

PRE-PRINT Version (September 12, 2004)

© All Rights Reserved. Will be Published in HICSS-38 (Hawaii International Conference on System Sciences), 2005.

Bounded in Cyberspace: An Empirical Model of Self-Regulation in Virtual Communities

Karine Barzilai-Nahon

Assistant Professor, Information School, University of Washington

karineb@u.washington.edu

Seev Neumann

Professor, The Leon Recanati Graduate School of Business Administration, Tel-Aviv University

neumann@post.tau.ac.il

Abstract

This article aims to focus on self-regulation mechanisms that are exercised by virtual communities in order to maintain their boundaries of autonomy while directing behavior of their members. An empirical examination of policy and enforcement of discarding inappropriate messages was conducted on 715 virtual communities. The study extracted lessons regarding the efficiency and efficacy of self-regulation in virtual communities. Moreover, it is argued that social capital and the history of users can serve as important factors while constructing self-regulation processes.

1. Introduction

Virtual communities face many challenges on a daily basis. Most challenges concern the desire of communities to preserve homogeneity and coherency. The literature names some of these challenges [1-3], for example, dealing with users who aim to flame and harm the norms of the community and quelling situations of anarchy and untidiness aiming to destroy the community. Communities exercise different mechanisms to maintain order and control in their boundaries, such as filtering and blocking software, security mechanisms, self-regulation mechanisms and others. This paper seeks to provide an explanatory model of activities of self-regulation.

The concepts of self-regulation, regulators and borders in the Internet are based on the presumption of the Internet as cyberspace. The metaphor of the Internet as a space conceptualizes the Internet as constructed by social practices between and among the objects of that space.

There is no meaning to regulators in virtual communities unless they have borders to protect and a space to operate in. Yet, the concept of regulation in Internet is different than regulation in the nation-state. The notion of national sovereignty is rooted in the ability to exercise authority in a defined territory over a defined population [4-6]. Authority is anchored in the ability to rule legitimately by controlling the means of violence and to attain obedience. There is no question that the classical meaning of sovereignty and authority undergo a transformation in the context of networks where information systems play a major role. The meaning of a defined territory and population assumes a new shape. The traditional sovereign, the nation-state, may not be the usual sovereign in cyberspace, and sometimes may not be a sovereign at all. Governments are currently in the midst of a regulation process over various issues relating to the Internet that may pose a threat to their sovereignty. The Internet enforces a frequent situation where increasingly governments must rely on mechanisms of self-regulation by non-governmental actors because of their inability or low efficiency in enforcing order and law in cyberspace [7-10]. Even in authoritarian regimes, where control is tighter, it is still not absolute [11-14].

In democracies, even when virtual community enablers declare that they operate according to state law, only few laws are actually enforced and applied. In most cases the daily business of virtual communities is mostly autonomous. Occasionally, the enablers exercise authority reactively, as a reaction to an event that forces them to act and intervene in the affairs of a community. Virtual communities are aware of this challenging situation and are constantly endeavoring to maintain an autonomous virtual existence together with keeping an order within

their borders. This is done through controlling the any information flow in any format (e.g., simple messages, pictures, sound track), which enters and circulates in the virtual communities' infrastructure via the mechanisms that we have described before. At this junction of the virtual community, the enablers, and state law, state sovereignty has much less of its original meaning, whilst the communal virtual space is being generated by the users, who trigger control unexpected alongside the regulation exercised by the managers.

Virtual communities face the challenge of having no physical borders. They cannot rely on frontier guards, since there are no physical borders to protect, and are required to develop and maintain mechanisms that protect the virtual borders of their communities. Protection is required in three areas: protection from an *unwanted* penetration from the outside, protection from an exit from the inside, protection for a *satisfying* interaction without the risk of chaos or internal strife. In this paper we try to identify major elements that influence one type of activity of self-regulation in virtual communities, that is, discard of messages.

2. Self-Regulation Mechanisms in Virtual Communities

In general, self-regulation is activated by delegation of responsibilities to private or semi-private bodies. It is often achieved through the collaboration of institutions and organizations in both the private and public sectors, and enforced by professional authorities over members of networks. In virtual communities, self-regulation may be considered as one of the major leverages to maintain a regular flow of information in any form, and refers to collaboration between the platforms enablers, the community managers, the members of the community and sometimes even official regulators of the state. So, we claim that self-regulation should be conceived through several levels of referral. On the one extreme one can find a more external governance type of control, which is more formal and in many cases more alienated to the community. Many times regulation in this sense is exercised via traditional mechanisms of regulation and not through self-regulation mechanisms. The state regulator and enablers frequently play this role. On the other extreme one can find a more internal governance type of control, which relies mainly on self-regulation mechanisms, and is more contingent upon the community context and culture, and therefore is less formal. Respectively, self-regulation in virtual communities may appear in various forms [8, 10, 15]. On the one hand we can look at external cooperation of regulators and generation of measures against users that breach the law,

while on the other hand we may look at policies enforced upon members of the community, construction of educational kits or designation of members to serve as mentors to new comers. Following, we will dwell into the roles of the community enabler versus the community manager, in order to try and construct a model of self-regulation.

A virtual community enabler provides a platform for conducting a virtual community, namely, a *virtual settlement* [16]. Enablers are not only technological facilitators; usually they also outline a policy and provide a general code of virtual behavior and norms to follow on their platform, procedures and rules that determine much of the online activity of the virtual communities. In many instances virtual communities override the general rules provided by the enablers and provide more specific and fitting regulation mechanisms to their members. Usually, the specific policy guidelines provided by the community do not contradict or conflict with the general ones specified by the enablers, and serve mainly to emphasize or elaborate certain aspects that are important to a specific community. For example, an enabler may prohibit any unlawful activity on the site that infringes on state law, and a specific virtual community engaged in games may add an internal rule stating "any discussion of serial numbers of stolen software will be automatically deleted".

A virtual community enabler is a regulator primarily in the sense of controlling the platform and creating the virtual settlement where the community runs its virtual life. Usually, enablers do not intervene operationally in content management, and leave this to the virtual communities' managers. They seldom act as gatekeepers of content, and when they do it is mainly in reaction to events or to a specific call asking for their intervention (e.g., a court ruling that requires them to remove some content from their platform). Thus, the everyday operational regulation of behavior is under the responsibility of the managers. They have the full authority to act upon the information that enters and circulates within their communities.

Depending on the enabler's policy of intervention, tension and conflicts may develop between the two hierarchies of self-regulators: the enablers and the managers, the more external oriented versus the more internal oriented. These arise, since in some cases enablers act contrary to the general will of the virtual community as reflected in the community-specific code or policy guidelines. Additionally, it happens because of different goals of the two entities: enablers are interested in increasing their revenues and therefore strive to add new users to their platforms, or to increase the popularity

of a certain community that has a high potentiality to be prominent or a high ratio of users and traffic; On the other hand communities are interested in preserving their quality and focus on their core subject (e.g., maintaining small communities, selecting engaged members). During these types of conflicts, managers of virtual communities serve as representatives of the community and as gatekeepers vis-à-vis the enablers and try to protect the content which is injected from outside by the enablers into their communities. In most cases this content is extrinsic to their norms, and sometimes can constitute danger to the existence of the community. Thus when an enabler invites members of a certain community to join an event in another community, the community at this point risks the exit of some of its members to another community, and also risks the shift of loyalty from one community to another.

Conflicts do not arise only between the enablers and managers, that is between the external governance and internal one. Tension may also arise internally, between the managers of the community and its members. Within the context of the current empirical analysis, the role of the managers is of special interest because of the expected tight relations between them and members of the communities. Managers use several mechanisms to control and direct behavior in their communities: censorship; editorial; channeling and regulation mechanisms [2]. This study analyzes a self-regulation model in forums, while focusing on one element of the censorship mechanism, the ability of managers to discard messages. In forums not all information passes through the manager for approval. Forum managers usually serve in this role as volunteers and need to be virtually present in the community only for a few hours a day, so they cannot serve as a focal point for all the information flow that passes in the community. Consequently, most tasks of the managers are *ex post facto* and are based on obedience of the members to the self-regulation procedures of the community. The role of self-regulation and the compliance and cooperation of the members of the community have a significant weight on the ability of the community to survive. Therefore, information control is performed mostly after members have posted the information, and mainly through discarding messages which are conceived incompatible with the norms of the community. Discarding messages is mostly executed whilst considering specific policy guidelines that belong to a community.

Additionally to the *ex post facto* nature of information control in forums, also the anonymity of the members creates a challenge for the regulator/manager [17]. Anonymity makes it easy for undesired users to enter the community. One may argue that in some cases the voice

of marginal members or undesired users might be useful for the community. This paper, however, does not deal with the normative question that concerns the health of the community, but rather with the ability of the community to execute regulation mechanism inside its boundaries. So *Ex post facto* information control aside along with anonymity allows users and members of the community to exploit it and post messages that might harm the community. When information control is eventually exercised, it might be too late, as the offense and turmoil may have already taken place [3]. Moreover, since the manager is not always present during the communication processes, community members occasionally take on the role of regulator in an attempt to maintain and enforce the internal rules of the community and to preserve what they consider to be its general will. It is important to note that when we refer to “undesired users” or “harming the community”, we refer to it as the community perceives it. Undesired users in some cases may be users that hold different views than the main stream of the community or the hegemonic discourse. The community might perceive these users as a threat, although the open discourse might actually strengthen the community health.

Next we will introduce our empirical study that focuses on an explanatory model which explores to what degree self-regulation, through discarding messages executed by managers of the community, is efficient in generating communal discipline.

3. Methodology

The study examined 715 virtual communities using data of up to three years back regarding the characteristics of the users, forums and messages. Most of the data needed for this research was provided by the community enabler, and some was constructed specifically as part of this study. The study used a data mining technique based on logistic regression combined with content analysis to construct the explanatory model. Therefore, it combined qualitative and quantitative methods simultaneously.

Data mining techniques may be regarded as bottom-up approaches that at their core lies the process of constructing a model to represent a dataset “Data mining is the analysis of (often large) observational data sets to find unsuspected relationships and to summarize the data in novel ways that are both understandable and useful to the data owner” [18] Choosing data mining as the main methodology strategy included the following reasons:

One, data mining technique may be more appropriate to apply to exploratory studies. There is no firm model regarding self-regulation in virtual communities. Many

cases were examined qualitatively, without a possibility to generalize and suggest a solid theoretical framework. Moreover, the focus on the deletion process as a self-regulation mechanism is new. In such cases it is more appropriate to use data mining because “the essence of data mining is the possibility of serendipitous discovery of unsuspected but valuable information. This means the process is essentially exploratory” [19]. Data mining extracts and discovers relationships from the data itself, a bottom-up approach, rather than forcing a relationship framework, a top-down approach.

Second, data mining is one of the few techniques that can easily deal with large amounts of data. In this research, a major part of the data was extracted through automatic tools and programs specifically created for the purpose of this study, together with manual data created through content analysis methods. The data set contains large volume of data in the form of a matrix of 80 variables and 1.385 million records, so that the number of overall observations is very large ($N = 112.185M$). According to the literature, data mining as a method has an advantage over other research methods when there is a need to analyze extensive volumes of data: “when a small data set is involved with only several predictors, one can manipulate the data set manually using statistical methods to search for the combination of predictors and their transformations that best fit the data, according to some statistical criteria.” ... “But with a large data set, containing hundreds of potential features and tens of thousands of observations, the number of possible combinations of features to explore is enormous, and beyond the capacity of any given individual, even a group of statistical experts, to handle in any reasonable amount of time” [20].

Third, reliability of results with data mining is higher in such a type of study. The percentage of messages that were discarded relative to the population size is very small (in this research – 5.48%). In such cases it is more reliable to depend on existing and observable data than using other methods of data collection such as interviews and surveys.

The relevant optional variables participating the data mining process (80 variable) were chosen after a process of 47 semi-structural interviews with managers and enablers. The data contained all messages that passed in the examined 715 forums during one month. Finally, the dataset was constructed with variables that belong to three dimensions: characteristic message variables (e.g., time, length of content, etc.), user characteristics (e.g., user’s history, personal details, etc.), and forum characteristics (e.g., number of messages in the forum, forum type, etc.).

This study has adopted the approach of data preparation suggested by the CRISP-DM model. According to this approach, preparation for data mining includes the following steps [21]:

- Select data – decide on the data to be used.
- Clean data – improve data quality to the level required by the selected analysis techniques.
- Construct data – constructive data preparation operations such as the creation of derived attributes and of new records, or the transformation of values of existing attributes.
- Integrate data – information is combined from multiple tables or records to create new records or values.
- Format data – primarily *syntactic* modifications to the data without changing its meaning, made necessary by the modeling tool.

To create and verify the model, the original data was divided into two datasets: a training dataset, containing 67% random sampling of the original dataset, and a validation dataset with 33% of the original dataset. The goal of the training dataset is to construct a model while using data mining techniques, in our case using logistic regression. The data mining software, GainSmarts, sift the important and significant predictors out of the many variables (80 variables). Once a model is offered using the training dataset, the same process is being executed on the validation dataset. Then the two models are compared with each other. The study offers a reliable model since the validation model is almost identical to the training model. The data mining process in this study illuminated 11 variables out of the 80 and suggested they were important in explaining the deletion process in forums.

4. Empirical Model for Self-Regulation

Overall, only a small portion of messages was deleted (5.48%). Most of the deleted messages (91%) were deleted due to three reasons: irrelevance of the message to the community (38.58%), attempt of a user to infringe on community culture (45.45%) and commercial information that was posted in the message (7.32%). Other reasons for message deletion were sedition for violence, unlawful activities, vulgarity, racism and outing. The examination included various types of virtual communities: expert groups, support groups, discussion groups, service providing groups, salon and acquaintance groups, interest groups and fan groups. The forums were also distributed randomly over different subjects (e.g., current events, culture, economy).

The model which was constructed through a data mining process while using logistic regression, suggested

that only 11 variables among the 80 may serve as predictors, and may explain the occurrence of deletion of messages. All 11 predictors were significant. Adjusted RR for the model is 0.3819. I will focus only on three predictors among the 11 that are relevant to the discussion of self-regulation. Theoretically it was justify to select these three variables for this paper since they enable a better understanding of self-regulation. Note, that also the other eight variables were significant and important, but less relevant to the understanding of self-regulation. The three variables we are going to discuss are: number of virtual communities in which a user participated during one month before the deletion occurred (Variable No. 41); proportional user seniority in a community (Variable No. 58); normalized number of deleted messages during the entire history of the user on the platform (Variable No. 79).

The results are as follows:

Number of virtual communities in which a user participated during the month before the deletion occurred:

Messages posted by users affiliated with one or two communities a month before deletion are less likely to be deleted than those posted by users belonging to three or more communities.

Proportional user seniority in a community: The more senior a user is, the lower the likelihood to experience deletion. Users who have been members of the community since the beginning of it are four times less likely to experience a gatekeeping event than new users. Senior members who are part of the community since the mid point of its life span are three times less likely to post a message that will be deleted than new users.

Normalized number of deleted messages during the user's history on the platform: This variable counts the number of deleted messages throughout the entire history of a user in all the communities with which s/he was affiliated proportional to her/his seniority. This variable is a very strong predictor. Users with a relatively larger number of deleted messages in the platform during their history are significantly more likely to experience a deletion event.

5. Discussion and Analysis: Self Regulation in Virtual Communities

5.1. Social Capital as Leverage of Self-Regulation

The explanatory model leads us to claim that social capital is critical in constructing self-regulation frameworks. There is an extensive literature on the issue of social capital and virtual communities [22]. While presenting evidence of the decline in social capital in the US in the physical world, Robert Putnam states in *The*

Other Pin Drops [23] that “the Internet may be part of the solution-- or it may make the problem worse... The Internet is about communication, and communication is central to community, not just etymologically but substantively. You don't get community without communication. The Internet certainly provides an opportunity to build what I would call social capital -- connections among people and, possibly, a deeper sense of reciprocity and trust”.

When discussing the issue of social capital in the context of the Internet, there are studies that support the view of the Internet as strengthening and supplementing social capital [24-27], while others see the Internet as decreasing social capital [28-31]. This study does not aim to compare or examine the role of the Internet in increasing or decreasing social capital relative to the physical world, but investigates the process by which virtual communities preserve or constitute social capital, a sense of inner sociability, civic engagement, and involvement.

Social capital in the Internet is mainly expressed through what can be called *the community core*, which determines much of the sovereignty of the community in cyberspace. A community that has a united and engaged core of members may be able to confront threats more effectively and help regulators or play the role of regulators by controlling information and directing behavior. The self-regulators need the discipline, compliance and support of members of the community to preserve the sovereignty of the community and therefore their power is not absolute. Variable 41, number of virtual communities in which a user participated during one month before deletion occurred, and variable 58, proportional user seniority in a community, and content analysis of messages and interviews reveal that communities without a structural core of members are more likely to rely on mechanisms of punishments (i.e., deletion of messages) and less on mechanisms of policy (i.e., community code) in context of self-regulation. Communities with a large number of veterans, or with members that do not participate in other communities and therefore devote all their time and energy to the community, who constitute a core of social capital, are less likely to experience actions of enforcement, and in most cases have a lower number of deletions in the community proportional to the total number of members there.

The results extract two main issues as important indicators in discussing social capital and self-regulation: seniority and overlapping loyalties. As mentioned above, the results show that members with seniority in the community are less likely to experience deletion of

messages (see variable 58). Senior and permanent members constitute the core of the community and are the main leverage of social capital in virtual communities, and leverage of an efficient self-regulation framework. In many cases they use their seniority to stabilize discussions and to calm contentions, and they convey their community knowledge and norms to new comers by sharing their experience with them. Often senior members take on the role of gatekeepers and protect the boundaries of their community. They remind members when a social norm has been breached and point them in the right direction of behavior. Since senior members are committed to the community and know its practices very well, they are less likely to experience sanctions from the managers. Through content analysis we could see that senior members used various methods to protect the community from outside intrusion and from an internal disorder. For example, if outsiders joined discussions with the intent to hurt the community, the senior members would direct other members either to ignore the intruders or to attack them verbally and drive them away.

Nevertheless, when analyzing messages that were deleted due to irrelevance of information, senior members are twice more likely to experience gatekeeping compared to newer members. This phenomenon can be explained by the sociability that may have developed over time among the seniors. Although many virtual communities are created around a specific issue that unite their members, over time relationships develop and in some cases senior members tend to shift the discussions to issues that are not relevant directly to the core issue around which a community was created and their messages are then deleted. Managers claimed in personal interviews conducted for this study that the objective in deleting messages because of irrelevant information is different from the deletion of harmful material that can hurt the community. They claimed that they had deleted irrelevant material mainly to clean the archives and maintain a sense of order for outside viewers. On the other hand, harmful material is deleted immediately without allowing the conversation to take place. This means that managers allow seniors to conduct discussions that are irrelevant to the community, and in most cases ignore these conversations and after a while delete them.

A second issue we examined in the context of self-regulation while exploring the issue of social capital is loyalty to the community. The results show that the number of communities in which a user is a member can predict the likelihood of that user to experience deletion of messages (see variable 41). Liphart [32] suggests that the presence of a source of overarching loyalties brings unity to a community. Which means, that a shared source of loyalty can create the *community core*, we have

discussed above. Is there such a source in virtual communities? We claim that there is. The narrative and subject of community is the source for loyalty of its members. But contrary to real offline life, where members are associated with a relatively small number of communities, in the Internet a member can be engaged with many virtual communities at the same time, which sometimes results in divided loyalties between them.

The results show that an increasing number of community engagements by individuals leads to an increasing likelihood to experience punishments and enforcement of regulation (see variable 41). This result supports the literature that deals with user attention on the Internet [33-36]. Following that literature and the results of this study one can argue that there is a limit to the attention and consequently engagements that one can devote to Internet communities. In most cases users choose to be members in only a few virtual communities. There is a suspicion that users who belong to many communities do not serve the *general good* of the community but exploit it for other purposes. For example, many users who advertise commercial information and experience deletions post their messages to many communities. A cross-posting of the same message requires them to sign up with many communities in parallel. Variable 41 shows that users with a large numbers of memberships are more likely to experience deletions because they are not actually a part of the community, and certainly do not constitute the core of the community.

To sum up, two contradictory forces related to social capital and boundaries protection of communities operate in virtual communities. On the one hand, senior members help to create a stable environment, contribute to the community's social capital and enhance self-regulation, even sometimes replace the regulator. On the other hand, the opportunities the Internet provides to members to be simultaneously associated with many social networks may cause potential harm to the social capital of the community and may create conflicts, which involve sanctions by the community as part of the self-regulation it poses on its members.

5.2. Looking for a “Clean” History of Users

From the results, we can extract important notions regarding users that experience enforcement of self-regulation. First, the history of members of the community is significant and may provide an indication of the users' future virtual behavior. Variable 79, i.e., normalized number of deleted messages during the entire history of the user on the platform, shows that users with more deleted messages throughout their history within all

communities with which they were associated, were more likely to experience sanctions.

The study shows that it was not enough to look at the behavior of a user a month before a deletion event in order to predict the occurrence of gatekeeping, except in cases in which the cause of deletion is commercial information. In these cases, one month was sufficient to provide a good indication of the chances of a user to experience deletion. Users who invade a community with the purpose of distributing advertisement and commercial material usually try to post the same message in many communities; after being discovered their usernames are blocked. Therefore, in most cases of commercial information, there is simply no history from which one can learn about the gated user; if there is one, it is usually short, such as one month.

A complete historical profile could give the enabler, managers and community a good indication as to whether a specific user can be regarded as a potential troublemaker or a troll. Note, however, that it is not enough to look at the history of a user only within a specific community but it is necessary to check the horizontal history of the user, in all the communities with which the user is affiliated.

Usually managers do not have access to the history of every individual member in their community through the information systems accumulated by the enablers, so why does the history of a user affect the chance of a message to be deleted? As communities become more mature, the core of the community and the managers assign an image to each member according to the accumulated activity within the community. This activity is also a big part of the process of self-regulation they construct. Managers handle a member who is considered to be a frequent troublemaker rigidly. In contrast to conventional expectations, the results demonstrate that despite a certain measure of anonymity that exists in the Internet, with members choosing to change their identity and online behavior more easily than in the real world, members of virtual communities adopt consistent patterns of behavior. Thus, users whose messages are constantly being deleted are likely to be irritants, and as long as this pattern continues, their messages will probably continue to be deleted. The results indicate that regardless of anonymity, virtual identities tend to exhibit a consistent behavior in cyberspace.

6. Conclusion: Efficiency of Self-Regulation

Self-regulation mechanisms may have advantages over state regulation by providing a more efficient way to

enforce norms and direct behavior in virtual communities [37]. The fact that only 5.48% of the messages and not more were discarded to begin with, may serve as local evidence to the efficiency of self-regulation. It seems that self-regulation, based on a high level of social capital, may hamper direct sanctions against community members. By creating the laws of the community from the bottom-up by the members of the community, there is more likelihood that members will comply with them, and in some cases, as demonstrated here, veterans will also help to enforce the policy. Through looking into deletion of messages activities in different forums, we were able to detect that in order to create an efficient process of self-regulation there is a need to construct a social capital in the community. That is, to create a core of members who are committed to the community, who exploit previous experience to elaborate policy, guidelines and norms important to the community and who convey these to new comers, and direct their behavior. In the study we found out that less social capital, i.e., less community core, resulted in exercising more sanctions as part of self-regulation. While more social capital strengthened the efficacy of self-regulation and discipline with no sanctions.

Control of a community and regulation may be facilitated by knowing the history of sanctions exercised against a member. Nevertheless, the procedure of self-regulation entails a danger. Self-regulation usually promotes the majority and not the minority, and may encourage dominance and suppression of other voices. A community, which promotes certain norms, may be efficient but also may become a very hostile space to users that think differently.

7. Bibliography

- [1] H. Rheingold, *The Virtual Community: Revised Edition*. Cambridge, Mass.: The MIT Press, 2000.
- [2] K. Barzilai-Nahon, "Gatekeepers and Gatekeeping Mechanisms in Networks," in *Management of Technology and Information Systems Program*. Tel-Aviv: Tel-Aviv University, 2004.
- [3] S. C. Herring, K. Job-Sluder, R. Scheckler, and S. Barab, "Searching for safety online: Managing "trolling" in a feminist forum," *The Information Society*, vol. 18, pp. 371-383, 2002.
- [4] L. Lessig, *Code and Other Laws of Cyberspace*. New York, NY: Basic Books, 1999.
- [5] M. Weber, "Politics as a Vocation," in *From Max Weber: Essays in Sociology*, C. Wright Mills, Ed. New York: Oxford University Press, 1946, pp. 77-128.

- [6] G. Zekos, "Internet or Electronic Technology: A Threat to State Sovereignty," *Journal of Information, Law and Technology*, vol. 3, 1999.
- [7] M. Birnhack and N. Elkin-Koren, "The Invisible Handshake: The Reemergence of the State in the Digital Environment," *Virginia Journal of Law and Technology*, vol. 8, 2003.
- [8] E. Brousseau, "Internet Regulation: Does Self-Regulation Require an Institutional Framework?," presented at DRUID Summer Conference, Copenhagen, 2002.
- [9] Y. Benkler, "From Consumers to Users: Shifting the Deeper Structures of Regulation Toward Sustainable Commons and User Access," *Federal Communications Law Journal*, vol. 52, pp. 561-579, 2000.
- [10] D. MacLean, "Herding Schrodinger's Cats: Some Conceptual Tools for Thinking about Internet Governance," Background Paper for the ITU Workshop on Internet Governance, Geneva February 26 2004.
- [11] Amnesty, "People's Republic of China: State Control of the Internet in China," Amnesty International AI Index: ASA 17/007/2002, November 2002.
- [12] A. Hermida, "Behind China's Internet Red Firewall," in *BBC News*, 2002.
- [13] S. Kalathil and T. Boas, "The Internet and State Control in Authoritarian Regimes: China, Cuba and the Counterrevolution," Information Revolution and World Politics Project, Carnegie Working Papers 21, 2001.
- [14] Reporters Sans Frontieres, "The Internet Under Surveillance: Obstacles to the Free Flow of Information Online," Reporters Without Borders July 2003.
- [15] H. Perritt, "Cyberspace Self-Government: Town Hall Democracy or Rediscovered Royalism?," *Berekeley Technology Law Journal*, vol. 12, 1997.
- [16] Q. Jones, "Virtual Communities, Virtual Settlements & Cyber-Archeology: A Theoretical Outline," *JCMC*, vol. 3, 1997.
- [17] J. S. Donath, "Identity and Deception in the Virtual Community," in *Communities in Cyberspace*, P. Kollock, Ed. NYC: Routledge, 1999, pp. 29-60.
- [18] D. Hand, H. Mannila, and P. Smyth, *Principles of Data Mining*. Cambridge, Massachusetts: The MIT Press, 2001.
- [19] D. J. Hand, "Statistics and Data Mining: Intersecting Disciplines," *ACM SIGKDD Explorations Newsletter*, vol. 1, pp. 16-19, 1999.
- [20] N. Levin and J. Zahavi, "Data Mining Digital Customers," in *Digital Marketing: Global Strategies from the World's Leading Experts*, V. Mahajan, Ed. New York, NY: John Wiley & Sons, 2001, pp. 256-281.
- [21] P. Chapman, C. Julian, R. Kerber, T. Khabaza, T. Reinartz, C. Shearer, and R. Wirth, "CRISP-DM 1.0," CRISP-DM Consortium 2000.
- [22] J. Preece, *Online Communities: Designing Usability, Supporting Sociability*. Chichester, UK: John Wiley & Sons, 2000.
- [23] R. Putnam, "The Other Pin Drops," in *Inc. Magazine*, 2000.
- [24] A. Kavanaugh and S. Patterson, "The Impact of Community Computer Networks on Social Capital and Community Involvement," *The American Behavioral Scientist*, vol. 45, pp. 496-509, 2001.
- [25] B. Wellman, Q. A. Hasse, J. Witte, and K. N. Hampton, "Does the Internet Increase, Decrease or Supplement Social Capital? Social Networks, Participation, and Community Commitment," *The American Behavioral Scientist*, vol. 45, pp. 436-455, 2001.
- [26] A. Kavanaugh, D. D. Reese, J. Carroll, and M. B. Rosson, "Weak Ties in Networked Communities," in *Communities and Technologies*, V. Wulf, Ed. Netherlands: Kluwer Academic Publishers, 2003.
- [27] P. Norris, "The Bridging and Bonding Role of Online Communities," in *Society Online: The Internet in Context*, S. G. Jones, Ed. Thousand Oaks, California: Sage Publications, 2004, pp. 31-43.
- [28] N. H. Nie, "Sociability, Interpersonal Relations and the Internet: Reconciling Conflicting Findings," *American Behavioral Scientist*, vol. 45, pp. 420-435, 2001.
- [29] J. Cummings, B. S. Butler, and R. Kraut, "The Quality of Online Social Relationships," *Communications of the ACM*, vol. 45, pp. 103-108, 2002.
- [30] S. Kiesler, R. Kraut, J. Cummings, B. Boneva, V. Helgeson, and A. Crawford, "Internet Evolution and Social Impact," in *The E-Business Handbook*, R. Watson, Ed. NY: CRC Press, 2002, pp. 189-201.
- [31] R. Kraut, M. Patterson, V. Lundmark, S. Kiesler, T. Mukophadhyay, and W. Scherlis, "Internet Paradox: A Social Technology that Reduces Social Involvement and Psychological Well-Being?," *American Psychologist*, vol. 53, pp. 1017-1031, 1998.
- [32] A. Liphart, *Democracies in Plural Societies*. New Haven: Yale University Press, 1977.
- [33] M. Van Alstyne and E. Brynjolfsson, "Communication Networks and the Rise of an Information Elite: Can Communication Help the Information Rich Get Richer?," presented at International Conference of Information Systems, Amsterdam, The Netherlands, 1995.
- [34] M. Goldhaber, "The Attention Economy and the Net," in *FirstMonday*, vol. 2, 1997.

- [35] E. Hargittai, "Standing Before the Portals: Non-Profit Websites in an Age of Commercial Gatekeepers," *The Journal of Policy, Regulation and Strategy for Telecommunications Information and Media*, vol. 2, pp. 537-544, 2000.
- [36] L. D. Introna and H. Nissenbaum, "Shaping the Web: Why the Politics of Search Engines Matters," *The Information Society*, vol. 16, pp. 169-185, 2000.
- [37] H. H. Perritt Jr., "Cyberspace Self-Government: Town Hall Democracy or Rediscovered Royalism?," *Berkeley Technology Law Journal*, vol. 12, 1997.