Would indirect liability reduce costly cyberspace externalities?

Holding Internet Service Providers Accountable

By Douglas Lichtman
University of Chicago

Internet service providers (ISPs) are largely immune from legal liability for the various forms of online malfeasance to which they contribute. America Online, for example, paid Matt Drudge $3,000 a month to write an online gossip column; but when Drudge used the column to accuse Clinton appointee Sidney Blumenthal of spousal abuse, a federal judge ruled that AOL bore no responsibility for the smear. Similarly, Verizon Communications today counts among its subscribers an untold number of peer-to-peer pirates, yet the firm faces no financial liability for copyright infringement online and, indeed, does almost nothing to help copyright holders defend their work.

This is surprising. After all, while it surely would be unwise to punish ISPs for every bad act committed by their subscribers and it would be equally foolish to force service providers to play policeman in instances where the costs of doing so would overwhelm any plausible benefits, legal liability can take more modest forms. With respect to copyright infringement, for instance, why not require ISPs to deliver warnings to accused subscribers? Infringers are anonymous Internet Protocol (IP) addresses to copyright holders, so copyright holders have a hard time delivering warnings themselves. But an ISP can easily match an accused IP address to a real-world name and billing address, and thus an ISP can easily deliver a warning that would remind an accused subscriber that piracy is illegal and that "the complaining copyright holder might take his evidence to court, where we will be forced to reveal your identity and provide further evidence of your alleged bad acts." Imagine the shiver that would go down an infringer's spine upon finding that note in his mailbox, complete with a specific accusation that he downloaded Madonna last Tuesday at midnight from his bedroom computer.

The copyright and defamation immunities to which I allude have been in place for years and would be difficult to displace. I therefore want to focus instead on what is shaping up to be the next immunity battle: the recent push to immunize Internet service providers for their role in the propagation of worms, viruses, and other forms of malicious computer code. Drawing analogies to copyright and defamation, courts have in recent years read the relevant statutes and interpreted common law principles such that ISPs are today almost entirely unaccountable for issues of cybersecurity. But, as I will argue here, immunity in this instance is hard to defend on policy grounds and it is sharply inconsistent with the conventional logic of indirect liability.

YOUR BROTHER'S KEEPER?

Indirect liability is said to attach in instances where the law holds one party liable for a wrong committed by another. A familiar setting is the employment relationship, where an employer can be held liable for torts committed on the job by his employees. But other examples abound. Bars are sometimes held liable when bartenders serve alcoholic beverages to patrons who later harm others while driving under the influence. A motor vehicle owner can be held to account if a driver to whom he loans

Douglas Lichtman is a professor of law at the University of Chicago. He may be contacted by e-mail at dgl@uchicago.edu.

This article is drawn from a paper co-authored with Eric Posner entitled "Holding Internet Service Providers Accountable" that is forthcoming in the *Supreme Court Economic Review.*

his car ends up causing an accident. Landlords are sometimes on the hook if they take inadequate precautions against criminal activity that harms tenants. Even product liability law has this same basic structure: A buyer might use a dangerous product in a negligent manner and cause injury to a third party; if the victim can show that the accident would not have occurred had the manufacturer employed better product design, the victim might be able to recover from the manufacturer instead of (or in addition to) the buyer.

LIABILITY BY CONTRACT Conventional economic analysis suggests that an explicit rule imposing indirect liability is not necessary when two conditions are simultaneously met: first, the relevant direct actors are subject to the effective reach of the law, which is to say that the employees, drivers, and criminals discussed in the previous examples are easy to identify and have assets that are sufficient to pay for any harm caused; and, second, transaction costs are such that the direct actors can use contract law to shift responsibility to any party that might otherwise be an attractive target for indirect liability. The intuition is that, when those conditions are satisfied, the various parties can create indirect liability by contract and, albeit subject to some minor constraints, they will do so where that would be efficient.

To see this, consider the employment setting in more detail. If the driver of a delivery van is himself easy to identify and, further, the driver has adequate resources to pay for whatever harm he might cause in the event of an accident, then there is no strong argument for imposing liability on his associated retailer. No matter what the legal rule, the driver and the retailer will efficiently allocate liability through the employment contract. Thus, if the optimal rule would impose on the retailer the obligation to inspect every delivery van each morning or to test employees randomly for drug and alcohol abuse, the driver and retailer will agree by contract to those monitoring activities. Similarly, to the extent that driving the truck poses an unavoidable risk of injury to others, the driver will either shift that risk to the employer through an indemnity clause or assume that risk himself and demand higher wages in compensation. The legal rule in this situation is just a default; where transaction costs are low and employees have adequate resources, contracts allow private parties to shift and divide legal responsibility efficiently.

LIABILITY BY LEGAL RULE Things change when either of the conditions identified above fails to hold. For instance, where contracts are easily negotiated between the driver and the retailer but the driver lacks resources, the absence of indirect liability would tempt the retailer to leave tort liability on the shoulders of the driver, in essence using the driver's financial limitations as a cap on legal liability. Similarly, in a situation where contracts are possible but a negligent employee's identity cannot be ascertained — for example, witnesses report that a firm's van hit a pedestrian but no one saw the driver — again the absence of indirect liability would limit victim recovery, putting the driver and retailer in a position where, taken together, they would not take adequate care. Where the driver has adequate resources but the parties cannot contract effectively, the legal rule clearly matters as well, this time because the inability to contract would make it impossible for the parties to shift responsibility as needed.

The interesting cases are, therefore, those where either the relevant bad actors are beyond the reach of the law or transaction costs make reallocation by contract implausible. For those cases, economic analysis identifies two additional considerations: first, indirect liability might be attractive where one party is in a good position to detect or deter another's bad act; and, second, indirect liability might be attractive where liability would serve to encourage a party to internalize some significant negative externality unavoidably associated with its activities.

Start with the first consideration, that indirect liability might be attractive where the potentially liable party is in a good position to detect and deter bad acts. That is, for example, one of the main reasons why employers are responsible for torts committed by their employees. An employer can control his employees. He can monitor their behavior, screen them before entrusting them with dangerous equipment, develop compensation schemes that encourage them to exercise due care, and otherwise beneficially influence their on-the-job decisions. The prospect of indirect liability pressures employers to make use of those mechanisms and, in that way, to minimize the expected cost of accidents.

Turn now to the second consideration. Even where a retailer can do nothing more to ensure that the drivers of its delivery vans take appropriate care, it is likely efficient to have the retailer pay at least some fraction of the costs of any delivery accidents. The reason is that this forces the retailer to account for accidents when deciding the price and frequency of deliveries. If accidents are unavoidable, price will rise and quantity will fall, which is exactly what should happen given this unavoidable harm. This is referred to in the literature as an effect on "activity level," which emphasizes that

Economic analysis suggests that indirect liability might be attractive when one party is in a good position to detect or deter another's bad conduct.

the purpose of liability here is not to encourage precautions but instead to influence how often the harmful activity in question takes place.

These factors — call them "control" and "activity level" — help to identify cases where indirect liability might be attractive. The actual question of whether liability should be imposed, however, typically turns on other, often setting-specific, considerations. Thus, while the telephone company surely has the ability to deter crank phone calls by more carefully monitoring calling patterns, it is unlikely that indirect liability would be attractive, both because of obvious privacy concerns and because of worries that, in its attempts to address the problem of crank calls, the telephone company would inadvertently interfere with substantial legitimate telephone activity. To reject indirect liability in this situation is to announce that the costs of crank telephone calls are not sufficiently high compared to the costs of indirect prevention. Similarly, the mere fact that an airport provides a venue from which airlines generate pollution and noise does not itself justify imposing liability for that harm. The reason is that private parties who own property near the airport themselves make decisions that increase and decrease the importance of airport externalities; in a world where the airport absorbed the costs in full, neighbors might inefficiently decide to use their properties to raise livestock or care for the elderly, two uses so sensitive to noise and pollution that they likely should be disfavored given the proximity of the airport.

That said, the control and activity level factors do helpfully sketch the contours of efficient indirect liability rules. For instance, these factors make clear that employers should not typically be held accountable for torts committed by employees acting outside the scope of employment. The employer has no special advantage when it comes to stopping its employees from abusing their spouses or picking fights at bars. Moreover, neither activity is rightly understood as a consequence of the employer engaging in its core business. Whether the employer is in its current line of business or another, the employee is probably just as likely to commit those bad acts. Thus, except in exceptional circumstances, neither the control nor the activity level rationale fits, and liability for torts committed outside the scope of employment is therefore inappropriate.

INDIRECT LIABILITY APPLIED TO ISPS

The conventional economic account makes clear that private parties cannot create the optimal liability regime on their own in instances where the party directly responsible for the bad act is beyond the effective reach of the law, or in instances where transaction costs make contract negotiations implausible. The conventional account further stresses that liability should be considered in instances where one party has the ability to deter or detect the bad acts of another, and also where liability can serve to encourage a party to internalize some significant negative externality associated with its activities. As I will argue here, violations of cybersecurity take place in a setting where most or all of those conditions seem likely to hold.

BEYOND THE LAW'S REACH Individuals who originate malicious computer code are typically far beyond the reach of conventional law. For one thing, they are hard to identify. Sophisticated saboteurs use the Internet's topology to conceal their tracks by routing messages through a convoluted path that is difficult for authorities to uncover. Moreover, by the time a computer virus or worm is detected, the trail often is cold. Internet pests like worms and viruses are routinely programmed to sit idle for a period of time before triggering. That allows mischief-makers to time their attacks to coincide with important world moments — the start of the new millennium, for example — and also affords the troublemakers time to disappear. The fact that many hackers reside overseas only exacerbates the problem, introducing issues of jurisdiction and the need for international cooperation.

Indirect liability might also be attractive if it encourages a party to internalize some significant negative externality associated with its activities.

Even if caught, individuals who create malicious computer code rarely have sufficient assets to pay for the losses they impose. Prominent Internet worms and viruses impose billions of dollars worth of damage. Obviously, hackers will rarely have resources sufficient to pay up. Criminal liability could, in theory, substitute as a deterrent; however, where the risk of apprehension is sufficiently small and the magnitude of the loss is sufficiently large, criminal punishments often cannot be made high enough to deter adequately. Juries, after all, are reluctant to impose too large a sentence for non-violent crime and, besides, long-term incarceration is expensive to the state.

Interestingly, concerns about bad actors being beyond the reach of the law do not apply to the individuals and entities who, instead of creating an Internet pest, inadvertently propagate one. An example might be a firm whose server is run in such a way that an outside party can easily take it over, or an unsophisticated user who installs a malicious program when prompted to do so by an anonymous e-mail solicitation. There is no reason to believe that careless firms and users lack the resources necessary to pay for whatever share of the harm they cause; moreover, neither would likely be that hard to track down. Computer users who fail to exercise appro-

priate caution when opening e-mail attachments are hardly sophisticated enough to cover their tracks in the event of a problem. The only sense in which those bad actors are beyond the reach of law is the practical concern about the costs of identifying and suing them as compared to the fraction of the damages for which they might be held legally responsible. Beyond that, parties who propagate but do not create malicious code are not beyond the reach of the law; although, as will become clear below, there are other reasons why indirect liability might be warranted even in those sorts of cases.

CONTRACTS AND TRANSACTION COSTS A second consideration identified in the baseline analysis was the ability of the relevant parties to write contracts. Applied here, the point is that ISPs in theory can use contract law to create for themselves system-wide liability; each would agree to be liable to the others for any harm caused by its subscribers. So why are those obligations not in place, and why should the law respond by imposing them?

An intuitive answer is that there are so many ISPs in operation that the transaction costs of negotiating the necessary web of contracts would be prohibitive. But that explanation is only marginally satisfying, in that ISPs are already all part of a complicated and fully inclusive network of contracts, specifically the "peering" and "transit" agreements under which the various private owners of the Internet backbone agree to carry traffic one to another. A more satisfying explanation is that any network of contracts focusing on issues of cybersecurity would be perpetually out of date, and updating such a complicated web of interdependent security obligations would be all but impossible given the number of parties involved and the complicated questions any update would raise regarding the appropriate adjustments to the flow of payments.

Still, there are puzzles lurking. Microsoft has licensing agreements with a huge percentage of home computer users, and thus the firm seems to be in the perfect position to ensure that users take sensible precautions like updating their virus software and downloading system patches. Microsoft could even make those obligations self-executing by blocking Internet access for any computer whose software is (say) more than 10 days out of date. That would be a minimally intrusive way to ensure that users keep their precautions current, a bit like mandatory vaccinations for school children. Instead, Microsoft merely offers updates to its customers and allows each customer to decide whether the private benefits of a given update warrant the private costs in terms of time and inconvenience. The result might very well be a classic case of externalities leading to suboptimal behavior: Microsoft's customers as a group would be better off were each to update regularly but, without coordination, each customer opts to update less frequently. This suggests that there must be a bigger problem with contractual solutions — public relations? privacy concerns? security? — although in truth the explanation might be that Microsoft is at the moment in too precarious a position vis-à-vis worldwide antitrust authorities to do anything that might be perceived as the use of its market power to foist additional software on unwilling consumers.

DETECTION AND DETERRENCE Indirect liability is primarily attractive in cases where the indirectly liable party can detect, deter, or otherwise influence the bad acts in question. ISPs seem to be a natural choice under this criterion. Consider, for example, an ISP through which a troublemaking user obtains access to the Internet. Such an ISP can detect criminal behavior by analyzing patterns of use, much as a bank can detect credit card theft by monitoring each customer's pattern of purchases. Easiest to catch would be patterns that are intrinsically suspicious, such as a continuous stream of communications from a home user or the repeated appearance of identical computer code attached to a large number of outgoing e-mail messages. But an ISP could also detect patterns that are suspicious because they represent a radical departure from the user's ordinary behavior. The ISP would need only maintain a profile that captures in broad strokes each subscriber's rough practices, and then evaluate recent activity against the historical backdrop. Again, credit card companies actually do this, and ISPs could do it too.

Another option might be to record a subscriber's data stream and store that information, ideally in encrypted form, for a period of time. Many offenders could be traced if ISPs were to record traffic in this manner. But ISPs do not routinely record traffic today, both because of privacy worries and because of the enormous volume of communications. Legal rules, however, could ease those concerns. For instance, the law could require that ISPs store the information securely and release it only to law enforcement officials, thus lessening the worry that stored information would leak out by accident or be used for impermissible purposes. The law could also require that ISPs record information about the data communication — size, duration, timing, and so on — but not its substance, thus protecting privacy and reducing volume. The law could even require ISPs to record information only when particular triggers raise suspicion, or perhaps only in response to specific government requests.

I could go on for some time with ideas along these lines. The goal for now, however, is not to describe the precise precautions that ISPs should or will take in response to liability, but instead to simply make clear that ISPs are in a good position to influence the number and severity of cyber-attacks. Indirect liability would pressure them to take that role seriously, thereby encouraging the people who have the proper technical expertise —not me — to first identify and then implement whatever turn out to be the most effective precautions.

ACTIVITY LEVELS In theory, indirect liability can be attractive independent of its role in encouraging detection and deterrence because it encourages the responsible party to account for negative externalities unavoidably associated with the relevant product or service. In practice, however, I doubt that I would favor ISP liability on this argument alone. My hesitation does not derive from any doubts over whether ISPs impose negative externalities as they enroll new customers and offer new services. Of course they do, given that any new subscriber can turn out to be a careless user, and any

new service can quickly devolve into a portal for Internet contagion. My hesitation instead derives from the fact that there are drawbacks to imposing liability solely because of negative externalities, and those drawbacks are significant in this particular application.

One drawback associated with the activity level rationale is that it might distort behavior by forcing parties to internalize negative externalities while they ignore equally important positive ones. As applied here, the negative externality is the aforementioned concern that each new subscriber could materially reduce cybersecurity by engaging in unsafe practices or intentionally introducing an Internet pest. The comparable positive externality is that each subscriber can just as plausibly turn out to be a homemaker who makes significant purchases online or a college student who contributes to the development of open source software. Liability that encourages ISPs to take precautions is one thing; but a legal rule that relentlessly brings home negative externalities while completely failing to account for positive externalities has no claim at creating optimal incentives.

A second drawback to the activity level rationale is the concern that imposing liability on one party almost inevitably discourages another party from taking adequate precautions. Applied here, the worry is that imposing liability on ISPs might inefficiently reduce subscriber incentives to install virus protection software and maintain adequate backups. That is a concern associated with indirect liability no matter what the rationale; but the concern resonates with particular force in cases where indirect liability is being used solely as a means by which to influence the liable party's activity level. The reason: these are cases where (by assumption) the liable party cannot take additional cost-justified precautions, so reductions in the level of care taken by other parties warrant considerable weight.

A third argument against imposing strict liability solely because of activity level concerns is that activity levels in this setting are already sufficiently suppressed. Worms, viruses, and the like reduce the allure of Internet access and thus discourage Internet use no matter what the liability rule. This is a natural reduction in activity levels and, while there is no reason to believe that it leads to efficient levels of activity, the existence of this natural disincentive does combine with the concerns discussed above to make any additional reduction seem not only less important, but also more difficult to calibrate.

All that said, activity level concerns can be important, and hence I harbor some uncertainty over where to draw the line. Consider Microsoft again. Even if the software giant cannot take additional precautions against Internet contagion, the price increase that would likely result from an increase in liability would itself have social benefits in that the resulting price would better reflect the relative value of the Windows operating system as compared to alternatives like Apple Computer's operating system, Mac OS. Many computer enthusiasts believe that Mac OS is more stable and secure than Windows. If so, that benefit is not today adequately captured in the products' relative prices. By increasing liability and hence disproportionately increasing the price of Windows software, however, an indirect liability rule would help to solve the problem, ultimately driving business toward the more secure and efficient alternative. More generally, in situations where several competing products are each capable of generating a comparable positive externality, it might be attractive to use indirect liability as a way of pressuring firms to select prices that accurately reflect each product's unique negative externalities.

Any indirect liability regime needs to be created by law, rather than by contract, because transaction costs are a serious obstacle to contractual solutions.

OBJECTIONS

My argument thus far is that indirect liability is attractive primarily because ISPs are in a good position to deter the various acts associated with cyber-insecurity, and perhaps secondarily because liability would force ISPs to internalize some of the negative externalities they impose. Further, I have argued that any indirect liability regime needs to be created by law, rather than by contract, because many of the relevant direct bad actors are beyond the reach of law and because transaction costs are a serious obstacle to contractual solutions in any event.

Consider now the two primary objections to this analysis: first, that liability will cause ISPs to overreact and thus exclude subscribers who should be online; and second, that liability will inefficiently interfere with subscriber efforts at self-help.

OVERZEALOUS PROVIDERS The most common objection to ISP liability is that it would cause ISPs to raise prices, and — while those higher prices might better represent the real costs of Internet access — the higher prices would also drive marginal subscribers out of the market. That end result is inefficient, according to this argument, because advertisers, merchants, friends, and various other Internet entities might in the aggregate prefer that the marginal customers remain. The problem is thus just an externality: a mismatch between the private incentive to subscribe to Internet service and the social benefits made possible by each new subscription.

My first response is that this concern, while plausible, seems overdrawn. Many of what at first sound like externalities turn out to be influences that are already accounted for in a subscriber's decision of whether to subscribe. For instance, I certainly benefit from the fact that my mother is regularly online and hence available for easy e-mail correspondence, but that is not an externality because my mother and I have a rich relationship through which I can indicate to her how much I value her presence and, if necessary, contribute in cash or kind toward the monthly cost of her subscription. So, too, the online bookseller Amazon.com benefits from Mom's Internet access, but Amazon also has ways of helping her to internalize that effect, for instance by rewarding her with free shipping on her purchases. This is obviously not to say that all externalities are internalized, but only to suggest that the problem is not as stark as it might at first seem, and not all that different from a million other markets where incidental positive externalities slip through the decision-making cracks.

Second, even if there are non-trivial positive externalities at play, note that it would be counterproductive to respond to the problem by reducing ISP liability from its otherwise optimal level. Restaurants, for example, create positive externalities by drawing crowds that in turn patronize neighboring businesses and stimulate the local economy. Yet no one suggests that, in response, local authorities should stop enforcing the health code; that response would just drive customers away. Similarly, inventors produce devices that stimulate further innovation; society in return rewards them by granting valuable property rights called patents. Does anyone really believe that society should instead shield inventors from liability if their inventions cause harm? In short, positive externalities are not typically compensated by legal immunity, because even an entity that produces positive externalities should still take due care while engaged in its beneficial activities. There is nothing special in this respect about the Internet. There are many mechanisms that might sensibly be used to encourage ISPs to create positive externalities — tax breaks, infrastructure subsidies, and so on — but immunizing ISPs from indirect liability is unlikely to be one of them.

SUBSCRIBER SELF-HELP A second objection to ISP liability is that it will reduce subscriber incentives to buy antivirus software, install firewalls, and similarly engage in prudent self-help. That undoubtedly is true. However, the logical implication is not complete immunity for ISPs. Instead, liability should be tailored in light of this possibility, the goal being to encourage service providers to adopt the precautions that they can provide most efficiently while leaving any remaining precautions to subscribers and other market actors. This is a standard scenario. Pedestrians can exercise care in crossing the street. They can also stay at home rather than venturing near the roads, and they can wear unfashionably bright attire so as to increase the odds of being seen at night or during inclement weather. Yet no one suggests that, because pedestrians can engage in their own forms of precaution, automobile drivers should be immune from tort liability.

The same intuitions apply here. The fact that multiple parties can take precautions against malicious computer code might argue for some form of a balanced liability regime that leaves both subscribers and ISPs with some incentive to take care, but that fact does not in any way argue for complete immunity for ISPs. There are precautions in which ISPs can and should engage, and shifting the full costs of accidents to Internet subscribers would inefficiently reduce each ISP's incentive to do so.

CONCLUSION

I began with reference to the immunities that ISPs enjoy with respect to defamation and copyright infringement. Let me briefly conclude by identifying some differences between liability in those instances and the liability at issue here.

In the context of defamation, it is important to remember that judgments are unavoidably subjective and fact-specific, and thus it might be unreasonable to ask an ISP to identify defamation on its own. A malicious computer program, virus, or worm, by contrast, can be more readily and less intrusively identified. Besides, the social costs of a system where a few innocent programs are accidentally delayed by an overly cautious ISP seem much less onerous than the social costs associated with an equivalently imperfect filter that might delay socially important free speech.

As for copyright infringement, meanwhile, note that, while the possibility of worms and viruses reduces the average subscriber's interest in Internet service, the possibility of copyright infringement likely increases it. Indeed, in many ways, music piracy is the "killer app" that is today driving the deployment of broadband Internet service. Infringement therefore has a silver lining — it is a camouflaged subsidy to broadband — whereas malicious computer code has none. This might mean that policymakers ought to be more interested in imposing liability for cyber-insecurity than they are in imposing liability for music piracy. Or it might mean the opposite, as ISPs already have a strong incentive to improve cybersecurity (subscribers favor it) whereas ISPs face no similar incentive when it comes to fighting copyright infringement.

All that said, my own view is that ISPs should not be immune from liability in any of these three settings. AOL surely should have been called to account for the Drudge Report, just as the *New York Times* would have been had one of its columnists printed that very same slur. And Verizon should be required to take steps against online piracy, primarily because the firm can enforce copyright law at low cost and with high efficacy. But, as compared to those two, liability for malicious computer code represents the strongest case. There is room to disagree over the details of legal liability — whether it should sound in negligence or strict liability, whether it is best implemented by statute or via gradual common law development, and so on — but it is hard to understand how complete immunity could possibly be the right answer. R