Reconceptualizing the Role of Security User

L. Jean Camp

Abstract: The Internet is not the only critical infrastructure that relies on the participation of unorganized and technically inexpert end users. Transportation, health, waste management, and disaster preparedness are other areas where cooperation between unorganized citizens who lack experience with the domain has increased resiliency, reduced social costs, and helped meet shared goals. Theories of community-based production and management of the commons explain this type of cooperation, both offline and online. This essay examines these two complementary approaches to organizing the cybercitizen for cybersecurity. Cybersecurity discourse has reasonably focused its attention on centralized parties and network operators. From domain name registrars to network service providers, solutions are sought through incentives, regulation, and even law enforcement. However great the ability of these centralized entities to implement change, the end user plays a crucial role. The Internet must remain open to enable innovation and diffusion of innovation; thus, the end user will also continue to be important. What is the role of the citizen in cybersecurity? What socio-technical characteristics might enable a system that encourages and empowers users to create a secure infrastructure?

How can cyberspace be augmented or organized so that security is more widely produced at home by citizens who lack technical expertise? Answering this question is critical to the governance of the Internet. When one person or computer is not secure, any or all of the people connected to the Internet potentially pay a cost. The average user receives spam because other

average users allow their computers to host spammers. Securing cyberspace is an inherently cooperative venture.

A growing body of work illustrates how classes of goods are being constructed by a collective (*community-based production*) and how shared resources are managed (*managing the commons*). When viewed from the first of these two perspectives, security is a good that can be cooperatively produced. When viewed from the second, computer security seems a shared common good that can be consumed and preserved but not produced cooperatively. The Internet as a commons can be compromised if too many people accept a high level of insecurity. In both cases, requirements for the nature of the good, whether public or private, must be defined. Cybersecurity is a good with significant private incentives; in the same way that no one seeks to become ill, no one wishes to be the victim of identity theft. Cybersecurity may also have a tipping point, after which a herd effect motivates action or adherence. On the network as in the realm of public health, herd immunity is needed to prevent epidemics.

In this essay, I describe both resource management as community-based production and as the management of the commons, and suggest how their underlying requirements may already exist or could be created.

Community-based Production

Community-based production refers to the self-organization of community members to create a good or service.[1] This model of social production differs from "crowdsourcing," in which a single entity, such as a firm, handles organization and management of a collective effort.[2] I begin by considering the case of community-based production of security information by self-correcting experts, then argue for computer security as a good that could be enhanced, though not fully implemented, by community-based production.

In his 1937 article "The Nature of the Firm," economist Ronald H. Coase explained why

individuals self-organize into firms with high degrees of specialization.[3] Adam Smith's *The Wealth of Nations* illustrated the value of specialization with the famed example of pin production.[4] Arguments for producing a particular good or service through this approach span centuries and cannot be adequately cataloged in this essay. In contrast, intellectual theories and reproducible analyses that explain and argue for community-based production as sustainable (and preferable in some cases) have emerged more recently.

Community-based production can drive the creation of any good that possesses the following characteristics: modularity, low capital requirements for entry into production, low marginal cost of production, and well-defined interfaces or interactions resulting in low cost of integration. The concept of community-based production was first defined in relation to the economics of the software running the Internet, because that software must be shared, readable by everyone, and have no cost. Otherwise the standards would be a barrier to connection. Software which can be read and used by everyone is called open source software.[5] Finance scholar Josh Lerner and economist Jean Tirole have described how individuals self-identify by contributing to tasks for which they are uniquely suited. [6]Lerner and Tirole argue that increased income is not the sole incentive for the production of open source software; rather, social intangibles such as respect, reputation, and the knowledge that one is the first to solve a particular problem are significant motivators. Contributions must be visible to the community for these incentives to function. In the realm of cybersecurity, neither the benefits of securing a computer nor the costs of failing to do so are visible, even to a computer's owner, thus reducing the incentive to take action. Nonetheless, there are multiple examples of loosely coordinated online communities that produce security information, including *vulnerability sharing* and *probe networks*.

A vulnerability is an error in coding or installation which enables unauthorized access to electronic resources. More specifically, a is a technical flaw allowing unauthorized access or use; where the relationship between the flaw and access allowed is clear; which has been documented to have been used to subvert a computer.[7] Information about vulnerabilities can be held secret, shared openly, or sold to a company. Vulnerability sharing refers to the second choice.

Probe networks are sets of networked computers that have no services, so there is no legitimate reason to connect to those computers.[8] Instead, the connections these receive are from massively parallel attacks, where the attackers are trying every feasible IP address. There computers are run by a social network of system operators who share statistics on these broadcast attacks, identify patterns as well as anomalies, and provide real-time distributed-network monitoring.

There are examples of cooperative distributed community efforts and for-profit firms for both vulnerability sharing and probe networks.

Vulnerability sharing has long been community-produced. The email list Bugtraq provides a mechanism to report vulnerabilities and rate the resulting report. Users of the list contribute knowledge, discuss vulnerabilities, and devise patches, and a large amount of computer security information is generated and validated by these volunteers with no obvious incentives; they do not have property rights to the information they provide or to the products that their critiques help improve. Like any informal system of social reputation, earning credibility is difficult, time-consuming, and uncertain. In contrast, when information production is managed by a firm, there is no basis on which to argue that such self-organization would occur. The HP service TippingPoint is in effect a firm created to purchase vulnerability information, with the result that distribution of this information is constrained. Thus, the

application of property rights to vulnerability information decreases social welfare [9] Yet Bugtraq's open system continues (with eight postings on April 26, 2011, for example).

The DShield project is an example of a probe network. DShield is a community of system operators who run network monitors or "probes" to detect intrusions and attacks. Information gathered by network probes is shared among the members, who then collectively analyze what is observed. The DShield project enhances the reputation of contributors by identifying them as such; increases the perceived value of its parent, Euclidean Consulting; and enables the identification of Internet Service Providers (ISPs) that are responsible for the worst attacks. DShield also provides a free self-help service that allows individuals to verify whether their IP addresses are associated with attacks. In this model, expert members of the community provide and generate security information at no charge and without claiming property rights.

Contributors to network probe systems are valuable to the entire Internet community and provide aggregate data that no individual could produce. However, there is no monetary incentive to participate. An individual who includes his or her own data in the aggregate may potentially benefit from a marginal increase in the ability to detect an early worm attack. Nonetheless, for the purpose of determining the pervasiveness of attacks or having an early warning system in place, the self-optimizing choice is to ensure that one's network neighbors participate without participating oneself.

While DShield provides a public list of malicious sites, other organizations have a more traditional, centralized production method. For example, WebSense offers security-based blocking of sites that it locates using its own honeypots. (A honeypot is a computer on the network that is built to be attractive to attackers by a combination of weak security and tempting file names. For example, the computer might be running old versions of software with known vulnerabilities and have fake files with names such as "creditcardnumbers.xls" or

"AccountsAndAddressess.db".) Individuals who do not subscribe to WebSense do not obtain the lists of compromised sites. Thus, WebSense offers security production in the more conventional model: that is, via a firm.

The previous section described community based production in two general cases when the participants are part of a community consisting of highly educated network administrators, engineers, and researchers These are examples of collective activities by security experts and network engineers committed to a more resilient Internet, not the creation of security information by non-experts. Is the latter practical? Which types of information goods can a community of non-experts feasibly create, under what conditions, and by whom? The question for the next seven paragraphs is the extent to which community-based production can be generalized. This begins with scholarly work, and then moves to possible examples where system design can enable the production of security by the non-technical computer user. Legal scholar Yochai Benkler's research on the production of trusted information describes the goods and services that can be effectively produced through community-based means.[9] In certain contexts, this model has many advantages over firm-based production. By altering the modularity, granularity, and cost of integration of the good that is produced, the cooperative model can shift the distribution of production costs to those most able and willing to bear them.

Challenges that can be effectively addressed by community-based production are modular rather than continuous; that is, they are characterized by clearly delineated decision points and explicit requirements for participants. For example, while prosecuting organized online crime is a broad objective best addressed by a highly organized governing body or firm, the task of identifying malicious websites is amenable to community-based solutions, because each web site is distinct, and the task is well-defined. Recall the task here is not to determine if the website is providing good and correct information but rather if it is a masquerade site (also called a

phishing site) or if it is distributing malicious code, and attacking visitors. Of course users can

currently recognize phishing sites and attempts to coerce downloads of

malware or such attacks would be flawless. Masquerading sites are recognized correctly by

non-technical users than by experts. However, if users are given simple information, such

as their own browsing history, and are told that this is a first visit to a site, and that site

claims to be Bank of America, then it is a phishing site. That is, users know that a site they

have never visited before is not one where they should enter their banking passwords. The

computer knows the site is not part of the users' history. The user does not. Telling the

user before the password is transmitted would be a straight-forward change in technical

communication.   Obviously you do not have a password for a site you have never before

visited. Similarly malware sites have very short lifespans, and telling the user that a site is a

day old, certified by a rarely used certification site in a remote locale, and that the user has

never taken such a risk could enable the user to make an informed decision.

The cost of integrating individual distributed efforts must also be limited for this model to

work. Benkler mentions software as an example in which integration can occur with a well-

documented application program interface; his counterexample is aircraft construction, which

requires an exacting physical integration of many components. What elements of cybersecurity,

like aircraft, require regulation and coordination? Which can be enabled from the bottom up?

Design parameters change when systems are implemented to enable peer, as opposed to firm or

governmental, production. In terms of participant requirements, considerations include lifestyle

(for example, whether the user must be connected to the Internet for 85 percent of the day) and

the type of response required (for example, whether a system simply reacts to alerts or requires

constant monitoring and an endless attention span).

The identification of malicious websites illustrates the potential uses and limits of

community-based production. Whereas legitimate banks can be appropriately identified in a centralized fashion, by the Federal Depository Insurance Corporation and the National Credit Union Administration, the rate by which unknown websites are increasing prevents any centralized entity from identifying all of them. Yet online behavior in the user community suggests a potential solution. The vast majority of sites that individuals visit in a browsing session are the ones they visited in previous weeks. One study[10] of browsing history over a four-week period found that within a social network of only ten people, more than 99 percent of all participants' clicks led to previously visited sites. Only one in a hundred clicks brought individuals to a site identified as unknown. (For the average individual in the study, 95 percent of clicks led to familiar places). By reconceptualizing the global issue as a community problem, the study uncovered new and potentially untrustworthy sites without compromising user privacy.

Attackers have long user social networks to enhance attacks. The "I love you" virus was the first malicious use of address books; today attackers harvest Facebook and blog comment sections. In contrast, today's defenders use centralized solutions that must address the vast heterogeneity of the entire Web. Given the sheer scope of the challenge, the identification of individual websites as new, and thus suspect, is best done by community members. Yet the long-term efficacy of the community model depends on the capacity of centralized institutions to coordinate the identification and takedown of malicious sites in a timely manner–that is, before the sites build reputations. Both community-based and centralized production are necessary; neither is adequate on its own.

In assessing whether web sites or individual users are trustworthy, community-based production can incorporate implicit, behavior-driven ratings or explicit, personal recommendations or selections from trusted parties. The amount of time an individual spends using, and therefore contributing to, various resources is another implicit measure of

trustworthiness. Social trust reduces technical complexity[11] and can alter the nature of cumulative risks taken, in terms of system failure, privacy, and even threats inherent in system operation. Community production achieves a governance system that either could not be accomplished by a centralized agency or could not be done accomplished without very large-scale investment of capital.

Community production recognizes the incentive individuals have to maintain their own information. Individuals' ability to protect themselves against the risks they take–that is, the capacity to shift costs to those with the greatest incentive to bear them–applies to malicious sites and many types of computer subversion. However, one difficulty of cybersecurity is that while the individual bears the cost of some computer subversions, in other cases the costs are borne by others. Again, the creative use of social networks and incentives can be applied here to help develop a more robust and resilient infrastructure.

The Commons Without Tragedy (Or Government)

Today large numbers of computers, called "zombies", are brought under the control of a single centralized entity, forming a "botnet." Zombies do not necessarily steal the information of affected computer owners. Rather, access to the computer provides a launch pad for attacks, a storage resource, and a safe space away from virtual home. The fact that the malicious controllers of these computers are centralized has triggered centralized solutions, such as coordinated law enforcement. However, the subversion of these computers works from an economic perspective because of the large supply, high connectivity, and very low marginal cost of hitting tens of millions of computers to find hundreds of thousands that can be subverted. One possible solution for the cyber-commons is a ground-up approach that would induce secure behavior in communities and subnetworks.

A critical component of any ground-up approach will be the use of preexisting social trust and risk communication to create subnetworks that seek to change individual and group behaviors. A range of powerful authentication technologies have yet to be applied to the challenges of securing devices (including proximity authentication, for example, which works only for devices that are physically collocated). Virtual neighborhoods created from secure group formation and physical neighborhoods authenticated by proximity are examples of possible subnetworks where effective interaction design combined with social transparency can enable neighborhood self-defense.

Elinor Ostrom has illustrated that effective governance of shared resources requires can emerge under given conditions. In her 2003 summary in Science, she lists five conditions: 1) the monitoring of resources and their use, 2) moderate rates of change in social and economic conditions as well as user populations, 3) the existence of social capital in the community, 4) the ability to exclude outsiders from the community, and 5) user-supported enforcement of norms.[12] In theory as well as in practice, creating these conditions requires the development of secure systems that are designed as social networks.

It is also important to consider the relationship between social networks and herd immunity. If policy cannot change the behavior of all users, what category or number of people must be encouraged to change? Can visibility of low-security choices be leveraged to create the transparency necessary for self-governance of an Internet-security commons?

Monitoring Resources

The first of the five conditions listed above, monitoring resources in a shared domain, is one of the simplest but most underused governance mechanisms. Information monitoring has a trivial effect on the information infrastructure, and simply providing information can be a potent agent of change. Individuals are rarely capable of monitoring their own network experience, and yet

there are few available interfaces for monitoring network resources. Even organized efforts have difficulty measuring resources available to individuals, with some comparisons of national broadband networks measuring nothing more than the parameters set by the computers for sharing bandwidth.[13] Individuals often do not know their own resources, and arguably, never have. The 2003 spread of the Slammer worm is an example of the challenges home users face in monitoring their own resources. As Slammer attacked SQL servers, most people were unaware that they might be vulnerable even if they knew of the worm's existence. Announcements specified that the worm attacked Microsoft SQL Server 2000, but how many users knew that their PCs, in fact, ran an SQL server? Any technically useful report could have been construed by the average user as acknowledgment that the worm did not apply to him or her. Today, few individuals who have broadband would be concerned that they have a web server in the house. Most would argue that they have no such thing. Yet every wireless hub has a simple web server to enable the owner to initialize and configure the device. The home owner does not link the fact that he or she set up the wireless router using a browser to the fact that there must be some server code running on the device. Transparency may have improved in terms of individual knowledge of their own resources but there is not evidence that awareness has.

Currently, some ISPs may notify individuals when the ISP detects or is notified that a computer on the ISPs network is subverted. It is commonly estimated that 5 percent of all computers on the Internet have been subverted[14], yet even the most aggressive ISP responses have offered recovery services to just 1 percent of subscribers.[15] For this reason, both end users and network service providers have limited awareness of the existence of botnets. One problem for a network service provider is the heterogeneity of the network population.

Concerns about individual privacy and close monitoring of network behaviors also limit network monitoring, as organizations that have been found to practice unwarranted monitoring

have faced, at the least, a media backlash.[16] For this reason, close monitoring by an agency may be impractical, but the homogeneity and consistency of individual behavior is an asset to the extent that home users can observe the actions their computers take on their behalf. Detecting a change in website visitation behavior across the entire network is very difficult. Detecting a change by a single computer which has very few (human and therefore non-random) users is a problem that is easier to solve. The individual computer is fairly consistent while the network is not. And when a computer becomes suddenly inconsistent, a home user will have better information on why that may or may not be suspicious. For example, a network service provider may observe changes in traffic behavior without knowing if it is because there is a family reunion taking place (and thus a dozen teenagers are on the wireless) or a computer has been subverted; the individual user, on the other hand, can easily distinguish between the two scenarios. Thus, while recovery services may need to be centralized, network monitoring and communication with individual users function best when decentralized. The global network is dynamic, rapidly changing, constantly reconfiguring routes, and profoundly heterogeneous; while the individual is a notoriously poor source of entropy. Monitoring resources at the end point means leveraging the innate homogenous humanity of the single user, as opposed to simply bemoaning the fact that humans produce bad, nonrandom passwords.

Home users face a plethora of add-ins, add-ons, and an ever-expanding lexicon of attacks and defense. A more productive approach would be to provide individuals with a single narrative and a clearly marked path to risk mitigation and recovery. Users could be informed of radical changes on one computer in the house either by other computers in the house or by the computers that participate in the social networks I describe below. Current technology is not designed to communicate, in an effective, carefully timed, and educational manner, the particular risks to which a user might be exposed; nor does it automatically change settings to respond to

personal contexts (work, play, banking) or technical ones (public wide area network, protected workplace, patched or unpatched).

Market forces, property rights, and even assigned identifiers can solve some of the incentivization problems related to computer security. However, an effort to control and enforce behaviors on the population in a "war on computer insecurity" risks being both ineffective and expensive. In contrast, making risks and decisions visible to individuals, thus enabling them to monitor their own computers, is a technological challenge that can be met without violent metaphors or intrusive monitoring.

Moderate Rates of Change

The second requirement for community management of shared resources is for moderate rates of change in the population. This requires stable communities and exclusion of others, not on the Internet as a whole but within communities. The notion that the Internet is open to all is a canard; exclusions is now and has historically been practiced online. The earliest exclusion was in email lists. The very existence of some of these lists was effectively secret (particularly from professors in some schools.) Some lists added members by invitation only; others allowed open subscription as well as banning; and still others embraced a simple no-holds-barred approach. For example, the early lists of mothers based on birthdate (e.g., August1990) have followed their offspring for twenty years now and are highly exclusive. The same models apply today in the blogworld. Blogs can be completely open, allowing unmoderated and immoderate anonymous comments. More often, they are slightly restrictive, allowing open readership with member comments or moderated anonymous comments. Many reserve access for members only, especially when members have strong shared experiences (such as surviving abuse) or have tired of defending the existence of the group itself, and simply want to be able to discuss the topic at

hand (for example, feminist blogs that exclude men's rights activists). Such closed blogs can be read only if an application for inclusion is accepted.

The second-oldest form of the closed online community is arguably the chat room, a service built on the idea of the single-identity provider. The chat room functioned primarily because of a large, AOL-installed user base and AOL's centralized governance ability. Chat rooms are also called "pull technology," meaning the individual must actively log in to participate. Before the chat room, there were closed mailing lists.

The Internet has long been applauded for its openness. Yet the network enables the creation of spaces that can be closed or even invisible to others. Social networking has enabled exclusion since the first days of email, when reply-all became reply-to-sender for the occasional snark. The implementation of stable Internet communities is widely managed by those communities across an array of platforms. Even with difficult interactions and exploitive privacy requirements, the story of the Internet is one of community formation. The ability to form communities, whether bound by physical location, shared interest, or sheer random selection of the moment, illustrates that the second condition for management of shared resources can be met. There are an unknowable number of groups, from high school classmates to their mothers who have been chatting online since the first positive pregnancy result, that meet the requirement of "moderate rates of change in social and economic conditions as well as user populations".

Social Capital

Consider the third characteristic necessary for shared governance of resources: the existence of social capital in the community. Discussions of security in economic terms–that is, as financial capital–have been active for the past ten years.[17] Security as social capital, however, is rarely considered. Most related technologies define security as an individual effort and presume that information is an individually owned resource.

A 2000 paper introduced the idea of computer security as a good with externalities.[18] Since then, models of various components of security-related externalities have been widely explored. In economic terms, the current crisis in computer security is a market failure. There is some agreement that components of cybersecurity are a public good.[19] Private security decisions have a public externality, as the cost of an insecure system is accrued by other systems that are subsequently infected as malicious computer code such as viruses and worms spread. Various solutions to this problem have been proposed, including liability,[20] insurance markets for business risks,[21] and enforcement mechanisms for ISPs.[22]

Security defined as a good has both public and private elements, but security proposals have tended to focus on the private aspect. In terms of the public good, solutions have emphasized monetary liability or insurance. Although proposals for liability could function for larger actors in the security market, where decisions (at least in theory) are driven by cost-benefit analysis, this approach would likely backfire in the realm of vulnerable home users. In fact, some proposals could increase the potential risks for individuals without providing any mechanisms for enabling them to avoid these risks. Increased liability, as a means to encourage individuals to cease insecure behaviors, is unlikely to be highly effective if individuals are unaware of being engaged in such behaviors. Consider that 5 percent of home computers may be infected. A regime that criminalizes users for having an insecure home computer would immediately transform some 5 percent of the online population from law-abiding citizens to enemies of the state. It is difficult to imagine an external attack that would similarly hinder or harm so many Americans.

A commonly proposed solution to an externality is the creation of a property interest that would enclose the information security space. A functional market would require adequate information, the ability to process this information, and a sufficient attention span. Currently,

security technologies provide none of these requirements. For example, when a public key certification is identified as invalid, the user receives as information two incomprehensible[?] "hash" values for calculation and comparison but no information about the source of the warning or the certificate. Only the rare brilliant mathematical genius could consider the option of comparing these two values as potentially useful. Yet this is the information helpfully presented to the end user to determine if the certificate should be trusted. Is there indeed a human on earth who can compare these numbers on the fly?

Fingerprints

SHA1 Fingerprint	6E:73:A8:C4:72:C2:97:74:68:9B:BC:F8:95:9D:02:1D:3A:54:5C:E8
MD5 Fingerprint	67:05:A8:45:B5:FB:9F:F0:D7:07:36:33:0B:58:DB:5E

Is the certificate invalid because it is signed by a university rather than a more widely recognized corporate provider of certificates? Is it invalid because it expired yesterday? Or is it signed by a previously unknown, and therefore likely to be malicious, party? Is it signed by a leading certificate provider for large corporate entities, or a rarely used provider favored by marginally legal organizations? Certificate providers have social capital and reputations that are well-known by those who read security literature and manage networks. For experts in the field, it is known that some certificate providers are simply more trustworthy than others. Good reputation is a form of social capital, but unless this is visible to normal users, it cannot be an effective part of collective decision-making.

Similarly ISPs have widely varying records of protection of individuals. Some ISPs simply let user computers drop onto blacklists, never contacting the owner. Other ISPs notify and assist customers whose computers are apparently infected. There is no place to locate this information, no consumer reports. Not only individuals, but also organizations have histories and social capital on the network. This information should be made much more widely available, and it certainly exists. To paraphrase Warren[23], that which is whispered in the halls of North

American Network Operators Group[1] should be shouted from the rooftops.

Better information monitoring would allow individuals to know whether their computers are responsible for spam. Were this information more readily visible, neighbors and friends would know, too. Computer insecurity, were the costs to others apparent, could become as socially unacceptable as littering: that is, it would exist, but to a much less egregious extent. Given that computer crime is driven by profit more than pride, making insecurity anomalous rather than ubiquitous may be adequate to stem the tide.

Exclusion

Recall that there are five characteristics that must exist for a resource to be, in theory, managed by the community. One of these is that individuals can be excluded from the community. For policy makers, this implies permanent exclusion. Permanent exclusion requires permanent, single identities. In this section I argue for and illustrate that exclusion has long been possible without the certain costs and stochastic risks inherent with single identifiers. (Of course, this complements the discussion of the first requirement on stable communities above.) I further argue that single individual true names or real names, in theory implemented by government, is in no way the best approach. Instead, stable pseudonyms in specific communities are more than adequate. Social capital requires a social context, and the nation is too large a context to be workable for any but the most famous or most dangerous criminals.

The use of social networks and the explosion of social communication illustrate the capacity of those on the network to implement change through reputation mechanisms of all levels of openness. While Facebook consistently alters user's control of their own profiles, there has consistently been a wide range of Facebook mechanisms that provide user control. Anyone who

---

[1] NANOG is the organization of ISPs. Members know each others' corporate habits and practices. Not only network standards but also network norms emerge in groups of network operators and engineers.

has ever "unfriended" another person has experienced the power of exclusion on the Internet. Similarly the new social network system from Google, G+, enables more detailed control, with groups of people in categories such as "friend", "acquaintance" or "family".

The major functions of security must include exclusion and control. As noted above, the most straightforward solution involves a centralized entity that has the authority to issue–and therefore revoke–accounts, enabling access control. Proposals for federated or trusted identities all follow the same logic: a single identity will enable access control and ensure responsible behavior. A similar logic applied to the recent cancelation of G+ accounts that were not based in names that Google determined were adequately true and real. Others argue that requiring individuals to use a single identity will create another tragedy of the commons, whereby identities are the overused and under-protected resource. A second line of objection is social rather than economic, that individuals who would be threatened in employment or communities for having unpopular views (e.g., feminists in Texas, fundamentalist Christians in the Bay) should be able to speak without their coworkers and employers knowing. This line of objection was the dominant one with the cancelation of G+ accounts based on lack of true names, with reasons for pseudonym use including the following:

- "I am a high school teacher, privacy is of the utmost importance."
- "I publish under my nom de plume, it's printed on my business cards, and all of the thousands of people I know through my social networks know my by my online name."
- "I have used this name/account in a work context, my entire family know this name and my friends know this name. It enables me to participate online without being subject to harassment that at one point in time lead to my employer having to change their number so that calls could get through."
- "I do not feel safe using my real name online as I have had people track me down from my online presence and had coworkers invade my private life."
- "I've been stalked. "
- "I'm a rape survivor."

• " I am a government employee that is prohibited from using my name IRL." (in real life).[24]

Is the price of effective governance global names? Or can exclusion exist in smaller communities? Certainly, the dual threats of privacy violations and misuse by (authorized or unauthorized) parties have not been adequately addressed by any federated or single-identity proposal. Another challenge with an Internet-wide or national solution is that security is a social activity in which the most malicious participants work to be the least visible. Identity thieves will not be hindered by the introduction of a new identity infrastructure any more than foxes would be hindered by an increase in the chicken population.

Exclusion already exists in Internet communities. Just as the Internet is the network of networks, it is the network of communities. Engineering solutions that enable governance in small communities can make a difference in the Internet as a whole, arguably more efficiently and certainly with better privacy. The Internet experience is one of physical disconnectedness but of social connection. For every highly verbose commenter who weighs in on a blog post, there are orders of magnitude of more silent readers. The anonymity of the Internet is easily violated, yet users act as if privacy were protected by social contracts communicated via website design.[25] These social contracts could be leveraged with security systems that address the shared costs of security, and encourage cooperative governance. Social contracts require exclusion for those constraints to be both binding and enabling.

Social contracts such as those enacted by the users of Facebook[26] existed in the "real" world long before the Internet came into being. In the real world, however, individuals can use the visual, geographical, and tactile information embedded in physical interactions to evaluate the safety, competence, and trustworthiness of those who control a physical space. For example, merchants offering high-quality products can charge a premium based on reputation and invest their profit in retail spaces that reflect their wealth and standing. These cues are not available

online because they have not been integrated into the network–not because they are impossible to engineer. Blacklisting and blocking are usable tools in email, on blogs, chat, and within social networking, and recommendation systems. Better engineering that enables self-organization for purposes other than superior advertising targeting is therefore possible; further, it may prove less costly and more effective than a punitive legislative approach.

Social Norms

The fifth and final requirement for the referenced discussion of management of communal resources is that there exists user-supported enforcement of norms. In the immediately preceding paragraphs I argued that exclusion is possible on the Internet, and that further it can be more effectively implemented at the smaller scale than the national or global. Before that I identified digital domains with social capital; preceded by the illustration that there are many communities with low rates of change using email, chat groups and blogs as examples. The first condition, monitoring of resources use, is the most simple on the information infrastructure. Here I join Lessig[27] is arguing for the existence and importance of norms in cyberspace.

Certainly there are community norms that can be altered or enforced by centralized control. In several domains–for example, online communities such as Facebook–information has multiple stakeholders. Acceptable use of knowledge is determined by implicit social and explicit corporate policy norms established through perceived imagined communities.[28] By contrast, current security technology is based on the mental model of the security expert who develops usable security controls for the individual, informed user. This approach limits the effectiveness of such controls by discounting both the social context of use[29] and the mental model of the end user.[30] Considerable research on the creation of norms has examined the development of social conventions in different online communities; material studied ranges from explicit sexual

interactions in Second Life (a three-dimensional virtual world where users can socialize) to pattern sharing through Ravelry (a site for knitters and crocheters). Norms must be widely accepted in online communities, but they need not be perfectly enforced. Norms that are violated on rare occasions remain norms, and just as challenges of cybersecurity are manageable without flawless enforcement. Actions that simply are not taken (for instance, there is no norm against nailing one's own foot to the floor or setting your computer on fire with your hair) are not suitable for governance by norms. Security controls could enable individuals to set their own norms for interaction. This would require empowering an individual to filter and accept different risks automatically when in different virtual spaces and communities in terms of liability, ownership, rights, and acceptable use.

Recent work has examined how to nudge users toward positive behaviors through inter-action design rather than with inescapable defaults or tedious reminders. Yet even on social networking sites, these efforts target the isolated user sharing his or her own information. In contrast, tools that encouraged communal norms and makes those norms visible in a user-defined community can encourage members of social networks (as discussed above) to comply with these security and safety norms.

Some of the language used to address computer security may, in fact, discourage compliance. Being a "pirate" may seem desirable by end users because of its appeal in popular culture. "Zombies" are imagined as inherently villainous, but they are also popular (consider the films *Zombieland* and *Shaun of the Dead*). "Phishing" is a purposefully obtuse word, created as an inside joke. Such nomenclature may prompt users to ignore security, or not to take it seriously. Certainly the language of computer security does not encourage norms of security adoption. Nor does the one size fit all design, yet security designs assume that interactions and defaults should be uniform for all people and across the entire browsing or Internet experience.

The term *computer security* is also deeply intertwined with the self-interested protection of copyright holders–to the detriment of many users. The goal of policing the individual directly conflicts with the objective of recruiting the user to participate in securing cyberspace. Disentangling these goals through community involvement, encouragement, and communication can enable security-enhancing norms to emerge. Pursuing these goals through ever-less-functional devices and more expansive definitions of felonies may effectively choke the Internet as an engine of American innovation, without ensuring security for the novice end user. Norms must align with the interest of the community to be adopted. Disentangling very different risks of sharing copyrighted material without permission or hosting a botnet which serves organized crime can encourage norms of security and digital safety.

Closing

I have presented two major arguments in this work. First, security is in part a good that can be cooperatively produced. This implies recruiting end users into the production of security. The second argument is that the Internet is a common resource and security is a component that requires shared management. That shared management requires five characteristics (using the selected model) that are not obviously present on the Internet. For each of the five characteristics (resource monitoring, low rates of change, social capital, exclusion, and norms) I have argued that while these are not available on the global Internet these certainly exist in Internet communities.

I conclude that the challenge of securing cyberspace cannot be met without the cooperation and coordination of the end user, and the conditions to enable this cooperation and coordination exist. While the cybersecurity community, in both technical mechanism design and political discourse, has concentrated on large-scale centralized entities socially-aware security engineering can make a profound difference. Currently, home computer users who seek to

remain safe face a patchwork of standards and corporate products. Each user has a unique set of challenges embedded in geography, health, social context, service provider, platforms, and other variables, such as home layout. One size cannot fit all, and one user cannot be expected to face the world alone with no community support.

Similarly, with respect to roadways, water management, and public health, the behavior of nonprofessional participants is critical. Not everyone speeds; littering, smoking, and drinking and driving were all socially appropriate behaviors at one time. These behaviors have been radically reduced by a combination of public education and sporadic enforcement. Changes in norms, more than any other factor, has led to positive change in behaviors.

Tackling more fundamental engineering problems is an important first step. These include the need to build systems that clearly communicate security state, guidance to improve that state, and directions to recovery assistance when a computer is compromised. Even in the most optimistic theory or authoritarian regime, individuals cannot be given effective incentives to accomplish tasks beyond their capabilities. Technology to communicate with and empower individuals to protect their own digital assets is sorely needed. Yet the business models for such technologies are uncertain, and the research so interdisciplinary as to be an ill fit with traditional funding agencies.

Scholarship in community production and cooperative research management can be combined with security engineering research to create a socially, and thus technologically, resilient Internet. Research and exploration into which challenges are well-suited to community-based production and which require centralized coordination is critical. Engineers and institutions can provide decision-makers with the socially-aware technical tools to empower families, individuals, and localities to enhance their cybersecurity. Social engineering has enforced top-down solutions, such as the Trusted Identifiers initiative and other designs for

intensive monitoring and control. I am advocating for a different approach, one in which engineering courts community rather than seeking to control it.

Identification of the potential and applicability of community production for cybersecurity at home can make a significant contribution to the total social cost of implementing cybersecurity on a national scale. Strengthening the network will require understanding the potential for citizens to self-govern, with regard to the protection of their own home systems and their personal information and identity.

L. Jean Camp is a Professor in the School of Informatics, Adjunct Professor of Telecommunications, and Adjunct Professor of Computer Science at Indiana University. Her publications include *Trust and Risk in Internet Commerce* (2000), *Economics of Information Security* (edited with Stephen Lewis, 2004), and *Economics of Identity Theft: Avoidance, Causes, and Possible Cures* (2007).

Acknowledgements

Endnotes

[1] Yochai Benkler, "Sharing Nicely: On Shareable Goods and the Emergence of Sharing as a Modality of Economic Production," Yale Law Journal 114 (2) (2004): 273–359..

[2] Karen S. Cook and Robin M. Cooper, *Experimental Studies of Cooperation, Trust, and Social Exchange* (New York: Russell Sage, 2003), 209–244; Thomas Dietz, Elinor Ostrom, and Paul C. Stern, "The Struggle to Govern the Commons," *Science* 302 (5652) (2003): 1907; Jeff Howe, "The Rise of Crowdsourcing," *Wired*, June 2006; Jeff Howe, *Crowdsourcing: Why the Power of the Crowd is Driving the Future of Business* (New York: Crown Business, 2009).

[3] Ronald H. Coase, "The Nature of the Firm," *Economica* 4 (16) (1937).

[4] Adam Smith, *An Inquiry into the Nature and Causes of the Wealth of Nations* (London: Methuen and Co., 1776).

[5], C. DiBona, S. Ockman, and M. Stone, eds. *Open Sources: Voices from the Open Source Revolution*, Sebastopol, CA: O'Reilly, 1999.

[6] Josh Lerner and Jean Tirole, "Some Simple Economics of Open Source," *Journal of Industrial Economics* 50 (2) (2002): 197–234.

[7] L. Jean Camp & Catherine Wolfram, "Pricing Security *The Economics of Information Security*, Chapter 2, pp. 17-34. L. Jean Camp & Stephen Lewis (Eds), Springer-Kluwer (NY, NY), 2004.

[8] John Bethencourt, Jason Franklin, and Mary Vernon. Mapping internet sensors with probe response attacks. In Proceedings of the 14th USENIX Security Symposium , August 2005; Edward Balas and Camilo Viecco. Towards a third generation data capture architecture for honeynets. In Proceeedings of the 6th IEEE Information Assurance Workshop, West Point, 2005. IEEE.

[9] Yochai Benkler, "Coase's Penguin, or Linux and the Nature of the Firm," *Yale Law Journal* 112 (2002);

[10] Zheng Dong & L Jean Camp, "The decreasing marginal value of network size", *ACM SIGCAS Computers & Society*, Vol. 41 No. 1 (2011).

[11] Kiran Lakkaraju, William Yurcik, and Adam J. Lee, "NVisionIP: Netflow Visualizations of System State for Security Situational Awareness," *Proceedings of the 2004 ACM Workshop on Visualization and Data Mining for Computer Security* (New York: Association for Computing Machines, 2004), 65–72; Russell Hardin, *Trust and Trustworthiness* (New York: Russell Sage, 2002); Niklas Luhmann, "Trust: A Mechanism for the Reduction of Social Complexity," in *Trust; and Power: Two Works* (Chichester, N.Y.: John Wiley and Sons, 1979), 1–103; Helen Nissenbaum, *Privacy in Context: Technology, Policy, and the Integrity of Social Life* (Stanford, Calif.: University Press, 2009).

[12] T. Dietz, E. Ostrom, and P.C. Stern. The struggle to govern the commons. *Science*, 302(5652):1907, 2003.

[13] David D. Clark, "Window and Acknowledgement Strategy in TCP," *RFC 813* (Cambridge, Mass.: MIT Laboratory for Computer Science, July 1982).

[14] M. Fossi, D. Turner, E. Johnson, T. Mack, T. Adams, J. Blackbird, M. Low, D. McKinney, M. Dacier, A. Keromytis et al., \Symantec report on rogue security software," Whitepaper, Symantec, October, 2009.

[15] M. van Eeten, J.M. Bauer, H. Asghari, S. Tabatabaie, and D. Rand. The Role of Internet Service Providers in Botnet Mitigation An Empirical Analysis Based on Spam Data. *Ninth Workshop on the Economics of Information Security*, 2010.

[16] A recent example is Apple's location information compilation; see Brian X. Chen, "Why and How Apple is Collecting Your Location Data," *Wired* blog, April 21, 2009, http://www.wired.com/gadgetlab/2011/04/apple-iphone-tracking (accessed April 29, 2011).

[17] Bruce Schneier, "We Don't Spend Enough on Security," *First Workshop on Economics and Information Security*, Berkeley, California, May 2002.

[18] L. Jean Camp and Catherine Wolfram, "Pricing Security," *Proceedings of the CERT Information Survivability Workshop*, Boston, Massachusetts, October 24–26, 2000, 31–39.

[19] Science and Technology Committee, *Personal Internet Security, Volume 1: Report*, Fifth Report of Session 2006-07 (London: Authority of the House of Lords, 2007).

[20] Ross Anderson, "Why Cryptosystems Fail," *Communications of the ACM* 37 (11) (November 1994). Tyler Moore, Richard Clayton, and Ross Anderson, "The Economics of Online Crime," *Journal of Economic Perspectives* 23 (3) (2009): 3–20; Bob Blakley, "The Measure of Information Security is Dollars," *First Workshop on Economics and Information Security*, Berkeley, California, May 16–17, 2002.

[21] W Yurcik, Cyberinsurance: A Market Solution to Internet Security Market Failure, *Workshop on the Economics of Information Security*, 2002 May 16-17, Berkeley, CA.

[22] Rafael Yahalom, Liability Transfers in Network Exchanges, *Workshop on the Economics of Information Security*, 2002 May 16-17, Berkeley, CA.

[23] Warren, S. and Brandeis, L. 1890. The right to privacy. *Harvard Law Review*. 4: 193-220.

[24] Kirrily "Skud" Robert , "Preliminary results of my survey of suspended Google+ accounts", InfoTropism http://infotrope.net/2011/07/25/preliminary-results-of-my-survey-of-suspended-google-accounts/, Jul 25, 2011 • 2:18 pm

[25] Jens Riegelsberger and Martina Angela Sasse, "Trustbuilders and Trustbusters," in *Towards the E-Society: E-Commerce, E-Business, and E-Government,* ed. Beat Schmid, Katarina Stanoevska-Slabeva, and Volker Tschammer (Boston: Kluwer, 2001).

[26] Russell Hardin, *Trust and Trustworthiness* (New York: Russell Sage, 2002).

[27] Lawrence Lessig, *Code and Other Laws of Cyberspace* (New York: Basic Books, 1999)

[28] Alessandro Acquisti and Ralph Gross, "Imagined Communities: Awareness, Information Sharing, and Privacy on the Facebook," in *Privacy Enhancing Technologies*, Volume 4258 of Lecture Notes in Computer Science, ed. George Danezis and Philippe Golle (Berlin; New York: Springer, 2006), 36–58.

[29] Niklas Luhmann, "Trust: A Mechanism for the Reduction of Social Complexity," in *Trust; and Power: Two Works* (Chichester, N.Y.: John Wiley and Sons, 1979), 1–103

[30] L. Jean Camp, "Mental Models of Security," *IEEE Technology and Society Magazine* 28 (3) 2009.