CRIMINAL LAW IN CYBERSPACE[1]

Neal Kumar Katyal

The new millennium brings new crimes. Witness two of the most talked-about crimes of last year, the ILoveYou computer worm (in terms of economic damage, perhaps the most devastating crime in history, causing more than $11 billion in losses) and the denial-of-service attacks on Yahoo!, eBay, E*Trade, and other sites (which caused $1.2 billion in damage). These events suggest that a new breed of crime has emerged over the past decade: cybercrime. This umbrella term covers all sorts of crimes committed with computers—from viruses to Trojan horses; from hacking into private e-mail to undermining defense and intelligence systems; from electronic thefts of bank accounts to disrupting web sites. Law has not necessarily caught up with these crimes, as the recent dismissal of charges against the author of the ILoveYou worm demonstrates. How should the law think about computer crime?

Some academics see cyberspace as a new area in which first principles of law need to be rethought. David Johnson and David Post, for example, contend that existing legal rules are not suitable for the digital age and that governments should not necessarily impose legal order on the internet. Others, in contrast, believe that a computer is merely an instrument and that crime in cyberspace should be regulated the same way as criminal acts in realspace. The recent U.S. Department of Justice ("DOJ") report on cybercrime typifies this approach. I contend that neither view is correct and that each camp slights important features that make cybercrime both different from and similar to traditional crime.

Underlying the "cybercrime is not different" position is a worry about a unique form of geographic substitution. The concern is that disproportionately punishing activity in either realspace or cyberspace will induce criminals to shift their activities to that sphere in which the expected punishment is lower. For example, if the electronic theft of one million dollars warrants five years imprisonment, and the physical theft of one million dollars warrants ten years imprisonment, criminals are likely to opt for the electronic theft. Such analysis is, however, incomplete. [Some] have observed that the expected penalty for criminal activity is not only the sentence in the criminal code, but also a function of the probability that one will get caught. To the extent that cybercrimes are easier to get away with, sentences might be increased to compensate for this lower probability.

In addition to the probability of being caught, another variable overlooked by the "cybercrime is not different" camp is the perpetration cost of engaging in crime. A bank robbery in realspace, for example, consumes tremendous criminal resources. A robber has to hire lookouts and firepower, garner inside knowledge about the bank, and so on. Profits will be split among five, six, or even more people. A computer theft, by contrast, involves fewer resources and may even be accomplished by a single person sitting down at a computer. Because cybercrime requires fewer resource inputs and less investment to cause a given level of harm, the law might approach these crimes differently.

These variations suggest that cyberspace is a unique medium for three reasons. *First*, and most importantly, the use of computers and other equipment is a cheaper means to perpetrate crime. Criminal law must be concerned not only with punishing crime ex post, but with creating ex ante barriers to inexpensive ways of carrying out criminal activity. In this Article, this principle-which is generally applicable in criminal law-will be called "cost deterrence." The idea is that law should strive to channel crime into outlets that are more costly to criminals. Cyberspace presents unique opportunities for criminals to reduce their perpetration costs; the probability of success in inflicting a certain level of harm while holding expenditures constant is greater. Accordingly, the law should develop mechanisms to neutralize these efficiency advantages.

Some neutralization techniques, however, risk punishing utility-producing activities. For example, encryption has the potential to further massive terrorism (which leads many in the law enforcement community to advocate its criminalization) but also has the potential to facilitate greater security in communication and thereby encourage freedom (which leads many others to push for unfettered access to the technology). This is a standard dilemma that the law encounters in the regulation of technology—call it the "dual-use problem." The problem arises when an activity has both positive and negative uses and forbidding the act forfeits the good uses. To help solve the problem, I introduce a conventional tool, the sentencing enhancement, as a mechanism that can selectively target improper uses. Policymakers and academics have given little attention to sentencing enhancements and lack a theory of when they should be used. This Article endeavors to fill that gap, arguing that sentencing enhancements are suited to certain acts that are not inherently harmful to society and whose benefits depend on context. It shows, for example, how enhancements provide a solution to the encryption debate because they can be aimed at encryption's harmful applications.

Second, cybercrime adds additional parties to the traditional perpetrator-victim scenario of crime. In particular, much cybercrime is carried out through the use of Internet Service Providers ("ISPs"), such as America Online. Government should consider imposing responsibilities on such third parties because doing so promotes cost deterrence. Third parties can develop ways to make criminal activity more expensive and may be able to do so in ways that the government

[1] *Note from TK: This excerpted version of Neal Katyal's article omits footnotes and includes minor alterations to assist with readability. Any errors caused by my edits to this article are my own, not Mr. Katyal's. The complete article is available at* 149 U. PA. L. REV. 1003 (2001).

cannot accomplish directly. The same logic sometimes can apply to victims of cybercrime; law can develop mechanisms to encourage optimal victim behavior as well. As part of this discussion, this Article shows how victim self-help depends on changing police behavior and outlines a strategy to make police departments behave more like fire departments (focusing more on warning and prevention and less on chasing suspected perpetrators after they commit crimes).

Two features of cyberspace, however, suggest that these burden-shifting strategies will be difficult to implement. The first, which borrows from the New Economy theory of "network effects," contends that interconnectivity is an important goal that should not be sacrificed lightly. If potential victims and third parties like ISPs are forced to take precautionary measures-from building strong firewalls to forgoing communication with risky computer systems-these measures may diminish the value of the internet. A strong public law enforcement presence is necessary to prevent the net from fragmenting into small regions accessible only to subsets of trusted users with the right passkeys. A second feature that limits burden-shifting arises from the asymmetric incentives between ISPs and their users. Because an ISP derives little utility from providing access to a risky subscriber, a legal regime that places liability on an ISP for the acts of its subscribers will quickly lead the ISP to purge risky ones from its system. ISPs, as private entities, face no constitutional constraints and little public accountability; the results of ISP liability may be unfair and threaten the potential benefits of the net.

Third, and more generally, a host of thorny problems arise because most activities that occur in cyberspace are invisible to third parties-and sometimes even to second parties. In a space where crimes are invisible, strategies that focus on trying to prevent crime by maintaining public order, such as "broken windows" policing, are of limited utility (though some insights can be adapted to cyberspace). Social norms cannot operate as effectively to prevent crime on the net because its users are not necessarily constrained by the values of realspace.

On the other side of the ledger, the danger of overly aggressive law enforcement is multiplied in cyberspace. Each new major cybercrime leads law enforcement to push for changes to the technical infrastructure to create better monitoring and tracing. If these monitoring mechanisms are hidden in private hardware and software, however, some contend that public accountability may be undermined. A similar point can be made about enforcement by police: Because police are invisible on the internet, the potential for entrapment or other forms of police misconduct may be greater. The ultimate effect of this loss of police visibility may be to poison legitimate activity on the net because confidence in communication may be undermined. An internet user will not be sure that he is talking to a friend and not a government interloper seeking evidence of criminal activity. Because the technology of law enforcement is not well understood among the public, citizens will fear the net and its potential advantages will be stymied. Consider the public uproar over a third prominent news item from this year: the discovery that the Federal Bureau of Investigation ("FBI") has a system, with the poorly chosen title of "Carnivore," which allows it to examine private e-mails.

Nevertheless, the differences between crimes that take place in cyberspace and those that occur in realspace should not obscure their similarities. For example, if crime in cyberspace is easier to commit due to technical prowess, then the law needs to consider how to treat offline crimes that harness technical ability. Similarly, if acts in cyberspace portend criminal activity in realspace, then this dangerous complementarity can-if sufficiently strong-justify punishing acts in cyberspace (an example might be electronic stalkers who may graduate to stalking in realspace). This notion undoes the standard idea that criminal punishment should be reserved only for acts that are harmful; the point here is not that a certain act is itself harmful, but that its commission will lead to a harmful act. Preventing the former act is a mechanism the government may use to discourage the commission of the latter.

The problem of cybercrime is really a larger one of how the law deals with new technologies. Sometimes the law treats crimes that employ new technologies as different and deserving of special regulation (such as wire fraud, hijacking of airplanes, and grand theft auto) and at other times it does not (crimes performed with typewriters and most thefts, which carry the same penalty whether accomplished with James Bond-style panache or by a simple break-in). Lurking underneath this differential regulation is a complex symbiotic relationship between technology and law. Computer crime forces us to confront the role and limitations of criminal law, just as criminal law forces us to reconceptualize the role and limitations of technology.

After all, computer crime is not simply constrained by law. Before Bob Ellickson's and Larry Lessig's pathbreaking work, many scholars assumed that law was the primary mechanism for the regulation of conduct. Ellickson and Lessig helped introduce a second constraint, social norms. They showed how such norms can regulate as effectively as law, or even more so. Lessig's recent work has suggested a third form of regulation, system architecture, or code. Rather than relying on social pressure or legal sanctions, Lessig explains how physical and electronic barriers can prevent harmful acts. In realspace, installing lights on street comers can prevent muggings and other forms of street crime, and placing concrete barricades near inner-city highway ramps can prevent suburbanites from quickly driving in and out to purchase drugs. In cyberspace, internet browsers can be configured to prevent repeated password entry attempts for sensitive web sites or could be coded to prevent certain forms of encryption.

This Article suggests the presence of two other constraints, physical harm and monetary cost. The risk of physical harm in committing a crime is a rather obvious constraint, and one that

is lower with computer crime as compared to realspace crime. Monetary costs, by contrast, are not generally thought of by criminal scholars as a deterrent, and this omission is unfortunate. One reason that computer crime is so dangerous is because it is so cheap to perpetrate.

The legal system, I contend, should focus more on perpetration costs. After all, unlike the probabilistic specter of legal sanction, these costs are certain to be incurred by all who commit a crime. In some ways, the legal system's current focus on legal sanctions at the expense of monetary costs is ironic. Criminals tend to be gamblers-willing to speculate on the chance that they will not be caught-and yet the conventional wisdom is to set up a parlor from which to conduct the wager instead of relying on fixed costs that elude speculation and games of chance. Governments use the threat of jail time to deter offenses even though they know that the bulk of offenders discount the threat of long jail sentences because they have many years to live due to their youth. The lack of high perpetration costs is one factor that explains the rise in cybercrime. Indeed, the fact that some forms of crime are cheap to commit weakens the power of social norms; the ease of, for example, copying a CD leads many to think of it as an innocent act.

Monetary costs, in short, may deter a different stratum of the population than law enforcement-those with less money. Suppose, for example, that the majority of hackers are teenagers. Teenagers, with their lower levels of disposable income, might be particularly responsive to strategies that increase the monetary costs of crime. If dangerous software programs such as hackers' tools were more expensive, or if sensitive web sites charged low admission fees, these forms of regulation may deter criminal wrongdoing in a way that conventional law enforcement may not. This strategy also suggests that when sites such as Napster begin to charge fees for their use, those fees might deter more crime than the speculative risk of a legal sanction. Civil forfeiture of computers and equipment and postconviction restrictions on computer use may also increase perpetration costs and thereby prevent recidivism. Criminal law scholars should incorporate monetary costs into their calculations about optimal deterrence, just as they should recognize social norms and architecture. This multifaceted strategy of regulation is particularly important for those crimes whose offenders tend to be heterogeneous.

Put a different way, the emergence of computer crime threatens an implicit calculus that thus far has constrained realspace crime. Computers make it easier for criminals to evade the constraints of social norms (through pseudonymity and removal from the physical site of the crime), legal sanctions (the probability of getting caught may be reduced for similar reasons), and monetary costs (because the resource inputs necessary to cause a given unit of harm are much lower). The standard Beckerian solution to this problem is to increase the legal sanction, but situating cybercrime within these other constraints reveals other solutions. These other strategies might be more effective because it may be difficult to increase the sanction enough to compensate for a very low probability of getting caught.

Some examples of perpetration cost strategies have been given, so the point will be illustrated by final examples of architectural regulation. Government could redress the lowered constraints against crime by enacting regulations that would prevent pseudonymity by any of the following: (1) by regulating the Internet Protocol ("IP") and software manufacturers (increasing the power of social norms as a constraint on crime, as well as increasing the probability of getting caught); (2) by insisting upon mechanisms that ensure electronic tracing of computer signals to locate offenders (increasing the probability of getting caught); or (3) by requiring targets to use software-hardening measures to prevent hackers from interfering with web sites (increasing the perpetration cost of committing these computer crimes). Reasonable people can disagree about the wisdom of each of these solutions; my point is only that because the emergence of computers can reduce all five constraints to crime, our legal solution cannot be blind to these other constraints and focus willy-nilly on the legal sanction.

It is possible, indeed likely, that our blindness to these other constraints is related to the phenomenon discussed earlier, the subtle existence of second and third parties in crime control. After all, it is difficult for the government to increase the monetary cost of crime directly, and it is likewise difficult for government to modify architecture. It can do so at times by fiat, but government shies away from doing so because it is not situated to know which devices are optimal in preventing crime at the cheapest cost. Mistakes made by the government, by mandating the wrong device or strategy, can impose huge deadweight losses. This Article is designed to show how government, by modifying prosecution incentives and altering civil liability and payment rules, can promote cost deterrence and architectural solutions by harnessing second and third parties. These parties enable government to do indirectly what it often has trouble doing directly-change the perpetration cost of crime and modify architecture in ways that prevent criminal acts. . . .

The term "cybercrime" refers to the use of a computer to facilitate or carry out a criminal offense. This can occur in three different ways. First, a computer can be electronically attacked. We may further subdivide this category by distinguishing among acts that involve (1) unauthorized *access* to computer files and programs, (2) unauthorized *disruption* of those files and programs, and (3) *theft* of an electronic identity. An example of the first category is a break-in to Defense Department Computers. An example of the second category is the ILoveYou Worm. The third category, identity theft, occurs when a person's or entity's identity is wrongfully appropriated. A web page may be "page-jacked," for example, so that when you click onto a financial service to read investment news, you receive spurious information instead.

The above crimes involve situations in which a computer is the subject of an attack. A rather different type of computer crime occurs when a computer is used to facilitate or carry out a traditional offense. For example, a computer might be used to distribute child pornography over the internet or it might be used to create a massive number of copies of a popular and copyrighted song. Complicated insurance fraud, large check-kiting operations, and other sophisticated forms of white collar crime rely on computers to run the criminal operation. In these cases, computers make it easier to carry out a crime in realspace. In these circumstances, computers are tools that expedite traditional offenses.

As news reports suggest, cybercrime is becoming an increasingly common form of criminal activity. The numbers are staggering. The number of recorded computer security incidents grew from 6 in 1988 to more than 8000 in 1999. Theft on the internet caused $2 billion in losses in the year 1995, a number that is much higher today. One company has found 100,000 instances of illegal activity on web sites in one and a half years. New viruses are being launched at the rate of ten to fifteen per day and over 2400 currently exist. Last year, there were more than 22,000 confirmed attacks against Department of Defense computers. It is no surprise that the FBI's caseload has skyrocketed as a result of these trends. …

For several years, the dreams of technological promise and the specter of technology-driven disaster have threatened to collide. The net is becoming an engine of personal, professional, and economic growth, but, because of this growth, new dangers loom. The first months of the new millennium aptly demonstrated these dangers; two crimes that imposed some of the largest economic losses from crime in history were launched from a few private computers. Ironically, these attacks took advantage of what all of us like about computers: their speed, efficiency, trustworthiness, and low startup costs. As criminals become more sophisticated about such attacks, the incidence of these crimes will rise and criminals' escapes will multiply. Law must counter this trend by embracing new strategies that harness the legal and nonlegal constraints on crime.

This Article [proposes] four such strategies, although many more are possible. *First*, law must recognize that an unintended byproduct of computers is that they serve as substitutes for conspirators. Because conspirators sometimes provide benefits to law enforcement, by becoming informants or cooperating witnesses, the government must devise strategies that recognize the fact that these benefits are lost when this substitution occurs. One such strategy . . . is to treat computers as quasi-conspirators.

Second, law should recognize that certain technologies, such as encryption and anonymity, have dual purposes. Rather than postulating that they are entirely deleterious and punishing them wholesale, society must understand that these technologies can be used for both good and bad ends. To accomplish this balance, the law should develop sophisticated sentencing enhancements and other nuanced strategies such as specific exclusions, and forgo the blunt sword of total prohibition.

Third, the government must increase the financial cost of crime, and the skills necessary to commit it, by placing some responsibility on third parties, such as ISPs, and even on victims. But the government should also recognize that while victims and ISPs might be cheapest crime avoiders, able to prevent crime more cheaply than other actors, their prevention strategies may carry broad, systemic costs, such as balkanization of the net via systems of passwords and other methods that limit access. Law enforcement must have a strong presence on the net to steer victims and ISPs away from suboptimal self-help strategies; yet at the same time, the police must stress that these entities have a duty to take self-help measures.

Fourth, instead of treating all crime as equal, law enforcement should attempt to inflict disproportionately heavy punishments upon those crimes that create the most visible, or otherwise evident, social disorder in cyberspace. Doing so will avoid complementarity problems, such as copycat crimes or crimes committed because hackers' tools are easily accessible, and will help reassure the public and industry that cyberspace is safe.

These four strategies are calculated to help set up incentives that make crime too expensive to carry out, preserve the benefits of the net, and provide computer users with the assurance that the net is at least as safe as realspace. Yet the strategies do run risks, from trenching on privacy and freedom of speech to poisoning the free flow of ideas. Those risks cannot properly be addressed in this initial Article, but doing so is a requisite component of an effective plan to combat cybercrime.

Although cyberspace has unique particularities, the lessons we have learned are not confined only to the electronic world. A central theme of this Article, for instance, is that a crucial variable for preventing crime is perpetration cost. Law can and should develop strategies to make crimes more expensive. The government currently relies on the speculative risk of imprisonment to deter wrongdoing, but a strategy focused on raising certain costs associated with the wrongdoing itself may be more effective. If the majority of criminals are gamblers, or are at least less risk-averse than others, then the law should focus on raising the fixed, ex ante monetary costs that these criminals will pay to perpetrate a crime, not on merely enhancing probabilities of jail time that criminals will tend to ignore. Deterrence may be better served by increased monetary costs on all lawbreakers rather than by traditional strategies such as raised penalties for the few criminals unlucky enough to be caught.

[There is also] the need for a more nuanced solution to the problem of dual-use activities and [I argue] that sentencing enhancements can preserve the positive uses of a given act while attacking its negative uses. This theory of regulation applies

generally, although it may be particularly useful in the area of cybercrime, the hallmark of which may be a preponderance of dual-use activities. . . .

At issue in this treatment of cybercrime is a view of deterrence that differs substantially from that offered by economists and sociologists, one that is not fully focused on the mind of the offender at the last minute before she commits a crime. My account stresses the way in which legal rules promote deterrence in other ways, such as by encouraging products that prevent crime, building architecture that makes crime more costly to criminals, and harnessing individual conscience and public values in ways that make crime look less attractive. By manipulating variables besides legal sanctions, crime may be prevented even when criminals are not that responsive to legal sanctions.

Both realspace and cyberspace are rapidly evolving, and the way criminal law approaches these spheres today may soon be anachronistic. Still, while the approaches may need to be updated over time, the fundamental building blocks of successful anticrime strategies will remain constant. Law must strive to prevent great harm at cheap cost, and it must define costs broadly enough to include all of the negative effects of crime prevention (substitution effects, the social costs of suboptimal self-help strategies, and so on). Our system of criminal law should attempt to raise the perpetration costs of engaging in crime and should also provide enough enforcement to create the conditions under which trust flourishes and networks develop. At the same time, the government must avoid creating disincentives to utility-producing activities and must strive to surgically target harmful acts. These building blocks of criminal law apply to the brick-and-mortar world, as they do to cyberspace.