

The Law of Copyright Accidents

Oren Bracha

“The question is commonly thought of as one in which A inflicts harm on B and what has to be decided is: how should we restrain A? But this is wrong. We are dealing with a problem of a reciprocal nature. To avoid the harm to B would inflict harm on A. The real question that has to be decided is: should A be allowed to harm B or should B be allowed to harm A? The problem is to avoid the more serious harm.”

Ronald H. Coase, The Problem of Social Cost, 3 J. Law & Econ. 1, 2 (1960).

Assume you accidentally killed someone. As a rule, you would only incur civil liability for the harm you caused if you were negligent. Now assume that you accidentally infringed someone’s copyright (if you don’t understand how you can “accidentally” infringe copyright then—trust me—you can; and keep reading). You would incur civil liability for your infringement irrespective of your negligence, intent or even knowledge. Why the difference? Today the most common answer to this question is a just-so story. Copyright infringement, it goes, is by nature not dependent on the presence of fault, as if copyright has a nature or is based on some unquestionable truisms. Court opinions and commentaries from the time that this rule appeared in copyright law around the mid-nineteenth century occasionally offer two reasons for it. First, these sources observe that copyright is a property right and transgressions on property rights are actionable with no regard to the fault of the transgressor. Standing alone, this reason is what Felix Cohen called “transcendental nonsense” and (hopefully) not many would find it convincing today. Second and relatedly, it was sometimes asserted that in copyright the preemptive consideration is harm to the copyright owner and that such harm will always outweigh any countervailing considerations of benefits to the alleged infringer or to third parties. Again, few would find this reasoning convincing today, even among the minority (in the U.S. anyway) of those who subscribe to natural rights theories of copyright. The closest thing to a policy argument in favor of faultless copyright one can find is the observation in a handful of modern cases that the infringer is in a better position to know of the infringement and avoid it and therefore he should always bear the cost of the harm created by it. While having a grain of truth in it, this observation is doubly wrong. Not in all cases the copier is better positioned to avoid the infringement. And what’s more, even when in between the copier and the owner the

former is better situated to avoid the infringement, in many cases the optimal avoidance strategy is some combination of preventive measures by both the copier and the owner.

So we are left with the question: is the current rule of faultless copyright justified and if not what should be its alternative? To make things less abstract, let's restart with a specific example. Jack, a fairly known visual artist, buys a postcard displaying an old black and white photograph of his hometown. The postcard bears neither copyright notice nor any metadata about the photograph. Jack scans the photograph and digitally incorporates it together with other imagery and art work into a visual collage he calls "memories of my hometown." Jack duplicates the collage in a poster form and sells 10 copies for a considerable price at his daughter's elementary school fundraiser. At that point Jill materializes, explains that she owns the copyright in the photograph, that Jack infringed her copyright and that she suffered substantial harm as a result. This is a simple scenario of a copyright accident. Why? Because at the time of infringement there was no certainty of infringement and its resultant harm. A reasonable person in Jack's shoes would have suspected a probability of infringement. But probability is not certainty. Given the lack of information *ex ante*, for any number of reasons Jack's use could have been non-infringing: because the copyright term in the photograph had expired; because the photograph had fallen into the public domain due to another reason (e.g. failing to satisfy the strict formalities of the pre-1976 regime); or because some owners would permit this kind of use. When Jack engaged in his activity he knew (or should have known) of a substantial risk of infringing, but it was only a risk as opposed to a certainty.

Even conceding that it was an accident, couldn't Jack prevent it or at least reduce the risk? Of course he could. Jack could have invested in investigating the provenance and status of the photograph or in locating its copyright owner. But given the fact that there was neither notice affixed to the photograph nor registration with the copyright office this would have been a costly endeavor. Alternatively Jack could have refrained from his use altogether, incurring the opportunity cost. A crucial point is that Jill too could have prevented the accident or at least reduced the risk of its occurrence. She could have affixed notice to any authorized copy of the photograph, added additional metadata to the copy, registered the work with the copyright office and policed the market for availability of unauthorized copies. Thinking on the situation as an accident, the relevant question is not whether Jack could have prevented the accident or even whether he was better situated to prevent it compared to Jill. The relevant question is whether the

legal rules create the correct incentives to *both* Jack and Jill to invest optimally in reducing the risk of the accident.

The above is a stylized example. Lest you think that cases of copyright accidents are rare or trivial, let me hasten to say that there are many real world examples that are neither. Consider the following:

- A documentary film maker working on a tight budget wants to combine in her film a historical photograph of considerable relevance to the film's subject. The photograph bears no notice and its exact provenance is unknown.
- The Library of Congress wants to expand the American Memory project (<http://memory.loc.gov/ammem/index.html>) to include a digital collection of twentieth century advertisement prints available online to the public. The copyright status and ownership of many of the prints is unknown and costly to uncover. The Library considers using the prints subject to allowing copyright owners to "opt-out" by providing notice of their objections.
- A computer programmer incorporates in his own software an open code computer program offered to the public online via a General Public License (GPL). Many have contributed to this program and it's hard to establish whether at some point unlicensed "proprietary code" under copyright had "tainted" the program.

Perhaps in some other contexts we can assume that cases of accidents are rare and/or that the cost of avoiding them to potential injurers is trivial. In such contexts we should be content to optimize the scope of protected entitlements and then fully protect them even in cases when *ex ante* only a risk rather than a certainty of violation is present. If the cost of avoiding infringement accidents is truly trivial let potential injurers invest in prevention or (what amounts to the same thing) contract with entitlement owners over their transfer or, when feasible, even over assigning the risk of their violation. For several reasons, however, this is not the case with copyright:¹

- a. *No clear boundaries.* Since expressive works have no clear, easy to ascertain boundaries, it's often costly to know whether a particular action constitutes infringement or even copying. I do not mean here "legal accidents" where it is hard to know whether a known

¹ The following explication is based on Assaf Jacob and Avihay Dorfman, *Copyright as Tort*, 12 Theoretical Inquiries in Law 59 (2010).

copying is legally permissible or not (for example because it's fair use, or just copying of "ideas"). How to handle legal accidents against the backdrop of fuzzy law is a thorny issue that I want to bracket here. Even putting legal accidents aside, however, it remains the case that sometimes the very fact that boundaries have been transgressed is unknown or costly to discover. Think of the computer programmer above who uses a computer program potentially tainted by proprietary code. Or consider George Harrison who was found by a federal court to have copied another song "unconsciously."

- b. No self-evident legal status of the work (most importantly, whether it is in the public domain or not).
- c. No self-evident signal of the owner's preference in cases where it is plausible to assume that some owners would actually permit and even welcome the use. Consider full digital reproduction of old academic articles.
- d. Hard to ascertain identity of owners (for purposes of ascertaining other relevant information and concluding a transaction where necessary).

Whether or not in other cases (such as trespass to land) we can assume that accidents will be rare and that prevention cost by injurers will typically be trivial, this is not the case with copyright. This conclusion does not tell us yet whether an element of fault is justified in copyright accidents cases and how exactly the relevant rules should be crafted. But it does tell us that these questions should be seriously considered and it suggests that the familiar tort theoretical framework of the law of accidents could be applied to answer them.

Strict liability vs. Negligence

Consider the copyright accident from an ex ante perspective. There is an action by the copier (or the potential "injurer") that creates a risk of harm to the copyright owner (the "victim"). The expected private cost of this risk is the harm suffered by the owner from the infringement discounted by the probability of its occurrence. This is however not the expected public cost. The expected public cost is, rather, the extent to which the owner's expected private cost, suffered just prior to the accident (t_1) will reduce his incentive at the time of creating the work (t_0). It's this public cost that copyright policy cares about. Since this public cost will be proportionate to the private cost, however, I'll be referring here simply to the "cost of the accident" or even "cost to the owner." Finally we have preventive measures available to both

parties at various costs. This includes the preventive measure available to the copier of forgoing the risky action altogether, thereby suffering the opportunity cost.

Assume for now that the relevant normative criterion is efficiency and that the only available rules are strict liability or simple negligence. What should be the rule applicable to copyright accidents? The standard answer is that in unilateral cases---meaning, when reasonable preventive measures are only available to the injurer---the correct rule is strict liability. Under strict liability the injurer internalizes the full cost and benefits of his actions and therefore he will have optimal incentive to take just the right amount of preventive measures (for now we assume that the right amount is represented by the Learned Hand formula: $B < PL$). As explained above, however, in a significant subset of the cases it does not hold that the copier is the only party who can take preventive measures. In many cases of copyright accidents both the copier and the owner can take some preventive measures and often the optimal preventive strategy will be some combination of the two. In such bilateral situations strict liability is deficient because the owner has no incentive to invest in prevention. Moreover, under this rule when the copier faces alternative preventive strategies and the more efficient one requires cooperation from the owner, knowing he can expect no such cooperation he is likely to choose the less efficient one. A negligence standard, by contrast, incentivizes both the copier and the owner to optimally invest in prevention. The copier will bear the costs and benefits of his preventive measures up to the optimal point and therefore will implement them. The owner, who is not protected from the cost of the accident when the copier takes appropriate precautions, will bear the cost and benefits of any additional preventive measures taken by him and therefore will undertake them to the extent they are optimal.

Externalities. Intellectual works often involve externalities. The copier's use of the work may benefit many people thereby creating much value that he does not internalize, as in the case of the large digitization project by a digital library made available to others on relatively open terms. The copyrighted work too may sometimes involve positive externalities (perhaps less often and less significant). How do the two rules discussed above, fare when externalities are brought into the picture? Under strict liability the copier will fail to take into account positive externalities related to both his use and the copyrighted work. When considering the potential preventive measure of forgoing his use altogether the copier will consider only the opportunity cost of his forgone private profit, not the forgone social value of his use. Similarly, when

considering the value of preventing the owner's harm, he will only take into account his private loss, not the positive externalities of the work (unless, the court adjusts the damages to reflect the higher social value of the work but this is very unrealistic). By contrast under a negligence rule a court can adjust the standard to reflect externalities on both the copier's and the owner's side. If the use is of the kind that confers large public externalities the court should adjust the value of the use accordingly when considering the preventive move of forgoing the use altogether. Similarly, when considering any preventive move the court should adjust the expected harm to the owner to reflect any significant externalities of his work. These adjustments of the negligence standard will be reflected in the copier's behavior. The owner on the other hand, will not change his behavior in response to the adjusted negligence standard applied to the copier and will take into account only the private value of the harm in taking precautions. It follows that negligence is superior to strict liability in this respect too: to the extent there are substantial externalities of either the owner's work or the copier's use, a negligence standard properly adjusted, will cause the copier to change his behavior to reflect these externalities. Strict liability will not have such an effect.

Activity levels and other unobservables. The tort literature points out that ordinarily a negligence standard fails to cause the injurer to optimize his activity levels (e.g. how many times the car was driven per month as opposed to whether its brakes were in good repair). The reason is that activity levels are not easily observable by courts and therefore they cannot be incorporated into the negligence standard. By extension, the same applies to any unobservable preventive measures. The same analysis applies to the victim's behavior under a strict liability standard. The tort's dilemma is that there is no rule that can cause both parties to internalize activity levels (and other unobservables). Whoever is the residual bearer of the harm (i.e. the party who will bear the cost of the accident whenever it is not allocated on the basis of fault), will internalize his activity levels and optimize his behavior to reflect that, but the other party will not. Can this consideration enlighten us about copyright accidents? Not necessarily. Activity levels and perhaps other unobservables are sometimes relevant in regard to either the copier or the copyright owner. But it doesn't seem that this consideration is consistently more salient in regard to one of these parties. In the absence of a rule that can optimize both parties' behavior in this regard we are left at an impasse. The pockets of strict liability applied in tort to categories of "abnormally hazardous activities" are often described as a reaction to this problem. These carve-

outs presumably designate categories of “extra dangerous” activities where activity levels problems are particularly troublesome. Strict liability assures that injurers internalize their activity levels in such cases. It seems unlikely that similar general categories could be identified in copyright. Yet a similar approach could be applied on a more ad hoc basis in regard to *both copier and owner*. Under this approach, if the rule is negligence in specific cases where the copier is engaged in an extra-dangerous activity (e.g. digitally distributing a commercial music file bearing notice under the claimed assumption that the owner might approve, given the publicity he is getting) strict liability will be applied. Conversely, if the rule is strict liability there will be no liability at all when the owner engages in a particularly dangerous activity (see for example the facts of *Field v. Google, Inc.*, 412 F.Supp. 2d 1106 (D. Nev. 2006)).

Error Cost. Erroneous decisions by courts of two kinds are inevitable: about the standard of care and its application; and about the amount of damages. Under a negligence standard, consistent errors in regard to the standard of care will distort the copier’s behavior in the direction of the error. Random errors in regard to the standard of care will also distort the copier’s behavior in the direction of over-cautiousness. Obviously, strict liability does not entail such distortions. However, under strict liability consistent (as opposed to random) erroneous damages decisions will distort the copier’s behavior. Such consistent errors won’t distort the copier’s behavior under negligence, except where the errors are of a very large magnitude. In sum, these considerations don’t seem to point strongly at a particular direction. Since it is unclear whether the frequency and magnitude of courts’ mistakes is greater in regard to either the standard of care or damages, there is no clear reason based on reducing the cost of errors to prefer strict liability to negligence or vice versa. Even if there is such a knowable consistent difference it may not be very significant.

Administrative Cost. Negligence involves a higher administrative cost of deciding disputes because of the need to determine and apply a standard of care. Two factors mitigate this effect, however: a) to the extent that over time open-ended negligence standards can be converted into rule-like norms through precedent their cost declines; and b) negligence-based standards usually involve fewer cases requiring decision compared to strict-liability.

Negligence vs. its Cousins

Negligence and strict liability are not the only possible rules. Their alternatives include strict liability + contributory liability and negligence + contributory liability.² The analysis of the various factors discussed above may be applied to each of these alternative rules. Below I provide an extremely telegraphic version of this discussion. *If you are not already familiar with the tort literature you may want to skip ahead to the next section, but consider taking a look at the bottom lines summarized in Fig. 1 below.*

Let's try to factor in all of the considerations discussed above as applied to all alternative rules. Because a significant number of copyright accidents will involve a bilateral situation where both copier and owner can take some reasonable precautions, all three standards with a negligence component are preferable to strict liability in this regard. This is so because, unlike strict liability, these standards make the owner internalize the value of his precautions. When substantial externalities of either the work or the copier's use are present, strict liability fares worst because it incentivizes neither of the parties to consider the social rather than the private value of the work and the use. Negligence + contributory negligence is superior to the two other negligence-based standards because it allows courts to take into account externalities in setting negligence standards that optimize the behavior of both parties. To the extent we believe that externalities are more likely to be important in regard to the copier's use, negligence is preferable to strict liability + contributory negligence. This is so because negligence allows optimizing the behavior of the copier who is the only one who would consider the opportunity cost of forgoing his activity. Considering activity levels leaves us in an impasse since any standard will make only one party consider his activity levels. To the extent we believe that it is possible to incorporate activity levels into the negligence standard this should be done in regard to both negligence and contributory negligence. Also, perhaps a category of "extra dangerous" activities could be applied on an ad hoc basis making certain actions by either copier or owner trigger infringement or no-liability per se. Strict liability is preferable to all other standards in reducing administrative/decision cost, but probably not dramatically so. As for error cost, because it is not clear whether such errors are more significant in regard to damages than in regard to the standard of care there is no clear superiority in between strict liability and negligence. Both of those,

² I am bracketing here comparative negligence for good reasons.

however, seem preferable in this regard to strict liability + contributory negligence and negligence + contributory negligence because each of the two latter standards risks two separate distortionary effects. These effects are summarized in the table below.

	C internalizes value of precautions	O internalizes value of precautions	Activity levels internalized	Externalities internalized in precautions	Administrative cost	Error cost
Strict liability	Yes	No	C only	No	Low	moderate
Negligence	Yes	Yes	O only	C only	moderate	moderate
Strict liability + con. Negligence	Yes	Yes	C only	O only	moderate	high
Negligence + con. Negligence	Yes	Yes	O only	Both C and O	High	high

Fig. 1 Effects of liability rules

Alternative Consequence-Oriented Theories

How if at all does the above analysis change if we switch our normative perspective from efficiency to some of its contenders in the copyright policy arena? In the paper I will focus on three normative approaches: *self-determination*; *cultural-democracy*; and *human-flourishing*. I am unable to develop these approaches here, but let me point out central commonalities shared by them. Each of these approaches is consequence-oriented. Like efficiency these approaches normatively assess legal rules on the basis of a comprehensive consideration of all the consequences that follow from such rules. Unlike efficiency (and welfarist approaches more generally), each of these approaches does not simply try to maximize a total social sum of one good measured by a single scale (e.g. wealth, preferences, utility). Instead each of them starts

with an account of a political-morality ideal from which follows various desiderata. These desiderata cannot always be placed on a single scale so their quantities can simply be traded off against each other, especially not across individuals. In some instances some of those desired ends stand in a clear relationship of normative priority to others. One way of describing this common feature is saying that a “normative premium” is attached to certain effects that follow from the copyright rules (i.e. those affecting the conditions of genuine free choice by individuals, those affecting the level of semantic equality in the cultural sphere; or those affecting central dimensions of human flourishing). Positive effects of this kind may be justified despite overall net losses in individual welfare. It does not follow, however, that these approaches are free from the fundamental cost/benefit tradeoffs of copyright policy. This is so because the competing effects of copyright rules are often operate on the normatively prioritized goods themselves. The important point is that under certain circumstances the optimal point may be different than the one dictated by efficiency because certain effects are valued differently.

One goal shared by all three normative approaches referred to above is the broad open availability of creative opportunities to both amateurs and those who pursue such an activity on a professional basis. Each of the approaches, for its own reasons, is likely to support broadly extending such opportunities even beyond what is required for maximizing subjective welfare and even at the cost of falling short of such maximization.³ One implication of such a policy preference is rethinking the distribution of *risks* of copyright accidents. This means that in some cases, the interest in maintaining creative opportunities broadly open requires limiting the relative share of copyright accidents risk imposed on those who engage in activities that entail such risk. Consider the independent documentaries maker who often has to use materials of unknown provenance; the indie recording artist who runs the risk of “unconsciously” copying

³ Stated tersely, the three approaches would support broad access to creative opportunities (including as professional careers) irrespective of the content or quality of the creative product and even at the cost of not maximizing subjective welfare for the following reasons: a) Self-determination: keeping wide open and free of obstacles opportunities for those whose life plan involves as a main element a creative occupation; b) Cultural democracy: broad equal opportunity for meaning making power as a collective “decision” process; and c) Human flourishing: creative occupations as an important avenue for “meaningful activity.”

successful songs in her music; and the maker of a large digital library of historical images that is available to many small-scale authors and artists on relative open terms. The extent to which the entities involved are exposed to having to bear the cost of copyright accidents will affect the extent to which creative opportunities are widely open and distributed. All cases may merit a lightening of the relative share of the risk placed on the users *beyond the degree justified by efficiency considerations alone*. This is particularly true in cases where the practical implication of imposing the risk on the user is likely to be forgoing his activity altogether. The way to operationalize this normative interest is through a functional equivalent of an externalities analysis. Actions of users relevant for normatively preferred activities should be treated as having an additional social value beyond their private, subjective value to all parties involved. Accordingly, when an optimal prevention strategy is considered the value of the opportunity cost of forgoing the action should be adjusted upwards to reflect the normative premium attached to the activity. In structural terms, this is the functional equivalent of taking into account externalities. But the motivating force here is not any actual claim that externalities are present and the normative framework is not efficiency-based.

Application

Assuming that a negligence standard or one of the three standards with a negligence element is theoretically justified, how should it be implemented in doctrine? I will examine three alternatives:

- Create an independent negligence element as a requirement for copyright infringement.
- Incorporate the negligence inquiry (in regard to the copier, the owner or both) into the existing fair use doctrine.
- Impose a version of contributory negligence, embodied not in an open ended standard of care, but in a few “petrified” rules or presumptions such as whether the owner had affixed notice or registered the work.
 - This approach may or may not be combined with one of the two others.