online friendship seems to be a strong factor in deciding who to trade with on Shout!. I explain the implications of our findings and outline our plans for future improvements to the platform.

Introduction

May I borrow your social network?

In July 2015, Dharmesh Shah, the CTO of Hubspot (a technology start-up in Boston), posted an article on LinkedIn and told the marketing department to circulate it widely.[1] Employees were suddenly bombarded with emails urging them to retweet and share the link to the article. They even received messages containing pre-made tweets, so that with minimal effort, they could send the message out through their own social media accounts. Why did Hubspot make this big effort to get employees to share their message? Why bother employees instead of just promoting the message themselves? What was wrong with Hubspot's social media account?

The problem with Hubspot's social network, and with most companies' social networks, was its limited reach. With every employee that retweeted the article, the company was effectively borrowing that employee's social network and extending its own reach. The alternatives would have been to pay for a sponsored tweet or hire people (or bots) to actively share the message on social media.[2] However, by utilizing its own employees' social networks, Hubspot was taking advantage of a prescreened set of qualified users, each with their own valuable contacts. This extended network was far more likely to be interested in Dharmesh's LinkedIn article than the network that would have been generated by sponsored tweets or bots. Even if a similarly high quality network could have been generated through alternative means (like a marketing company), building such a network would have incurred both time and monetary cost. Instead, by borrowing its employees' social networks Hubspot was accessing the value in its employees' *social capital*.

For our purposes, social capital can be understood as a person's personal social network. Social capital has value that impacts society across every level, from close friend groups to large organizations. In Hubspot's case, a group of key technology executives viewing and circulating a LinkedIn article could boost the company's visibility, increase its growth, and contribute to higher revenue.

The Benefits of Social Capital

Though not popularized until more than 80 years later, the notion of social capital was first introduced by L.J. Hanifan, the state supervisor of rural schools in West Virginia in 1916.[3] Hanifan was very concerned with the importance of community involvement for the success of schools. He noticed that the accumulation and coordination of human capital, necessary components to form a strong business corporation, seemed to parallel the necessary components to form a strong social community. Hanifan defined social capital as "those tangible assets [that] count for most in the daily lives of people: namely goodwill, fellowship, sympathy, and social intercourse among the individuals and families who make up a social unit." The rural communities Hanifan worked in were plagued by problems such as low literacy rates, school absenteeism, and poor infrastructure, something he saw as a result of this virtually nonexistent social glue. Hanifan spent a year trying to build this social capital by implementing community center meetings, evening classes, school fairs, libraries, and athletics, with huge gains in the recreational, intellectual, and even economic conditions of the community. For example, during community meetings, members figured out how to reallocate inefficient spending and pooled together $250,000 to improve their roads.

Half a century after Hanifan's success story, Pierre Bourdieu also noted that it is economically beneficial to create and maintain an expansive social network.[4] Bourdieu, a renowned sociologist and public intellectual, was among the scholars who focused their work in social order and the dynamics of power in society. The traditional definition of capital as "wealth in the form of money or assets," is purely economic.[5] Bourdieu challenged this definition by presenting social capital as a legitimate form of capital. In *The Forms of Capital*, Bourdieu argues that the traditional notion of capital is too narrow, since it reduces the meaning of all human interactions to mercantile exchange. In his introduction of cultural and social capital, he argues that all forms of capital are related and can be derived from one another through transformation. In his work, Bourdieu asserts that capital is actually present in several forms (economic, cultural, and social):

> *"...capital can present itself in three fundamental guises: as economic capital, which is immediately and directly convertible into money and may be institutionalized in the forms of property rights; as cultural capital, which is convertible, on certain conditions, into economic capital and may be institutionalized in the forms of educational*

qualifications; and as social capital, made up of social obligations ('connections'), which is convertible, in certain conditions, into economic capital..."

The findings of Bourdieu, Hanifan, and other intellectuals are consistent with real-world evidence of the benefits of social capital. In a 2013 survey of over 2,000 adults either currently working in or actively seeking a job in the US Labor Force, 4 in 10 people reported finding their favorite job through a personal connection, and 64% of company recruiters agree that their greatest source of job talent comes from personal referrals.[6]

Even in less developed economies, social capital has been leveraged by microfinance institutions that hold communities of people accountable for each other. These institutions provide loans to people who otherwise do not have access to banking-related services; however, interest rates are often extremely high due to both the administrative cost of lending and significant default rates. In Bangladesh, microfinance institutions organized groups in communities that were liable for each others' loans. The social pressure to repay loans within these circles decreased the frequency of defaults, driving down the interest rate and increasing the economic benefit to the communities.[7]

Social Capital and Social Networks

There is ample support for the claim that social capital has economic value, but how is social capital acquired? Social capital is embodied in the set of personal connections that link people together. In other words, social capital is embodied in social networks. Links between people come in various forms, which directly relate to different types of social capital. Robert Putnam was a political scientist who studied social capital and the nature of civic society. Putnam defined two categories of social capital that correspond to different patterns in social networks.[8] He defined 'bonding' social capital as strong intra-group ties within a (usually homogenous) set of people, or "sociological superglue". This stands in contrast to 'bridging' social capital, which allows different social groups to link and share information. Bonding and bridging social capital operate in conjunction, as one without the other has negative effects (for example, pure bonding within a community without bridging can lead to the group being excluded from the wider society). Bonding and bridging capital can manifest themselves in different forms. Because bonding social capital comes as the result of strong intra-group trust, it can enable members of a tight-knit friend group to borrow money from one another, where the collateral is trust. In contrast, bridging social capital often encourages information flow and connection, such as learning about job opportunities.

The way people form and maintain social networks is continuously evolving. In the last two decades, the growth of the internet and the advent of social media have shifted a large fraction of our interactions online. These technologies have made the world more accessible, but they have also created differences in the social networks of today. The differences in dating patterns over the last few decades provide a powerful example of this shift. In 1932, a sociologist named James Bossard looked at data from five thousand marriage licenses to see where people tended to find their significant others. He discovered that a full third of married couples had lived within five blocks of each other, and one-eighth had even lived in the same building.[9] Contrast this with dating patterns of today, where social media and dating apps are leading directly to the construction of new relationships and bonds. A third of couples today meet through a dating app.[10] This evidence would suggest that social media has a powerful effect on our social networks, and potentially reduces the importance of physical proximity between people to form bonds. What do these shifts mean for social capital?

In Putnam's famous and controversial work *Bowling Alone*, he argues that social capital is plummeting in America due to these changing social networks. Putnam describes the advantages of living in a community rich with social capital:

> *"For a variety of reasons, life is easier in a community blessed with a substantial stock of social capital. In the first place, networks of civic engagement foster sturdy norms of generalized reciprocity and encourage the emergence of social trust. Such networks facilitate coordination and communication, amplify reputations, and thus allow dilemmas of collective action to be resolved..."* [11]

Putnam cited a variety of examples of how civil engagement was on the steady decrease. His iconic example was that Americans are increasingly going to bowling alleys alone, rather than with organized groups. He points out many areas where participation is declining, such as parent-teacher organizations, voting, and volunteer groups. Putnam explained that even familial bonds are looser than previous generations, and that much fewer people socialize with their neighbors.

In general, civil engagement within people's close geographical communities is certainly on the decline, as evidenced by Putnam's many studies. However, the impact of the Internet and online social networks on our communities may not be so clear-cut. Researchers have challenged Putnam's

claim that social capital is declining by showing that the Internet can facilitate the formation of communities, although the structure of these communities is changing. When people communicate online without needing to be in close physical proximity, a different set of communities emerges, centered around shared interests rather than geographical closeness.

Even before the advent of social media websites, in the early years of the Internet, people communicated with each other through online discussion groups and forums.[12] These groups were centered around particular topics, and were often a resource for specific types of information (for example, a group for hardware hackers to troubleshoot, or for doctors to share knowledge). However, information exchange is only a small piece of the benefits people received from online communities. There is evidence of people joining discussion groups not only to transfer knowledge, but to also to share in the support of a group. For example, many online groups emerged that provided some form of communal support, on topics ranging from alcoholism recovery, to companionship forums for women in computer science, to weight loss groups where group members encouraged each other to meet their goals. These groups provided its members a sense of community and belonging, even if the members never met in person.

Barry Wellman was a sociologist who focused the majority of his research in how computers and the internet impacted social networks in communities and organizations. He noted that the internet, with its ability to connect the world, gave rise to these new types of communities, which were different than traditional communities, but also beneficial. Wellman explained the structure of these emerging communities:

> *"People on the net have a greater tendency to develop feelings of closeness on the basis of shared interests rather than on the basis of shared social characteristics such as gender and socioeconomic status. So they are relatively homogenous in their interests and attitudes just as they are relatively heterogenous in the participants' age, social class, ethnicity, life-cycle stage and other aspects of their social backgrounds. The homogeous interests of virtual community participants can foster high levels of empathetic understanding and mutual support."* [12]

Wellman's claim indicates that social capital is actually present in these online communities, which provides a counterexample to Putnam's original claim that social capital is on the decline.

The advent of online social media networks created another way for people to connect with others over mutual interests. For example, Twitter provides an easy way for users to build their network by 'following' the people or topics that they are interested in. These online networks have evidence of embodying social capital. In "Tweeting Alone: An Analysis of Bridging and Bonding Social Capital in Online Networks" [13], the authors examine online Twitter connections formed across three major current events. They analyze the Twitter network structure with respect to the definitions of forms of social capital originally created by Putnam. The authors of "Tweeting Alone" found evidence for both bonding and bridging social capital in the online networks they studied. However, they found that social media groups were disproportionately successful at increasing *bonding* capital, or bringing together like-minded people to create small, dense groups. This finding is consistent with the changing community structure that Wellman had observed.

When people's social networks are centered more around interests than geographical location, the social capital in a person' network has increased value to others that share those interests. By sharing his LinkedIn article, the CTO of Hubspot was hoping to bring attention to his business strategies from other entrepreneurs and technology leaders. If the typical Hubspot employee had a social network consisting mostly of their local hometown community, sharing the LinkedIn article with such a network would be unlikely to help the CTO reach his goals. However, because Hubspot employees had connections to other professionals who shared their career interests, their networks held much more value.

Whether offline or online, people have specialized community ties to others and reach out to different people for different types of support. The CTO of Hubspot leveraged his corporate network to achieve a corporate goal, but he may have communicated with other people in his life for emotional support or advice on his hobbies. Wellman noted: "People do get all kinds of support from community members but they have to turn to different ones for different kinds of help. This means that people must maintain differentiated portfolios of ties to obtain a variety of resources. In market terms, they must shop at specialized boutiques instead of casually dropping in at a general store".[12] Since people have networks with connections to others that share their interests, leveraging a specific friend's network based on a particular topic or interest is highly valuable.

Yes, you may borrow my network

Given that the changing nature of communities and social networks provides increased value in accessing the networks of others, we wanted to create a tool that would help people access the networks of their friends. This tool would take advantage of two of the most powerful benefits provided by social media. The first is the ability to quickly share a message with your entire social network, simply by clicking a button. The second is the potential to reach an exponential number of people through getting your message re-posted by others (via shares on Facebook, and retweets on Twitter). We saw an opportunity to help users reach their friends' networks through retweets and shares, and built Shout!, a marketplace for friends to exchange retweets. Shout! would help users reach a wider audience and enjoy the benefits of a more expansive social network.

The current norms for retweet behavior do not inherently take advantage of the power of shares. The benefits of retweets are enjoyed disproportionately by celebrities and popular media sources (in 2010, 8 out of 10 of the most retweeted posts were from celebrities).[14] Retweets are also most commonly used to pass along news and links from popular media sites.[15] Additionally, while users might be happy to retweet content from their friends, whether they even see their friend's post on Twitter can be subject to factors like the time the tweet was posted and the the news feed algorithm. These factors can lead users to take more active role in soliciting retweets from their friends.

Currently, retweet requests and exchanges are coordinated through personal communication among friends, without a marketplace structure. This method has inherent limitations. For example, it is easy for friends to forget to retweet posts, and it can be difficult to coordinate posts being shared at a specific time (for a time-sensitive platform like Twitter, this is important). We developed Shout! with the goal of facilitating this access to a broader social network by making this process smoother. Shout! is an online marketplace for retweets, where friends can setup retweet trades with one another. Users have the flexibility to trade at different ratios (i.e. "I'll retweet 3 posts for you if you retweet 1 post for me"). Users can then directly trigger retweets through their friend's account when they so desire.

Shout! is designed for friends and the platform is built on the existing trust implicit in those friendships. In the design and implementation of Shout!, a large part of the research effort involved understanding how such a marketplace built on trust might function. Thus, the creation of the tool involved not only building the software, but also developing and honing the rules of the platform.

With Shout!, we wanted to provide value that could otherwise be difficult for users to achieve in an in-person conversation. Unlike both mainstream retweet exchange platforms and traditional in-person interactions, Shout! provides a way for users to bargain, and agree on, quantities of retweets to exchange. The differences in these quantities might differ based on the level of existing trust or friendship between users, the difference in popularity between users, or another factor. Thus, in developing Shout! we also have created an avenue to begin exploring research questions related to how people actually price, and value, popularity. While this topic was not the primary focus of this work, we hope to set up the framework for subsequent statistical analysis in this field.

Related Work

Shout! is a marketplace that helps friends work together to get their messages further. The basic mechanism through which this happens is the exchange of retweets. Exchanging retweets is not a new concept; many companies, organizations, individuals, and inevitably spambots, have previously been interested in leveraging the power of retweets to push their messages out to a wide audience. As a result, there are many retweet exchange sites in existence (as well as similar tools for different social media sites).

The vast majority of online retweeting services have the same underlying model. Users accumulate credit points by retweeting others' posts, and then offer this credit as a reward to other users who choose to retweet his or her own content. JustRetweet, a popular exchange site, describes how the site operates: "Earn credits by retweeting, liking, or give out +1s from retweets submitted by other members. When you have earned enough credits, then you can submit your own retweets into the system and have other members start promoting YOUR content."[16] Users can typically choose the number of credits to assign to a tweet themselves, so tweets that offer more credits will garner more retweets.

This retweet exchange model caters to use cases where users want to widely and indiscriminately send their message out. While this is powerful, this bears no resemblance to social networking in the real world, where people meet and connect with others through friends. These tools provide no mechanism to connect with people you know or to target any particular audience. As a secondary consequence, because this model is so indiscriminate, it naturally attracts a large number of spam bots, which degrades the user experience.

The instagram app Postforpost allows a way for users to engage in more personal sharing of content.[17] This app allows the user to choose another user they would like to trade a post with. Both users then choose a specific post they'd each like to share, and the receiving users approve it before sharing it to their own walls. In terms of targeting a specific audience through a retweet, Postforpost is a huge improvement over the standard retweet exchange platforms.

Shout!'s model provides a similar benefit to Postforpost of allowing users to select their target audience. However, Shout! builds on this by also leveraging the existing trust implicit in friendships. With Shout!, instead of setting up a one-time share, users actually set up longer-lasting trade relationships with their friends. Shout! also uses existing trust by providing an option for users to directly send retweets through their friends, streamlining the process further. And with Shout!, users can actively control both the audiences and the timing of the posts rather than simply passively receiving indiscriminate retweets.

Leveraging existing friendships is valuable, but there is also huge potential to build social reach by making new connections to people with similar goals or interests. Socedo is an application that finds users new social leads to help them expand their network.[18] The tool analyzes each user's Twitter activity and allows the user to specify the type of people they're looking to connect with (e.g. 'Software Engineers'), their goals ('Drive website visits'), and even the types of things their ideal connections tweet about (through keywords). Based on this information, Socedo generates suggestions for connections that users can approve, and Socedo provides options for users to automatically engage with their new connections (for example, there is an option to retweet a new connection's latest tweet immediately).

Software that intelligently suggests social leads, if combined with functionality to set up retweet exchange relationships, might provide the most explosive way for users to grow their social capital

and share their messages widely. In our first versions of Shout!, we focus mainly on the retweet exchange relationship aspect, though we acknowledge that a tool to suggest connections to users would be highly valuable. We dive more deeply into possible uses of this technology when we discuss future additions to Shout!.

Conceptual Design

We explain the motivation behind the design of Shout!. We then explain the major design decisions that went into the creation of Shout!, and finally present the user interface and the rules of the platform.

In creating Shout!, we saw an opportunity to lower the access barrier to social capital by allowing friends to borrow each other's social networks. Online social media networks were a perfect avenue to capitalize on, and thus we developed a retweet exchange marketplace with the following goals. First, we wanted to mimic a real-world interaction between friends as closely as possible. This would allow us to leverage the trust already existing in the friendship so users could send retweets to specific audiences, avoiding the spammy and indiscriminate nature of existing retweet exchange sites. Second, we wanted to provide added value by lowering the transaction cost to starting a retweet exchange with a friend. Third, we were interested in learning about the types of connections created over Shout!, and the level of trust implicit in these new relationships.

The major design decisions we made were choosing the social media platform to base Shout! off of, creating the structure of the trade interactions and marketplace, and developing the rules based on how much trust existed between users on Shout!. We describe these decisions in detail and relate each one back to Shout!'s goals.

Design Decisions

Why Twitter?

The idea of Shout! is relevant to any social media platform that uses sharing actions (retweets on Twitter, shares on Facebook, LinkedIn, or Instagram, etc). Having said that, the two most

well-established and widespread platforms are Twitter and Facebook. Of these two, Twitter was the favorable choice for three reasons. First, on Twitter, every user, regardless of their popularity, has the same type of account. On Facebook, celebrities often use fan pages, which creates more of a separation between these users, who enjoy the greatest popularity benefits, and everyone else. Second, Facebook has a more complex set of privacy settings, making it less feasible to share a non-friend's post; on Twitter, everything is public except if users choose to protect their tweets. It is a simpler model to work with. Third, average posting frequency is much higher on Twitter, and it follows that each post has less relative importance than a Facebook post. We see this by looking at the behavior of marketers, who have to be in-tune with the optimal frequency to post content on various social media sites.

Buffer, a social media sharing tool, has researched the rate of posting that appears to maintain strong connections with followers, without annoying them (the quantity of likes and shares on content are used to measure engagement). Based on this analysis, Buffer recommends that marketers post 14 times a day to Twitter, as opposed to twice a day to Facebook. Logically, the life cycle of a tweet is also much shorter than a Facebook post. After 18 minutes, a tweet has generally accumulated half of the retweets it will, as opposed to a Facebook post which takes 90 minutes.[19] Based on these differences in post importance on Twitter and Facebook, it is probable that the concept of reposting through another user's account may be more natural through Twitter, rather than Facebook.

What is a marketplace?

We prioritized creating a user experience that would mimic a real-world, in-person retweet exchange agreement as closely as possible. Thus, we created a model where a user first chooses one of their own recent tweets, and then chooses a particular friend to retweet that tweet. This model allows users the flexibility to control which audiences are targeted by each tweet, and is a direct real-world parallel to a conversation between someone and their friend asking them to retweet a tweet.

While we aimed to parallel the real-world experience of asking a friend for a retweet, we also wanted to provide value that is otherwise difficult to achieve in-person. As a result, Shout! not only allows users to select who to send their messages through, but it requires that users set up explicit trade relationships with other users in advance. Each trade relationship contains a contract containing the

number of retweets that will be exchanged, and this contract is not necessarily symmetrical. For example, one deal could be “I will retweet three tweets for you if you retweet five tweets for me”.

By explicitly creating trade relationships with friends, Shout! provides an avenue for users to layer a new type of network on top of their existing social media networks. This is the crucial add-value of Shout!- The value provided is the creation of a new social network, one that you can draw capital from on-demand, and one whose time and energy cost is too high to feasibly set up through offline interactions.

The network that Shout! provides is a social network where those who are connected share each other’s content. Certainly, this type of network existed informally before Shout!. Retweeting is a common activity, and often it is one that systematically occurs between friends. However, Shout! actually provides us an avenue to quantify this retweet exchange behavior. Through Shout!, we can examine the links created between people and start to understand why people trade at certain ratios. For example, the graduate students in the MIT Macro Connections Group have a fraction of the Twitter followers that their professor, Cesar Hidalgo has. While this popularity imbalance might make it seem unappealing to Cesar to trade at a 1:1 ratio with his students, two students thought a trade ratio of 5:1 (where the student retweets 5 tweets for Cesar) seemed more fair, and both parties were satisfied and benefited from the trade. Did they choose to trade at 5:1 because Cesar had five times as many followers as his students? Do the rates at which people retweet one another correlate so smoothly with the difference in their social reach, or are there other factors at play? If trade ratios are correlated with popularity imbalances, it is logical that celebrities essentially trade with average people at a ratio approaching 0:infinity.

We hoped that Shout! would provide a way to start exploring how people perceive the relative price of popularity, based on the trade relationships they set up.

How much trust is too much trust?

Shout! assumes a high level of existing trust between users. This assumption affords users the flexibility to trigger messages through friends’ accounts. During our initial design phase and while demonstrating alpha versions of the site, we noticed that the main hesitation from users stemmed from a feeling of a loss of ownership over their own social media account. Users feared being

impersonated or misrepresented on their own Twitter timeline. Twitter's prevalent bot population (The company disclosed that in 2014 approximately 23 million, or 8.5%, of its active users were automated[20]) also meant that engaging in a trade with a bot would increase the risk of spam retweets. Many users felt more comfortable with the idea of being able to approve retweets posted on their wall. Yet, others enjoyed the instantaneous nature of directly triggering a retweet to a friend's wall.

We understand the importance of not violating a key sense of ownership users want over their online profiles. Taking into account this importance, the variation users have in how much access they would be comfortable giving a friend to their Twitter account, and the need to protect users from spammers or bots, we made a few key design choices when creating Shout!.

First, we made the choice to have Shout! be a retweet exchange site instead of a tweet exchange site. In theory, a marketplace for tweets would have even more powerfully leveraged friends' networks, as a user's message would appear to come firsthand from each friend. However, the risk of impersonation would be too high for users to engage with such a platform. Thus, we opted for retweet exchange instead.

We made a major choice in designing the actual retweet process. We originally designed the platform to support only directly triggered retweets. That is, if a user Alice has a trade relationship with Bob, she can choose a tweet and retweet it through Bob's account, without him needing to approve the tweet or take any other action. This direct approach has two main advantages. First, it capitalizes on the time-sensitive nature of Twitter and gives the user more control of when his or her tweets are retweeted. Additionally, it saves the user who is retweeting each tweet the trouble of having to log in and approve every single request. This model is appropriate for users who have a high enough trust level with the person they are trading with. For users who are want to exchange retweets but do not have the level of trust necessary with their trade partner to allow directly triggered retweets, Shout! also offers the user an option to approve retweets posted through his or her account. This setting is specific to each trade the user is in, so the user can set up different types of trades depending on their relationship with each person.

Shout! is designed for friends to set up trades with one another, so the tool is best used to increase efficiency of retweet exchanges that could otherwise have been set up in person. Thus, we broadcasted a rule of thumb to our users, which is that users should not set up trade relationships

with people who they would not normally feel comfortable asking in person. Having said that, we acknowledge that mistakes can happen and a robust tool must have support for handling possible misuse. While the initial release of Shout! does not provide support for handling malicious users or bots, we will later discuss possible extensions and design changes that would need to happen in future iterations to protect users adequately.

User Interface

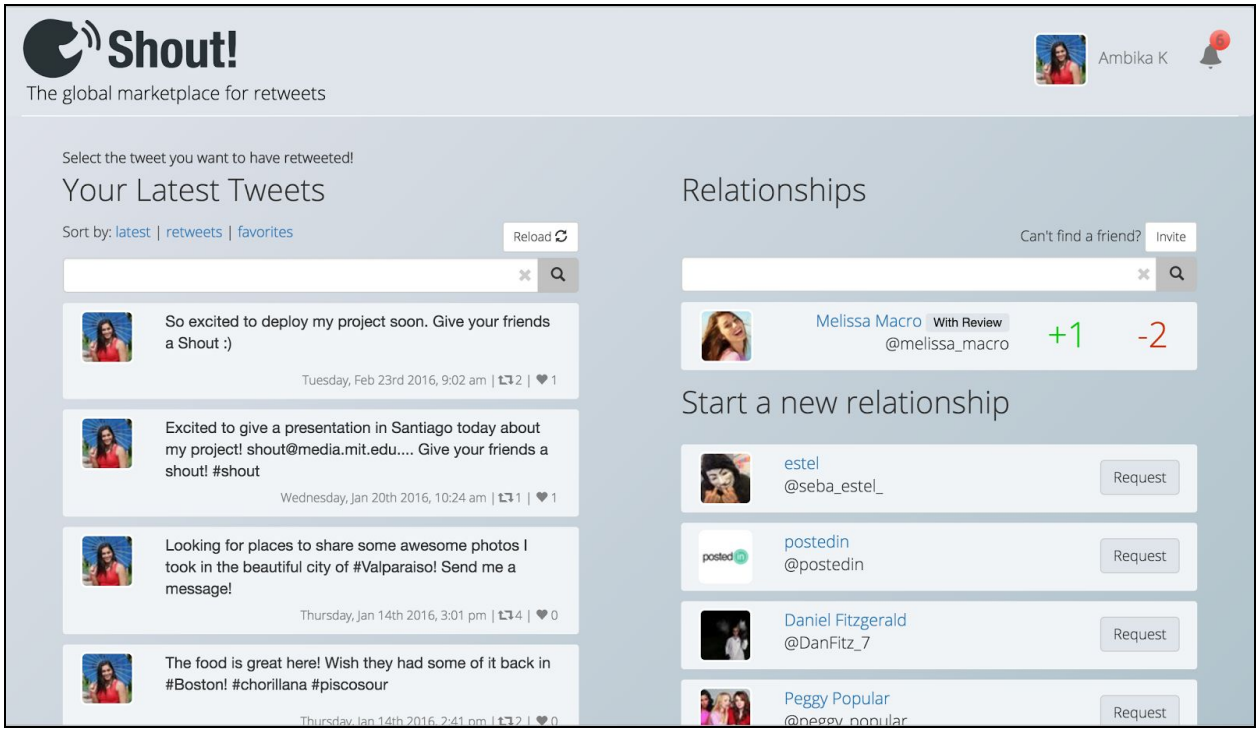

Figure 1: The homepage of Shout!.

We illustrate the design decisions we made by showing the user interface of Shout!. Figure 1 shows the main homepage of Shout!. The left column displays a searchable list of a user's recent tweets. The right column displays users. At the top appear the users which the logged-in user has a current trade relationship with. Underneath this set is a list of other users the logged-in user can request. To set up a new trade relationship, the user simply clicks on another user's name to propose a new trade with

them. The trade request includes the proposed trade quantities. The receiver of the request can choose to accept the trade, which creates a new balance of retweet credits between the two users. The receiver can also reject or modify the trade by adjusting the trade quantities. Like any traditional bargaining process, modifications can be passed back and forth until they reach a final accepted (or rejected) state. In the event that a user wants to set up a trade relationships with someone not yet on the Shout! platform, the user can send a Twitter direct message to invite them to the platform. Finally, to actually trigger a retweet through a friend, a user simply has to select a tweet, select the friend, and then confirm the action to send the retweet.

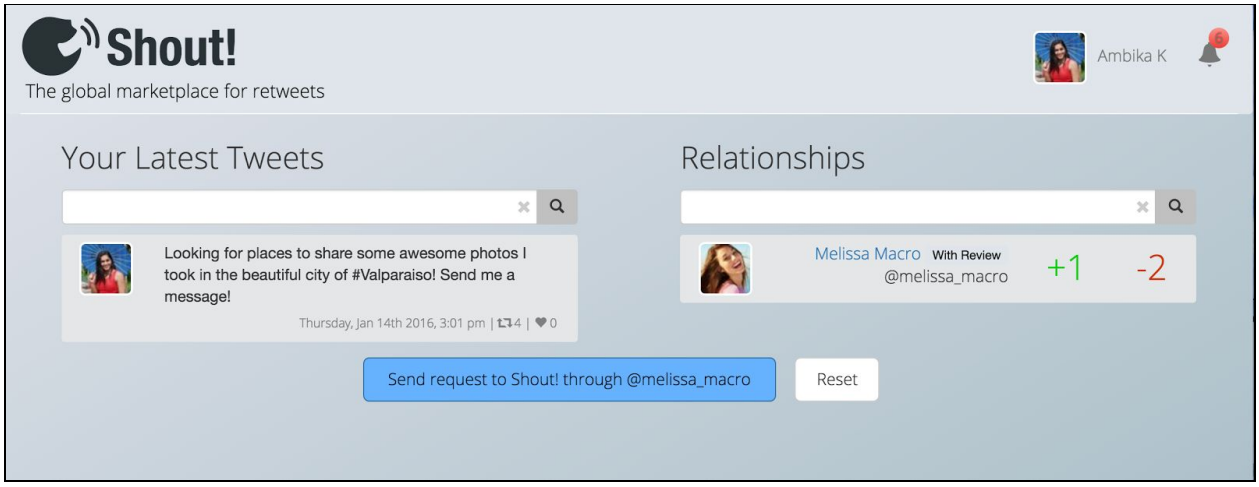

Figure 2: Retweet trigger screen.

Technical Design

We developed Shout! as a web application in 2015-2016. Here, we explain the main technologies used, the data model, and the noteworthy engineering decisions that went into the creation of the site. This section's intended audience is software developers.

Technologies Used

Shout! is a web application built using Meteor, a javascript web framework. We chose Meteor for its simplicity of use, excellent built-in security support, and seamless integration with Twitter; we connected Shout! to the Twitter API for user authentication, tweet retrieval, and posting retweets. For our production setup, we deployed Shout! to a virtual machine hosted by the MIT Media Lab. We

used the mLab database-as-a-service to host our MongoDB database. The mLab service provides first class monitoring, backup and failure recovery, and provisioning. We may switch Shout!'s production setup to Heroku, a deployment and application management service. For the early release we provisioned our own virtual machine instead, but we have a backup Heroku setup to move Shout! to in case we require better load-balancing support or have other related issues. For monitoring and analytics purposes, we use Google Analytics and Kadira, a Meteor-specific application monitoring solution.

Data Model

In Shout!, the main information we must manage are user information, the status of every current trade, and the trade requests that are passed back and forth. We present a simplified form of the data model that we use to handle this information. Meteor applications are backed by MongoDB, a NoSQL database made up of collections (like tables in a relational database) of documents (entries).

The main collections we have are *users*, *tweets*, *trades*, *trade requests*, and *recent activity*. Each of these collection's purposes and usage is summarized in Figure 2.

Collection Name	**Key Fields in Document**	**Purpose**
USERS	❖ Basic information (name, screen name, profile picture, email) ❖ Twitter access credentials ❖ Twitter friends/ followers	Store user information and information about their connections, for the user-list sorting algorithm.
TWEETS	❖ Similar information as a normal twitter tweet (text, date, user, etc)	Store our own tweets (so we can use our own search and minimize calls to the Twitter API)

TRADES	❖ Users involved in trade ❖ Trade quantities on each side ❖ Direct or indirect trade	Track the status of each current trade.
TRADE REQUESTS	❖ Sender of trade request, recipient of trade request ❖ Proposed trade quantities ❖ Whether this request is a counteroffer to another request	Track the status of each current trade request.
RECENT ACTIVITY	❖ User(s) involved in notification ❖ Notification content (e.g. accepting a trade request)	Used to create notifications and populate Shout!'s Recent Activity feature.

Figure 2. Description of the main Mongo collections that back the Shout! application.

We also have helper collections that store information like historic trades, historic trade requests, and already-retweeted tweets. For the most part, these collections are not used directly by the Shout! interface, but they are useful for gathering historical data about behavior on Shout!.

Technical Design Decisions

Many of our technical design decisions were motivated by the Twitter API developer rules. Shout! Is a third-party Twitter application (as defined by the official Twitter site, a third-party application is an "a product developed apart from Twitter.com or Twitter's official mobile apps, and that is used to access Tweets and other Twitter data"[21]). Twitter outlines certain rules that such third-party applications must abide by. These rules include limits on the number of times your application can query the Twitter API, and display requirements for tweets and other Twitter-related content. Twitter enforced stricter developer rules after June 2012, when they reduced the number of Twitter API calls per hour each endpoint could make from 360 to 60, required every third-party application to authenticate with Twitter, and began to revoke access to third-party applications that violated the Developer Rules.

These rules represent an effort on Twitter's part to bring the tweeting experience back to Twitter itself, rather than from outside third-party apps. Ryan Sarver, Twitter's Head of Platform and API, wrote in March 2011 that: "Twitter will provide the primary mainstream consumer client experience on phones, computers and other devices by which people access Twitter content (tweets, trends, profiles, etc). If there are too many ways to use Twitter that are inconsistent with one another, we risk diffusing the user experience."[22] Twitter finds greatest value in applications that provide business or analytics value, and lose value from traditional Twitter clients.

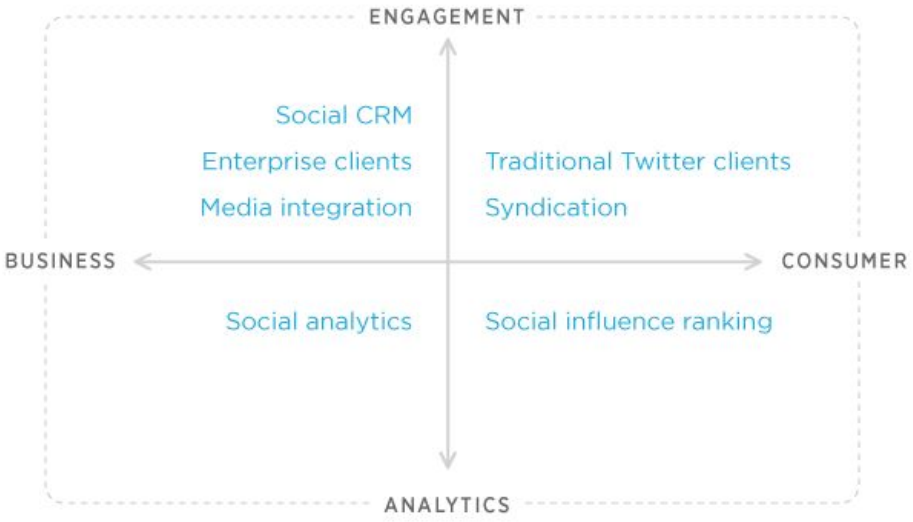

Figure 3. Twitter's application ecosystem. Twitter's more stringent developer rules restrict applications that fall in the upper right quadrant.

Twitter's developer rules represent a long-term company preference to have traditional tweeting and retweeting behavior occur through Twitter itself. If the developer rules pose problems for Shout! down the road, we will explore options such as becoming an official application under the Twitter umbrella, further optimizing the application to use fewer Twitter API calls, or switching the social media platform to one that provides greater flexibility. However, for the first release of Shout!, we worked on developing software that could provide a good user experience while minimizing the volume of requests to the Twitter API. This goal motivated many of our technical design decisions.

One of the decisions motivated by this goal was the way we load, store, and search tweets in Shout!. Shout! displays a list of the logged-in user's latest tweets on the left side of the interface. Because of the nature of timeline retrieval through the Twitter API, pulling the logged-in user's tweets on every

homepage visit would require multiple API calls to fetch the entire timeline, increasing loading latency and risk of running against the Twitter API rate limits. Additionally, we wanted to provide users the ability to search through their recent tweets to select one, and Twitter's Search API does not allow the developer to search content that is over one week old. These restrictions were suboptimal, so we chose to instead store tweets ourselves and provide our own search feature. Every time a user logs into Shout!, we query Twitter for any tweets they have posted since their last visit to Shout!. We then store each of these tweets in our own database, which we index to provide fast text-filter search performance. While the search feature at the moment is basic, we are looking to provide intelligent search and autocomplete in future iterations of the tool.

Another design decision we made related to displaying other users on the homepage of Shout!. Shout! shows a searchable list of users, so that the logged-in user can see who else is using the platform, and request other users to start trade relationships with. An unordered list of users would be unhelpful, especially as the number of users on the platform grows. Our early-stage ordering algorithm first displays users that the logged-in user follows, then users that *follow* the logged-in user, then other users. In order to store the information necessary to calculate this ordering on page load, we query the Twitter API for a user's follower list and friend list on login, and store this information. In future versions, we will have to investigate more complex user ordering algorithms that will increase users' rate of requesting others, while maintaining minimal traffic to the Twitter API.

In many of our design decisions, our choice to minimize calls to the Twitter API comes at the expense of fresh data. For example, once tweets are loaded in our database we do not revisit them, even if the tweets may have garnered additional favorites and retweets. We also do not do a live update on the user timeline if a user posts a tweet through Twitter during their Shout! session. However, we implemented a 'Reload' button so users could manually refresh their Shout! timeline to retrieve a new tweet. Lastly, we do attempt to keep the list of a user's friends and followers up to date, so this is refreshed on every user login.

Validation

In releasing Shout! and observing user activity on the platform, we had three main validation goals. First, we wanted to understand how user activity and retention rates on Shout! depend on its

features. For this, we did several releases with incremental feature additions between each one, to assess their impact. Second, we were interested in exploring the exchange rate between users to understand how users perceive the relative price of popularity. Finally, we wanted to understand the role that pre-existing online relationships between users plays on Shout!. Here, we discuss both our quantitative and qualitative findings during Shout!'s release journey with respect to each of our validation goals.

User Activity and Adoption

We first released Shout! for alpha testing in Valparaiso, Chile in January 2016. We demonstrated the platform to users in the B3 research camp, a four-week collaborative research effort with students from the MIT Macro Connections Group, the Center for Research in Social Complexity at Universidad del Desarrollo (UDD), and selected students from other graduate programs in Chile. We also put together a promotional video to inform people about Shout!, and at the end of the research camp, we presented both this video and a live preview of Shout! to three hundred students, professors, venture capitalists, and others interested in the technology developed at B3.

From the first release, we gathered useful feedback about people's comfort level with trust. While people were interested in the product, throughout the month many users were hesitant to sign up because they feared being impersonated or misrepresented on the platform. With these concerns in mind, we implemented the option to choose between indirect and direct retweets. We then officially deployed a beta version of Shout! in April 2016 to the MIT Media Lab during the Members Event, a two-day demonstration of projects in the lab. This release garnered about a hundred new first-time users as well as plenty of qualitative feedback through conversations with company representatives and investors. We noticed that users appreciated the option to choose between direct and indirect retweets, and that implementing this option helped alleviate concerns about impersonation on the platform.

	After the Media Lab release	After the MIT / Social Media Release	Total

Number of unique users	103	35	138
Total number of trade requests	81	58	139
Number of unique users who sent or received at least one trade request	53	23	76
Number of trades	35	29	64
Number of retweets	30	20	50
Number of unique users who triggered at least one retweet	18	13	31

Figure 4. User adoption and retention after each release phase.

We observed the following user activity after the Media Lab release:

103 users logged into Shout!

53 users sent or received at least one trade request

18 users triggered at least one retweet.

We noticed a 50% drop off from the number of users that initially logged in to the number of users that sent a trade request, and then a further 66% drop off to the number of users that actually triggered retweets. Both drop offs can be partly attributed to the nature of the Media Lab release. Many Members Event attendees were interested in learning about the application and would log in and look at the site interface, but they were not necessarily interested in using it immediately.

Nevertheless, we implemented three more features with the goal of boosting overall user retention rate. The first feature was an invite button, so users could more easily get their friends to join Shout! and begin trading. The second was a quick-tweet feature, so that users would not have to leave the Shout! platform to send a new tweet, but could instead send the tweet and then trigger retweets in one continuous process. The third was improved notification emails as well as alerts on Shout!, so users could see the details of a trade proposal and easily log in to approve it.

In May 2016 we did a third release of this new version, leveraging our own social media networks to generate user activity and excitement about Shout!. Through this release we picked up 35 new users, the majority of which were MIT students. These students were encouraged to try the tool on their own time (and were not exposed to it in the midst of many other demos and presentations, which

was the situation during the Media Lab release). Among the MIT community, Twitter is used less often, and many MIT students had not used Twitter consistently for several years. Thus, these users logged into Shout! to try it out but were less likely to stay engaged long-term (however, they may have been more likely to engage if had we built an equivalent tool for Facebook).

After this release, Shout! received the following additional user activity:
35 users logged into Shout!
23 users sent or received at least one trade request
13 users triggered at least one retweet.

During the third release, five users used the quick tweet feature and three users invited friends to join Shout! through the invite feature; these activities made a small positive impact on retention. The increased retention rate was also likely impacted by the difference in types of audience from the second to the third release. While rate of retention was improved after the third release, we still noticed significant drop-off both between the number of users that logged in to the number who sent a trade request to the number that triggered a retweet.

During the third release, we were also able to observe six users in person trying Shout! for the first time, in order to gain insight into why retention rates were not higher. We noticed that often users were confused about how Shout! worked, beyond the fact that it helped users get more retweets. Additionally, if the user did not have a clear idea of what Shout! was before logging in, he or she would often have difficulty navigating the site. Several users thought that once they entered a trade relationship, they were supposed to choose which of their friend's tweets to retweet, and they spent time looking for a way to select their friend's tweet instead of their own. Once these users realized that Shout's model is to directly trigger retweets through friends' accounts, this realization was shocking. Shout's model is powerful, but it is still extremely counter-intuitive, and its difficulty to understand can drive users away.

We attribute the second drop-off (between the number of users who sent trade requests and the number of users who triggered retweets) to two possible reasons. First, many trade requests are left pending on Shout!. Out of 139 total trade requests sent during the releases, 56, or a full 40%, of them were still pending over a day after the request had been sent. Some of the users we spoke with

explained that if they were not enthusiastic about the person sending a trade request (or did not know them), it felt easier and less confrontational to let the request hang, rather than to explicitly reject the request. Other users found that they were likely to forget about the trade request after even seeing the notification email, especially if they did not have immediate use for the retweets that would be provided in the new trade relationship. To boost the response rate on trade requests, in future releases we will follow up with users over email and possibly Twitter direct messages, reminding them of the trade request and encouraging them to accept the request quickly.

However, a significant dropoff in the number of triggered retweets is actually in line with our expectations for user behavior on Shout!. Shout! allows users to build up retweet credit with friends for later use. Then, once the user has a tweet they would like to reach a wide audience, they can take advantage of their trade relationships. Thus, it is likely that many users who tried Shout! did not immediately have a tweet they wanted to use their relationships to boost. We would have to do a longer-term analysis of retweeting behavior on Shout! to confirm this, which we leave to future investigation.

Overall, the rate of early adoption was slow. Part of the reason for this was release timing; the beta release had glitches in the website and the homepage of Shout! was susceptible to slow loading times. This led to subpar user experiences as we sorted out problems with the site, and we lost the opportunity to engage many users during the beta release. Additionally, Twitter is a platform that is more popular in certain communities than others.

Another way we will boost adoption is by improving the design of Shout!. A well-designed site must communicate both these aspects of the Shout! model (direct retweet triggers and the concept of saving future retweets) to the user without the user needing to watch a promotional video to understand them. We will have to incorporate this user testing feedback into a future iteration and improve the visual design of Shout! until it is more learnable and intuitive.

Lastly, we have received significant interest in the idea of Shout! for a corporate use case. Similar to the Hubspot case, company representatives we spoke to saw the inefficiencies in their current marketing strategies that Shout! could solve. They were also interested in ways that different companies could cooperate with one another using Shout!. Depending on the model of trust that

exists between companies, Shout! could serve as a solution for corporations as well as for individuals in the future.

The Relative Price of Popularity

One of our driving motivations in creating Shout! was to begin to understand how people value access to friends' social networks. By leveraging a friend's social network you can significantly expand your social reach; but what price can you put on that? More precisely, if two users enter a trade relationship with each other, they assign a relative 'price' to this relationship in the form of number of retweets on each side. How are these retweet numbers agreed on? If User A and User B start a trade relationship on Shout!, but User A has three times the number of followers as User B, does that mean an appropriate trade ratio would be some scaled version of 3:1? After we released Shout!, early adopters had an opportunity to begin to define these trading norms.

We make the assumption that the number of Twitter followers a user has is a good proxy for the strength of their social capital. This is not always the case, as all followers are not equal in the strength they contribute to the network; for example, a twitter bot counts the same as a celebrity. However, the audience for Shout!'s beta release consisted mostly of people whose Twitter networks reflected their real-world social networks. The presence of bots (or celebrities!) in users' follower networks was minimal.

Based on the trade relationships established over Shout!'s beta release, it appears that users may not have had a strong intuitive sense of what a 'fair' trade price is on Shout!. We imagined that in the process of establishing a relationship, trade requests would be passed back and forth several times through counteroffers until the quantities reached some equilibrium that both users were happy with. This behavior ended up being uncommon. By early June, there were 83 trade requests that had been acted on (and an additional 56 still pending). Of the 83 resolved requests, 64, or 77%, of them were approved immediately. 14 of the requests were modified, and five were rejected.

Based on our expectations for the bargaining process and the actual behavior we observed, we thought the proportion of users that accepted trade requests immediately (without modification) was disproportionately high. We suggest four possible reasons for this behavior. First, users may not have had an intuitive sense of what a fair trade price is, and thus were more likely to passively accept offers

sent their way; this would allow the sender of the request to dictate the final ratio. Second, users who sent trade requests tended to send reasonable requests. Because users traded with their friends, there was less incentive to 'rip someone off' than there might have been in transactions with strangers. Third, users may have been hesitant to reject or modify requests for fear of appearing rude or hostile. Social norms on Shout! were largely undefined and users prioritize maintaining their friendships. Finally, perhaps users simply did not care so much about the actual trade ratio (unless it was egregious). For example, Professor Cesar Hidalgo accepted every offer sent to him, regardless of the ratio. For him, the ability to trigger retweets through established relationships with friends was valuable, and it was not worth it to haggle back and forth about the trade quantities.

For a wider-spread release of Shout!, there should be a more uniform notion of what appropriate trade prices are, as well as what the social implications around modifying or rejecting requests are. These norms may evolve over time as more users engage with the platform. We could also accelerate the formation of these norms by providing users suggestions of quantities to trade at when they start a trade request. These suggestions could make users more comfortable with Shout! and might increase the number of trade requests we see. However, this feature would have biased our analysis into how users price popularity, so we leave it as a potential future addition.

Even though trading norms were not fully established during Shout!'s beta release, we examined the trades set up on Shout! to see if there was evidence of any relationship between users' popularity and the ratios they traded at.

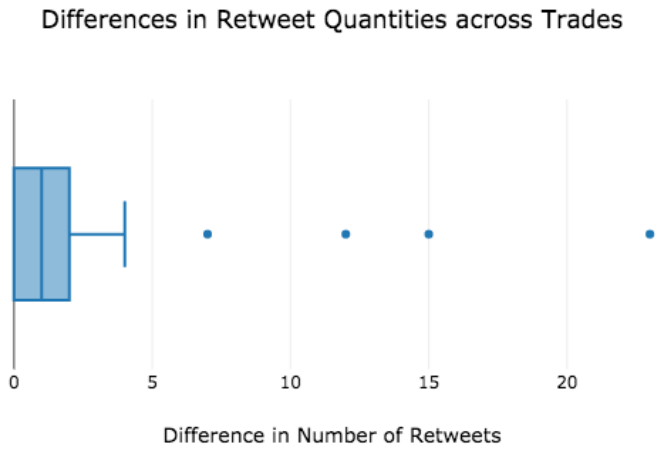

Figure 5. Distribution of the difference in quantity of retweets across all trades. The difference is calculated as No. Retweets User 1 receives - No. Retweets User 2 receives, where User 1 receives more retweets than User 2.

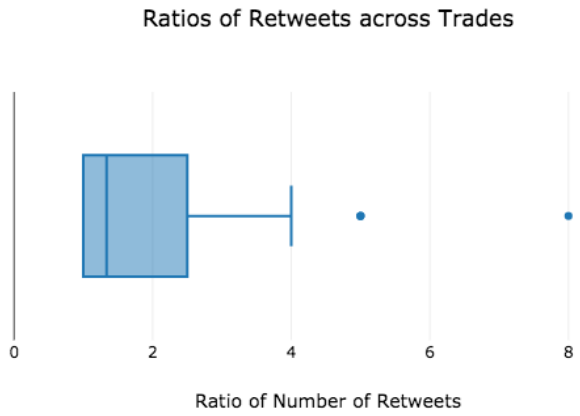

Figure 6. Distribution of the ratio between retweet quantities across all trades. Ratio is calculated as (No. Retweets User 1 receives) / (No. Retweets User 2 receives), where User 1 receives more retweets than User 2.

First we look at the distribution of quantities that users traded at on Shout!. We can measure trade quantities in two ways, as a difference or a ratio. If Alice and Bob have a trade relationship where Alice can send 8 tweets through Bob's account, and Bob can send 2 tweets through Alice's account, then the trade difference would be 6, and the trade ratio would be 4. Most trades happen at at a close to even ratio, suggesting that friends are unlikely to set up trades with massive differences in the quantity of retweets exchanged.

With this background on trades and their general retweet quantities, we now pose a basic question: In a given trade relationship, is the user with the more expansive network more likely to have been awarded more retweets in the contract? Of the 64 trade relationships started on Shout!, this was indeed the case for 27 of them, or 42%. However, we noticed a common trend of users sending trade requests with an equal number of retweets on both sides. 28 of the 64 relationships actually had this pattern. When users are not sure what trade ratio is appropriate, then assigning equal numbers on both sides is a safe default. Additionally, users might feel inclined to trade equally with people they consider their close friends. If we only look at the trades with uneven trade ratios, there were only eight cases where a user had more followers but received fewer retweets in the relationship. In these cases, perhaps the number of followers was an inappropriate proxy for popularity, or it might be a result of the fact that people tended to accept any offer sent to them. Nevertheless, of the trade

relationships with unequal ratios, 27 out of 36 or 75% of these relationships awarded the user with the larger network more retweets.

From these observations, it appears that the popularity difference between users does have an impact on the retweet quantities assigned in Shout!. What exactly is this impact? To explore this question, we consider the different ways users can compare the size of their network to that of their friend. One common comparison strategy is look at the difference in absolute sizes of the network (i.e. “My friend has 80 more followers than me”). Another strategy would be to consider the multiplicative difference between networks (i.e. “My friend has twice as many followers as me”). Similarly, when considering how many retweets to exchange in a trade proposal, users can think in terms of absolute differences (“My friend will be able to trigger five more tweets than me”) or in terms of proportions (“For every retweet I send through my friend, he will send two through me”). To assessed whether popularity differences were correlated with differences in retweet quantities, for each trade, we computed the popularity differences and trade ratio differences based on each of the comparison strategies we outlined.

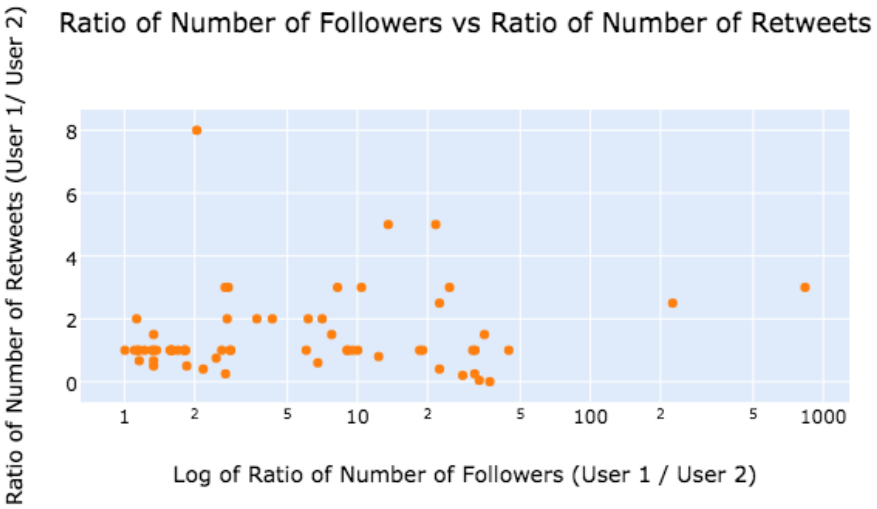

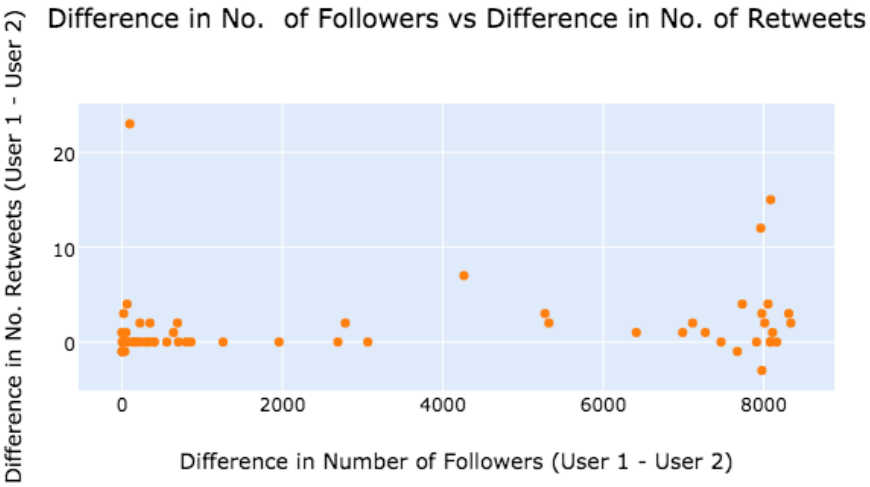
Difference in No. of Followers vs Difference in No. of Retweets
Difference in No. Retweets (User 1 - User 2)
20
10
0
0
2000
4000
6000
8000
Difference in Number of Followers (User 1 - User 2)

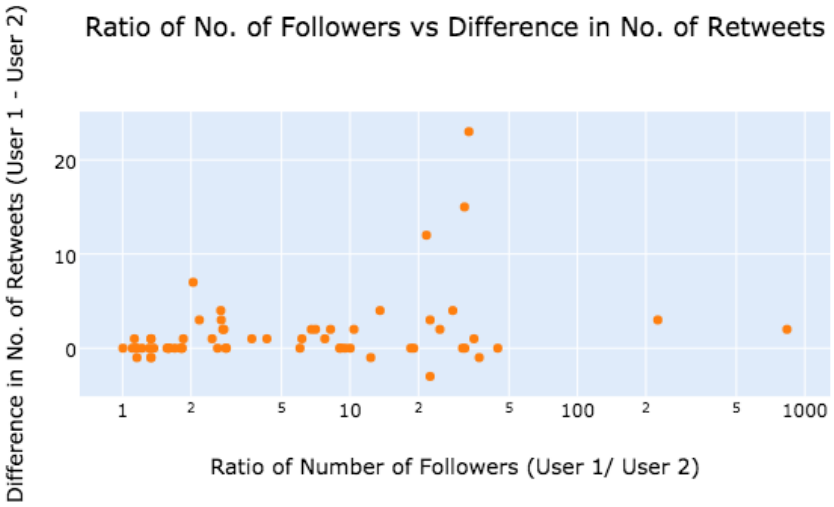
Ratio of No. of Followers vs Difference in No. of Retweets
Difference in No. of Retweets (User 1 - User 2)
20
10
0
1
2
5
10
2
5
100
2
5
1000
Ratio of Number of Followers (User 1/ User 2)

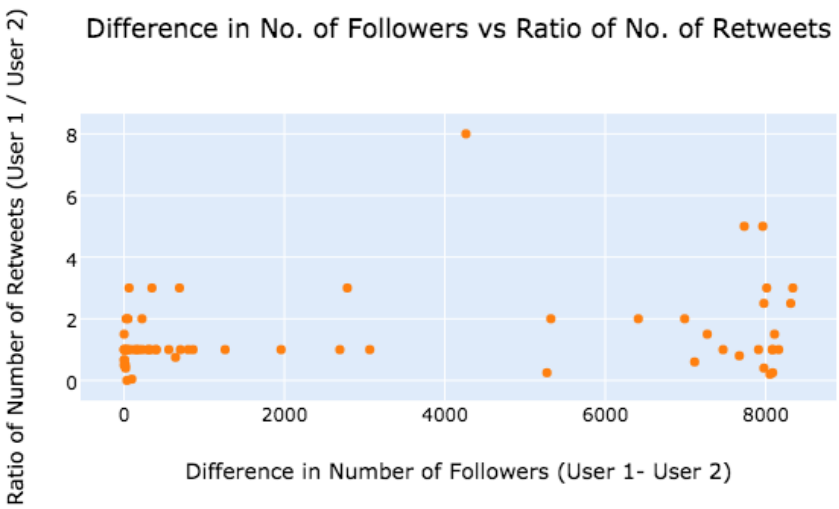

Figure 7. Four different angles to assess the correlation between the popularity differences of users in a trade request, and their trade quantities. In the graphs, User 1 has more followers than User 2.

We observe no significant correlation between popularity imbalances and the trade ratios they influence. However, we we do see that both the difference between the trade quantities and the ratio between them tends to stay low, even as the size of networks between users has a much greater difference. Out of 64 total trades, only four trades had a difference of more than five retweets per person, and only three trades had a ratio exceeding 5:1. This observation has positive implications for users who are looking to leverage the social capital of those who are more influential, in that they may be able to take advantage of a much larger network at a low cost.

We consider the lack of correlation between the two factors measured. It may be that there are other factors that more heavily inspire the trade quantities that users set in Shout!. Perhaps the differences in trade quantities are more strongly linked to a third factor, like the strength of the friendship between the two users in the trade. We do not have an adequate proxy for strength of friendship to analyze this, but we do notice that simply the presence of *a* friendship between users does affect trade quantities. Here, we take a deeper look at how existing online friendship affects behavior on Shout!.

The Impact of Trust and Friendship

Our final goal in validating Shout! was to understand the value of trust and friendship on the site. Shout! was built as a platform for use between friends. We assert that two users mutually following each other on Twitter is a reasonable proxy for friendship. To see how selective users were in choosing who to trade with, we look at the trade requests on the platform, and the proportion of trade request recipients who were friends with the sender.

	ACCEPTED Requests	REJECTED & PENDING Requests	ACCEPTED Requests with 1:1 trade ratios
Mutually follow	37 (58% of all accepted)	10 (16% of all rejected/pending)	18 (64% of 1:1 trades)
Sender follows receiver only	9 (14% of all accepted)	13 (21% of all rejected/pending)	2 (7% of 1:1 trades)
Receiver follows sender only	4 (6% of all accepted)	4 (7% of all rejected/pending)	3(11% of 1:1 trades)
Neither follows	14 (22% of all accepted)	34 (56% of all rejected/pending)	5(18% of 1:1 trades)

Figure 8. The number of accepted and rejected/pending requests, broken down by the type of Twitter relationship between the sender of the request and the receiver of the request.

We find strong evidence that users felt that a trade relationship on Shout! was appropriate mainly between friends. Of the accepted trade requests sent on Shout!, 58% of them were between users who mutually followed each other. More notably, of the requests that were rejected or left pending, only 16% of them were between users that mutually followed each other.

As we mentioned before, there were a high number of rejected and pending requests on Shout!. Our evidence suggests that users agree that trust is extremely important to engage in a trade relationship, but perhaps the high level of trust necessary is impeding new growth on Shout!, especially for use cases where users would like to form relationships with people they do not know as well. It is possible that in future iterations where we provide a more robust model to protect users from misuse, the trust threshold necessary to form a relationship will be lower.

Friendship seems to play a role not only in the decision to accept or reject a request, but it also impacts the retweet quantities in the trades. When we looked at the differences in popularity between users and how this impacted trade ratios, we noticed that there were a large quantity of 1:1 trades made on Shout!. We imagined that sometimes friends would choose to trade at an even ratio, regardless of a popularity difference they might have. The trades we see on Shout! support this theory. Of all of the 1:1 trades we observed, 64% of them were between users who mutually followed each other.

Once a user has decided to enter a trade relationship, Shout! offers a choice of whether or not to trigger retweets directly through his or her account, or to allow him or her to review it first. We found that out of all trades, 39% opted for the 'with-review' option. There was no notable difference in this proportion with regard to the type of the relationship between two users in the trade. This implies that the decision of whether to review retweets or allow them to be triggered directly can be a personal choice, determined by other factors such as how much the user values controlling their entire online presence. Our conclusion was that providing this feature was a good step, in that it expanded the audience of people who were comfortable trying Shout!.

From releasing and validating Shout!, we gained important insights about what factors are effective in retaining users, as well as the use cases that users would like Shout! to eventually cater to. We found only a weak connection between popularity differences and trade ratios, but we found evidence of both trades and trade ratios being strongly influenced by existing online friendship. Our findings confirm that friendship and trust are central components to the Shout! model.

Future Work

There are many ways that we are looking to improve Shout! and learn more about the value that users place on popularity and extending their social network. Our ideas for improvement fall in three directions: Conceptual changes, technical improvements, and analysis ideas.

On the conceptual side, we feel that Shout!'s conceptual model does not lend itself naturally to some of the use cases that people would benefit from. For example, users expressed interest in a corporate version of Shout!, where companies could leverage the professional connections of other companies

in their space. However, they were dissatisfied that Shout! had no model to handle malicious users or misuse (such as deleting retweets immediately once posted), and were wary about the impact another company posting directly through their media account could have on their social media image. The concern that someone will trigger a retweet on your wall that you do not like is a recurring concern, even between friends. This does not necessarily happen with abusive or malicious tweets; it could happen on a milder level (a friend triggered a retweet through your account on a topic you're not that interested in, and you like your twitter wall to have a certain theme). Based on this feedback, we think it might be valuable to provide the option to align trade relationships around certain topics. This option would provide a dual benefit - First, users could both more easily find people in their space of interest and could potentially expand their network that way, and second, users would feel greater control of the content being posted on their wall.

On the technical side, there are obvious improvements that will need to be made to the backend infrastructure to better support Shout! as it scales. We are also looking to improve both the load speed and the relevancy of the users Shout! displays. We are working on a recommendation algorithm to determine an ordering of users on the site that will allow users to quickly find their friends and discover new people to trade with. Building this recommendation feature is a significant project, and we will also consider hooking in an external service which has expertise in this particular area. Lastly, based on the current social media trends and the fact that Shout! as a concept is platform-independent, we may build and deploy Shout! for Facebook to reach more users.

On the Analysis side, understanding how people value popularity is an important social question. Shout! is one of the first tools that allows users to view popularity and social capital as tradeable; perhaps Shout! could pave the way for popularity to start being viewed as more of a currency than it is today. As Shout! expands potentially beyond the current main use case (which is between close friends), we would like to investigate user activity over a longer time period. This may lend more insights into the relative price of popularity. We would also like to assess Shout!'s longer-term effects on the user and their social capital. Over a long timespan, are users more successful in reaching wider audiences because of Shout!?

Finally, Shout! has the potential to redefine the conceptual meaning of a retweet, or more generally a social media share. In creating Shout!, we challenged the norm of 'how' to retweet, and provided a tool that takes advantage of the network-sharing that is implicit in the retweet action and allows

users to leverage retweets in an unorthodox way. We would be interested in exploring Shout!'s social impact on traditional social media sites, to see how it affects the concept of social media shares.

Conclusion

In this thesis I presented Shout!, a marketplace for friends to exchange retweets with one another. Shout! allows friends to set up micro-contracts with each other to exchange a fixed number of retweets, and provides users the ability to leverage their friends' Twitter accounts to have their messages reach a wider audience. Shout! essentially allows users to borrow their friends' networks, and redefines popularity as an asset that can be borrowed and traded for mutual benefit.

We designed the rules of this online marketplace and implemented the software, releasing Shout! to the public over three stages in the Spring of 2016. From our initial analysis, we found a strong connection between existing friendship and trade relationships. As we expected, friendship and trust appear to be strong requirements for engaging in a trade relationship on the platform. In our initial analysis, friendship had evidence of impacting not only the existence of a trade relationship, but the retweet quantities as well. On Shout!, friends were very likely to trade at even or close to even trade ratios, irrespective of their popularity differences. This observation has positive implications for users who are looking to leverage the social capital of their friends who are more influential, in that they may be able to take advantage of a much larger network at a low relative cost.

Shout! challenges the traditional notion of popularity and social capital as non-tradeable, private assets. It also challenges the norms of social media sharing. Over multiple iterations of user testing and data collection, we learned how such a platform could operate, and we have a better understanding of how to improve Shout! for the future. This conceptual design and first implementation of Shout! provide the groundwork for future tools that help friends leverage their networks, borrow each other's social capital, and get their messages further.

Acknowledgements

A big thank-you to Cesar Hidalgo, who was the original ideator of Shout! as well as the Principal Investigator, who provided immeasurable guidance and support throughout the year. Additional thanks to Cristian Jara Figueroa for his writing help, Joe Goldbeck and Anubhav Jain for the technical guidance they provided, Luisina Ardizzi for filming the promotional video, Nikita Pestrov for acting in the promotional video and providing background research for Shout!, and the entire Macro Connections and B3 crew for their promotion and enthusiasm for Shout!.

References

1. Lyons D. *Disrupted*. New York: Hachette Books; :114-117.
2. McGarry C. Twitter bots, fake retweets rake in big bucks. *PCWorld*. 2016. Available at: http://www.pcworld.com/article/2033766/twitter-bots-fake-retweets-rake-in-big-bucks.html. Accessed May 20, 2016.
3. The Rural School Community Center. *The Annals of the American Academy of Political and Social Science*. 1916;67:130-138.
4. The Forms of Capital. *Handbook of Theory and Research for the Sociology of Education*. 1986:241-258.
5. What is capital? definition and meaning. *BusinessDictionarycom*. 2016. Available at: http://www.businessdictionary.com/definition/capital.html. Accessed May 17, 2016.
6. 2014 Jobvite Job Seeker Nation Study: An authoritative survey of the social, mobile job seeker. *Jobvite*. 2014. Available at: http://web.jobvite.com/rs/jobvite/images/2014%20Job%20Seeker%20Survey.pdf. Accessed May 20, 2016.
7. Banerjee A, Duflo E. Poor Economics: A Radical Rethinking Of The Way To Fight Global Poverty. New York: PublicAffairs; 2016:195.
8. Keeley B. Human Capital. Paris: OECD; 2007:102-103.
9. Aziz A. *Modern Romance*. New York: Penguin Press; 2015:24-30.
10. John C. Marital satisfaction and break-ups differ across on-line and off-line meeting venues. *PNAS*. 2016;110(25):10135.
11. Putnam R. *Journal of Democracy*. 1995:65-78.
12. Wellman B and Gulia M. Net Surfers Don't Ride Alone: Virtual Communities as Communities. *Communities and Cyberspace*. 1997.
13. Sajuria J. et al. Tweeting Alone? An Analysis of Bridging and Bonding Social Capital in Online Networks. *American Politics Research*. 2016.
14. Bosker B. What Do We Retweet- And Why?. *The Huffington Post*. 2016. Available at: http://www.huffingtonpost.com/2010/12/16/what-do-we-retweetand-why_n_797369.html. Accessed June 7, 2016.

15. 5. Suh B. Want to be Retweeted? Large Scale Analytics on Factors Impacting Retweet in Twitter Network. 2010.
16. justretweet. *Justretweetcom*. 2016. Available at: http://www.justretweet.com/. Accessed May 20, 2016.
17. Kevin S and Mike K. *Postforpost Instagram*. Postforpost; 2016.
18. Socedo | B2B Demand Generation Solutions. *Socedo*. 2016. Available at: http://www.socedo.com/. Accessed June 7, 2016.
19. Lee K. The Social Media Frequency Guide: How Often to Post to Facebook, Twitter, LinkedIn And More. *buffersocial*. 2016. Available at: https://blog.bufferapp.com/social-media-frequency-guide. Accessed May 20, 2016.
20. Motti J. Twitter acknowledges 23 million active users are actually bots. *Tech Times*. 2014. Available at: http://www.techtimes.com/articles/12840/20140812/twitter-acknowledges-14-percent-users-bots-5-percent-spam-bots.htm. Accessed May 20, 2016.
21. Connecting or revoking third-party applications. *Twitter Help Center*. 2016. Available at: https://support.twitter.com/articles/76052. Accessed May 20, 2016.
22. Arthur C. Twitter horrifies third-party developers with warnings of cutoffs. *The Guardian*. 2012. Available at: https://www.theguardian.com/technology/2012/aug/17/twitter-third-party-developers-cutoffs?newsfeed=true. Accessed May 20, 2016.