From the book *Networks, Crowds, and Markets: Reasoning about a Highly Connected World.*
By David Easley and Jon Kleinberg. Cambridge University Press, 2010.
Complete preprint on-line at http://www.cs.cornell.edu/home/kleinber/networks-book/

Chapter 24

Property Rights

The final broad class of social institutions we consider is concerned with the allocation of resources in a society via *property rights*. Property rights give the holder of the right the ability to use a resource, the ability to exclude others from using it, and usually the right to sell or transfer the resource to another person. Property can take many forms, ranging from physical property such as a plot of land or a can of Diet Coke, to intellectual property such as a song or a manufacturing process. In this chapter we will examine how the existence and form of property rights, or the lack of property rights, affects social outcomes for each of these types of property. The central message of this chapter is that the property rights a society chooses to establish will affect the allocations that occur, and some property rights are more likely than others to result in socially optimal allocations.

24.1 Externalities and the Coase Theorem

In Chapter 17 we argued that the allocation of goods that arises in a market equilibrium (for an economy without network effects) is socially optimal. In a market equilibrium, the goods that are produced are assigned to the consumers who value them the most, and any unit of a good that is produced costs society less to produce than it is worth to the consumer who receives the good. This results in maximum total social surplus. The intuition for this fact comes from the observation that at a market equilibrium allocation, each person who consumes a unit of a good pays the cost to society of producing a unit of the good, and anyone who is not consuming the good is unwilling to pay the cost of producing a unit of the good. In this discussion, and in Chapter 17, we have assumed (implicitly) that: the cost of producing the good correctly reflects the true cost to society of producing the good; an individual's willingness to pay for a unit of the good correctly reflects the value to society

Draft version: June 10, 2010

of allowing that individual to consume the good; a producer of a good owns it (has the property right to it) and can sell it at the market price; and, in order to consume the good an individual must buy the good at the market price.

These are important qualifications to the social optimality of market equilibria. To see why getting the values right is so important, let's tell the story of production and trade in slightly different terms. When an individual consumes a can of Diet Coke the individual creates a personal benefit (otherwise she would not voluntarily consume it) and she creates a harm to the rest of society as there is now one less can of Diet Coke that could be consumed by another member of the society. However, if the price that the consumer pays for the can of Diet Coke is equal to the cost to society of producing another can of Diet Coke, then the consumer who buys and consumes the can of Diet Coke compensates the rest of society correctly for the harm she imposes. Well-defined property rights play an important role beneath the surface of this story. One important role that they play is that every good that is produced or consumed is covered by a clear property right. If property rights to the can of Diet Coke are clear, and if no one else is affected by the actions of the Diet Coke producer and the Diet Coke consumer, then property rights cover this transaction completely. If instead the actions of either the producer or the consumer of our can of Diet Coke affect others in a way that is not covered by some property right, then the resulting equilibrium need not be socially optimal. When the welfare of some individuals or firms is affected by the actions of other individuals or firms without a property right that requires mutually agreeable compensation, then we say that an *externality* occurs. Externalities can be negative, as we saw in Chapter 8 where we discussed traffic congestion, or positive as we focused on in Chapter 17 where we discussed goods with network effects. In this chapter we will discuss externalities at a more general level.

Externalities and Non-Optimal Allocations. Let's explore several examples to see how externalities might arise and why they can create non-optimal allocations. First, suppose someone decides to smoke a cigar in a restaurant in which there is one other diner. The smoker purchased the cigar at a price which presumably covered the cost of producing the cigar, so at least between the smoker and the producer there is no externality created by the sale of the cigar. But in the act of consuming the cigar in the restaurant the smoker imposes a harm on the other diner without compensating this person for the harm. Whether the resulting allocation is socially optimal or not depends on the amounts of harm and benefit created.

Suppose that the amount of harm suffered by the other diner is \$10; that is, if the other diner received a compensation of \$10 for the harm created by the smoke, then this person would be just as well off as he would be if the other diner didn't smoke the cigar in the restaurant. On the other hand, suppose the benefit that the smoker receives from smoking the cigar, above the price he paid for it, is only \$5. Then smoking the cigar reduces total social surplus by \$5: the difference between the harm of \$10 created by the smoke and the benefit of \$5 received from creating the smoke. In this case, social optimality requires a smoke-free environment in the restaurant. One mechanism that would achieve this goal is a law that prohibits smoking in restaurants.

An alternative mechanism that would achieve the same goal would be to establish a property right to smoke-free air in restaurants and to make this property right tradeable. In this case, the other diner can choose whether or not to allow the smoker to smoke his cigar by agreeing to abandon the property right in exchange for suitable compensation — that is by selling the property right. Since we're assuming the smoker only values the ability to smoke at \$5, there would be no smoking in the restaurant in this case, because the smoker would not be willing to pay enough to compensate the other diner for the (\$10) harm caused by the smoke. Of course, if the benefit the smoker received from smoking were \$15, rather than \$5, then there would be a trade. The smoker would pay the other diner an amount between \$10 and \$15 for the clean air right, the smoker would smoke his cigar, both parties would be happy with this outcome, and we would have a socially optimal allocation.[1]

In the smoking example, establishing a property right to smoke-free air results in a socially optimal allocation no matter what the individuals' values are for smoking and smoke-free air. Alternatively, a property right that allows the smoker to smoke will also work, because after negotiation between the smoker and the other diner, smoking will occur exactly when it is socially optimal for it to occur. The possible lack of social optimality arises when there is no clear property right or no property right at all. In this case, the individuals may simply disagree about whether smoking is allowed or not and negotiation to a socially optimal allocation seems unlikely to occur.

Finally, a law that prohibits smoking in restaurants will result in a socially optimal allocation if the values are such that the optimal allocation requires no smoking, but it will fail to provide an optimal allocation in the case in which smoking is optimal. In practice, smoking in restaurants is banned in parts of the U.S., and it is useful to relate this to the issue of optimal allocations. There are several possible motivations for a smoking ban, phrased in terms of the discussion above. Maybe the underlying values are such that the optimal allocation always or almost always requires smoke-free air; maybe policy makers believe that individuals consistently undervalue smoke-free air and so would make mistakes if they were allowed to trade; or maybe the costs of enforcing and trading a property right to smoke-free air in restaurants would be so expensive that it is better to just outlaw smoking in restaurants.

Let's explore the last of these motivations — the cost of establishing the property right

[1]In this discussion, and in the rest of this chapter, we will assume that individual's values are independent of their wealth.

— in a bit more detail. In our example there was only one other diner. What if instead there were many other diners and employees in the restaurant? Then no matter who owns the property right (the smoker or the others in the restaurant) a complex negotiation would be required, and if diners come and go, the negotiation would have to be conducted repeatedly. This could easily be so costly as to simply be infeasible, and instead establishing a law banning smoking could be the next best alternative. It is quite likely to be the best we can do if social optimality would typically result in a smoke-free environment anyhow.

Our example of smoking in a restaurant is a simplified story capturing some of the issues that arise in a broad and important domain where property rights play a role: the problem of environmental impact from activities like industrial production. Similar issues can arise here when property rights are not clearly defined and enforced. For example, consider a power-generating plant that pollutes the air and water. The power plant pays for many of the goods that it uses in the process of generating power; goods such as labor, capital equipment, and fuel are purchased on the market at prices that compensate the sellers of these goods for the harm they suffer in giving them up. But the power plant also implicitly uses up clean air and clean water in the process of production. If the power plant had to pay a price that reflected the harm it causes to others — both individuals and other firms — by converting clean air and water into dirty air and water, then the allocation of power, air, and water would be socially optimal. Just as in the case of smoking in a restaurant, establishing a property right to either clean air and water or a property right giving the power plant the right to pollute the air and water, would in principle work to produce a socially optimal allocation of both power and the degree of pollution of air and water. It is worth pointing out that social optimality is unlikely to mean there would be no pollution. Instead, it simply requires that the amount of pollution, like the amount of all other goods, is determined so that there is no reallocation that improves welfare. But also just as in the case of the restaurant, the transaction costs involved in negotiation between the power plant and all those affected by its activities may be prohibitive.

Mechanisms for Determining Socially Optimal Allocations. One difficulty with using property rights and mutually agreeable compensation to determine the socially optimal allocation in our power plant example is the following: how do we discover the true amount of harm created by the pollution? If we simply ask people how much harm they suffer from pollution, and attempt to use this to decide whether or not to allow the pollution, then each person who is harmed has an incentive to overstate the amount of harm. Similarly, the firm creating the pollution has an incentive to overstate the cost of reducing its pollution. However, there is a procedure that can be used to address this incentive problem and we have already analyzed a special case of it.

In Chapter 15, we demonstrated that the VCG procedure results in an efficient matching of sellers to buyers in a matching market (in the specific context of advertising slots and advertisers), even when the buyers' valuations for items is not known. This occurs because VCG pricing makes truth-telling a dominant strategy for buyers. A similar mechanism can be used to induce truth-telling for polluters and for those who suffer the harm created by pollution. The pollution setting is a bit more complex because the valuations of both the buyers (the polluters) and the sellers (the affected individuals) are unknown. Here we might imagine the government running the mechanism, collecting revenue from the polluter, and providing compensation to individuals harmed by the pollution. The goal of this mechanism is to determine the socially optimal amount of pollution; it is not to use the money collected from the polluter to fully compensate those who are harmed by the pollution. In fact, individuals may be better or worse off once the mechanism is run and the payments occur. In addition, the amount of revenue collected may not equal the amount of compensation, so the government may run either a surplus or a deficit.[2]

Actually running a VCG mechanism to determine the optimal amount of pollution would be difficult and costly. The first problem is to determine who is potentially harmed by the pollution and thus who should be included in the mechanism. Next, the mechanism would have to be re-run every time the group of affected individuals changes and every time the polluter wants to change the amount of pollution. This would have to be done for every polluter. The cost of running these mechanisms over and over would be large. Instead, some governments use a more market-based approach in which firms can buy the right to pollute at market prices. These are called cap and trade systems. The U.S. uses a cap and trade system for sulphur dioxide emissions [394]. In a cap and trade system the government provides a number of pollution emissions permits, allows the firms to trade these permits, and requires that any firm emitting pollution own a number of permits equal to the pollution it emits. If the initial number of permits is set correctly, then this too achieves a socially optimal allocation of pollution.

The central idea behind the use of property rights or tradeable pollution permits as a device to solve problems created by externalities is *Coase's Theorem* [113], which roughly says that if tradeable property rights are established and enforced, then negotiation between the parties affected by the externality will lead to a socially optimal outcome no matter who initially owns the property rights. For example, in our earlier scenario of the smoker in the restaurant, all that was necessary for social optimality was to establish and enforce the right to smoke or the right to have smoke-free air. Then trade between the parties would result in a socially optimal allocation. Of course, who owns the right will affect how well off each of the parties is in the resulting equilibrium and they will certainly disagree about who should initially own the right. But, no matter who it is given to, smoking will occur if and only if it

[2]For an accessible discussion of the issues involved in designing an optimal mechanism, see the Nobel Prize Committee's Scientific Background statement in support of the 2007 Nobel Prize in Economics [329], which was awarded for work in mechanism design.

is socially optimal. The same idea applies to pollution emission permits. If property rights are clearly established with someone as the initial owner then trade will lead to optimality. Again the initial allocation of permits will make some better off and some worse off, and it is sure to be politically contentious [81].

The one qualification that is necessary in Coase's argument (that initial ownership is irrelevant) is that it ignores transaction costs and simply assumes that negotiation beginning from any assignment of property rights will lead to an efficient outcome. As we noted in the smoking example this is not plausible when many individuals are involved in the negotiation. Similarly, in the case of pollution, establishing marketable pollution rights is more likely to minimize transaction costs and lead to socially optimal outcomes.

24.2 The Tragedy of the Commons

In a 1968 article in *Science*, entitled "The Tragedy of the Commons" [205], Garrett Hardin offered a compelling story about the inevitable "tragedy" of commonly shared resources. In his story there is a village commons on which any herdsman can freely graze his cattle.[3] Hardin noted that inevitably the commons will be overused to the detriment of all the villagers. He then argued that establishing property rights would solve the problem. These property rights could be privately held, the commons could be sold to some individual, or they could be publicly held. However, if the village continues to own the commons it must carefully limit the use of the commons if we are to have a socially optimal allocation.

A Model for the Commons. Let's construct a simple example to see how Hardin's story plays out. Suppose that there is a village with N people, for some large number N, and each villager owns one cow. If a fraction x of this population of N cows grazes on the commons, then the revenue generated per cow is equal to $f(x)$ for some function $f(\cdot)$. Hardin noted that the fewer cows there are on the commons, the more grass there is per cow, and so the greater revenue per cow. That is, the function $f(\cdot)$ is decreasing. Let's suppose for example that $f(x) = c - x$ for some number $c < 1$. This means that the revenue per cow remains positive until x reaches c, at which point it becomes zero: increasing the fraction of cattle using the commons beyond $x = c$ will cause a negative revenue per cow due to the crowding of the commons.

So the total revenue generated, if the commons is used by an x fraction of all cattle, is equal to $f(x) \cdot xN$, which in our example is $(cx - x^2)N$. Figure 24.1 shows the curve $y = (cx - x^2)N$. If the goal of using the commons is to generate revenue from the cattle that graze there, then the socially optimal fraction of cattle on the commons is the value x^*

[3]The term "commons" comes from the common use of a village green in Europe. Many old villages still have them, although generally not for use by cattle any more.

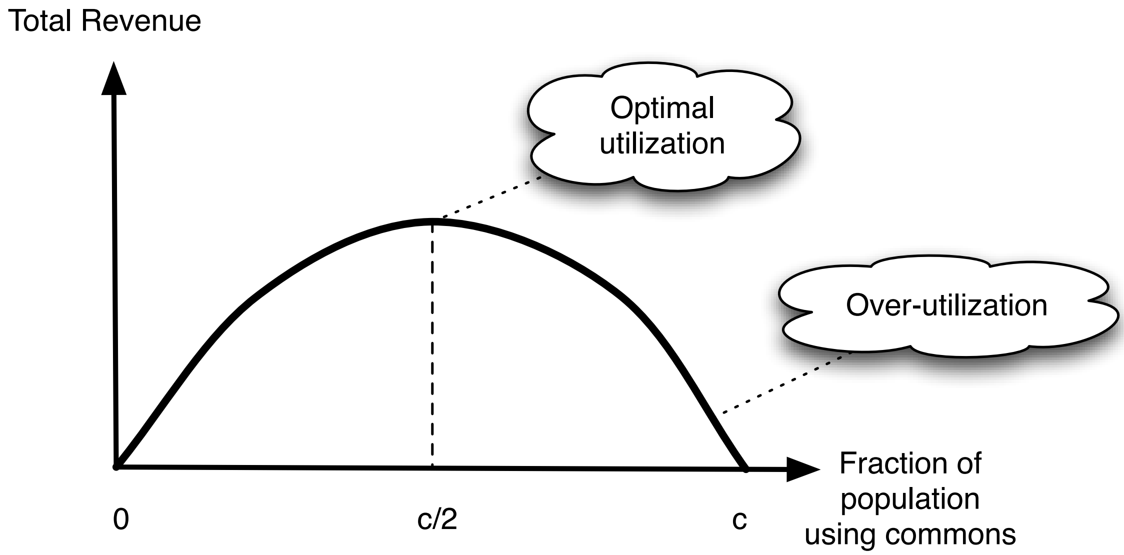

Figure 24.1: In the Tragedy of the Commons, a freely shared resource can easily be overused unless some form of property rights are established.

that maximizes the function $f(x) \cdot xN$; in our case, as we see from the figure, this maximum is achieved midway between the two points where the curve crosses the x-axis, at $x^* = c/2$. Thus the maximum revenue is

$$f(x^*) \cdot x^* N = \left(c - \frac{c}{2}\right) \cdot \left(\frac{c}{2}\right) N = \frac{c^2 N}{4}.$$

We analyzed a similar function in Chapter 17 to describe the maximum price that users would pay for a good with network effects when an x fraction of the population is using it. In that setting there is also an externality, since each individual's willingness to pay for the good is affected by the total number of users. There are key differences, however, between the setting of network effects and the issues involving the commons. With network effects, each additional user had a positive effect on the people already using the product, due to the positive externality. Here, on the other hand, increasing the population on the commons has a negative effect on each individual already there, due to the crowding. Revenue is maximized when the increased crowding due to extra cattle is traded off against the greater revenue from having more cattle using the commons — and this maximum occurs at a fraction $x = c/2$ of the population in the case of our simple example.

This is what happens if we are able to choose a quantity of cattle on the commons so as to maximize the total revenue. But what happens if the commons is freely available to all the villagers? In this case, villagers will add cattle to the commons as long as the revenue from doing so is positive. This will eventually lead to the revenue per cow being driven to zero (as long as there are enough cows available, which in our example is true). To see why this must happen, note that if there is presently a fraction x of cattle on the commons and $f(x) > 0$, then there are villagers whose cows are not currently using the commons, and it would be in the interest of any one of these villagers to add his cow to the commons in order to reap even a small bit of positive revenue. This can stop only when the fraction of cattle on the commons reaches the number $\bar{x}$ for which $f(\bar{x}) = 0$, which in our example is $\bar{x} = c$. At that point, there will be twice as many cows on the commons as is socially optimal and the total revenue from the collective use of the commons is $\bar{x} \cdot f(\bar{x})N = 0$. This is Hardin's tragedy. The village owns a resource which is clearly valuable, but because its use is not restricted, everyone who uses it receives zero reward from using it.

Avoiding the Tragedy. A further "tragic" aspect to this loss of social optimality is that the village could easily solve the problem. There are two obvious ways to do it and several variations on each approach.

One approach is for the village to continue to own the commons jointly but somehow limit the quantity of cattle to the socially optimal number. This could be done either by charging a price for grazing on the commons, or by simply setting the fraction of cattle that are allowed on the commons at the optimal value x^*, which in our example is equal to $c/2$. If the village charges a price, the optimal price to charge is $c/2$ per cow. To see why this is optimal, note that a villager will add his cow to the commons if and only if the revenue from grazing a cow on the commons is greater than the price. Thus, in an equilibrium, the revenue from grazing a cow on the commons must be equal to the price. So the equilibrium will be the value of x which solves $f(x) = c/2$ and this is $x = c/2$. Alternatively, the village could sell grazing rights for x^*N cows. Again, the maximum price per cow the village could charge for such grazing rights is $c/2$ per cow, as we just found above. Either variation on this method results in the socially optimal use of the commons and a revenue to the village of $c^2N/4$.

Instead of owning it jointly, the village could sell the commons to a large owner of livestock who has many cows. This livestock owner who buys the commons would put x^*N cows on the commons too as this is the number of cows that maximizes his revenue. The maximum price that the village could sell the commons for is the revenue that the buyer would receive from optimal use of the commons, which is again $c^2N/4$. So in either approach — joint ownership by the village with an appropriate price charged per cow for grazing; or outright sale to an individual — the village receives a revenue of $c^2N/4$, and the commons is used by the optimal fraction of cattle.

In Hardin's village commons example all that is needed to avoid the tragedy is to establish a property right. It can either be a property right held by some form of government which optimally limits use of the resource or it can be a privately held property right. Just as with Coase's argument it does not matter (for social optimality) who owns the property; all that matters is that someone owns it. Hardin uses this story to argue why too much pollution is produced, why national parks are overused (if there are no admission charges or limits on the number of users), why there is over-fishing, and, more controversially, to argue that over-population and thus overuse of the world's resources is inevitable. In each of these examples there is an externality as the actions of any one firm or individual affect others, and if there are no property rights, there is no compensation, and thus no reason to use the resource optimally.

24.3 Intellectual Property

The property that we discussed in our exploration of Coase's Theorem or Hardin's tragedy existed independently of investment by individuals or firms. The air in the restaurant, and the air or water that might be polluted by the power plant, were there regardless of what property rights society creates. Similarly, once the village green is established, it exists independently of any actions by the villagers. But there is a difference between the green and the air or water; and, as we will see there is yet another difference between either of these examples and intellectual property.

To this point in our analysis we have focused on efficient use of resources without regard to where theses resources came from. For air and water this seems reasonable. These natural resources are not created by any human effort and no effort is needed to make them useful. The village green is also a natural resource, but human effort can reasonably make it more or less valuable. Cutting the grass regularly, removing weeds, and applying chemical or natural fertilizer can all make the green more productive. What property rights are established will matter for whether these activities are carried out. If no one owns the green, then no one can fully benefit from undertaking these costly but productivity-enhancing investments, and it seems likely that they will not occur. Solving the property rights problem by assigning the property rights to someone solves the problem of inefficient use of the resource and it also solves the problem of a lack of incentives for productivity-enhancing investment. If an individual owns the green, then that individual reaps the reward of investment in the green, and this provides the individual with the proper incentive to make exactly those investments that are worthwhile. The individual will undertake any investment in the green that generates more revenue than the investment costs. Similarly, if the village owns the green jointly, but sells rights to use it, the village also has the proper incentives to maintain the green. So assigning property rights to the village green is even more important for social optimality than our earlier discussion indicated.

Rivalrous and Non-Rivalrous Goods. Perhaps surprisingly given our focus to this point on the desirability of assigning and enforcing property rights, there are commodities for which assigning property rights can lead to less efficient use than would occur if there were no ownership at all. Consider, for example, the result of a creative process. This could be a book, a song, a new computer program, a new variety of crop, a new drug to treat cancer, or a new manufacturing process to produce batteries. The use by one individual or firm of the outcome of the creative process (not the physical object, if one exists, but the idea that is captured in the physical object) does not affect the ability of anyone else to use it.[4] The same process for making a battery or making a drug, for example, can be used by an unlimited number of people. Everyone can listen to a song downloaded from the Internet, or read a book on-line, without affecting the ability of anyone else to listen to the song or to read the book. In contrast, the can of Diet Coke that one person drinks cannot be consumed by anyone else. The grass on the village green that one villager's cow consumes cannot be consumed by another cow. Goods such that use or consumption by one user precludes use or consumption by any other potential user are called *rivalrous goods*; goods that can be used or consumed over and over are called *non-rivalrous goods*.

For non-rivalrous goods, establishing property rights can interfere with efficient use of the goods. The owner of the good can charge for use of the good, and any non-zero price may result in some potential users not purchasing the good, and thus not using it (or at least not using it legally). This results in an inefficiency as there is no cost to society in allowing everyone to use the good, so prohibiting anyone from using it is wasteful. This contrasts with the socially optimal use of the village green in which it was necessary to limit the use of the good (the green) in order to achieve social optimality. In that case the good was rivalrous. But for non-rivalrous goods, Hardin's ideas do not apply.

This is, of course, not the end of the story about property rights for non-rivalrous goods. We also have to ask where the good came from and whether it would exist without property rights. If the creator of a non-rivalrous good does not own it, then the creator's ability to profit from his creation is limited. Even so the creator may receive some benefit from his creation. Humans discovered many useful ideas (using fire to cook meat, for example) long before there was any legal protection for the creators of the results of these ideas. These early inventors received a direct benefit from their own use of their inventions. Plato wrote and Mozart created music with little or no protection for their works. They too received benefits both directly and from users of their work. So some creative activity will exist and creators will receive some benefits from creation without property rights that protect intellectual property. But it is not obvious that the socially optimal amount of creative activity will occur without these property rights, nor is it obvious what form those rights

[4]Use by many people may of course affect how profitable it is to use the result of the creative process, but it does not affect the feasibility of use by others.

should take in order to generate the optimal amount of creative activity.

The relevant question is the tradeoff between providing incentives for creative activity and allowing efficient use of the creation once it exists. No one will have the creation unless the creator releases it, so the creator at least can benefit from the first use of a valuable good. If there are no protections then eventually the good will become publicly available and the creator's ability to charge for it will vanish. Thus, without property rights, the financial reward for creative activity may be small, and in the modern economy with the advent of fast and inexpensive copying and communication that ability may be very small indeed. Instituting property rights to creative activity increases the creator's incentives, but it does so at the cost of inefficient use of creation once it exists.

Copyrights. Rather than attempt to determine how this balancing should occur in the abstract we will look at several examples. First, let's consider the case of books, songs, plays, television shows and movies. In the U.S. all of these creations are covered by the *copyright law* (currently the Copyright Act of 1976) which gives the creator of the work the exclusive right to copy the work, to distribute it, to modify it, or in the case of songs, plays, television shows and movies to perform the work. This right lasts for 70 years beyond the lifetime of the creator of the work. The owner of the copyright has the right to transfer this right to others.

It is important to be careful about what use of copyrighted works is allowed, and what uses are illegal without permission of the holder of the copyright. First, the copyright holder's exclusive right to copy the work does not actually prohibit all copying. The doctrine of *fair use* has evolved over time to permit limited copying for non-commercial use of parts of a copyrighted work. For example, it is permissible to quote a copyrighted work in a review of the work, in a scholarly article, or in a classroom. The doctrine of fair use is outlined in Section 107 of Title 17 of the Copyright Act of 1976 (the entire law can be found at http://www.copyright.gov/title17/). This law does not define exactly what is or is not fair use; instead whether a use of copyrighted work is fair or not is to be determined case by case with the intent of the copier a critical factor in making the determination. Second, copyright law does not prohibit the owner of a copy of a work from reselling that copy. (This is distinct from making another copy and selling the new copy, which is prohibited.) So although it is illegal to copy a book or CD and transfer that copy to someone else, selling a legally obtained copy of the book or CD is permitted.

Copyright law grants a monopoly (single seller) to the holder of the copyright. Generally monopolies are harmful as the monopolist artificially restricts the use of the good by setting a price above the socially optimal price. For copyrighted works the socially optimal price would clearly be zero if this price had no effect on the creation of the work. The extent to which the protection afforded by a copyright is necessary to provide sufficient incentives to produce goods that are currently copyrighted is unclear. Some authors, such as Boldrin and Levine [66] for example, argue that copyrights should not exist at all as they believe that copyrights are not necessary for innovation and they prevent efficient use of works once they are created. The more common view is that copyrights are a necessary evil: they do prevent efficient use, but without them the amount of resources devoted to creative activities would be inefficiently small.

Patents. Next, let's consider inventions such as a new drug, a new manufacturing process, or a new piece of computer hardware. The inventor can apply to the United States Patent and Trademark Office for a patent on the invention, and if the patent is granted, the inventor has the right to exclude others from using the invention for a fixed period of time, usually 20 years. The web site of the U.S. Patent and Trademark Office[5] provides a description of the law.

The economic role of patents is much like that of copyrights. They increase the reward to inventive activity at the expense of inefficient use of the patented invention once it is created. However, patents are different from copyrights in several ways. First, granting of copyrights for original work is automatic; the creator only needs to indicate that the work is copyrighted. Patents are obtained by filing an application with the U.S. Patent and Trademark Office which reviews the application for originality of the work. Second, enforcement of copyrights and patents is generally left to the holder of the copyright or patent. The primary exception to this is copyright infringement of songs and movies over the Internet, and the creation of devices or procedures that can be used to circumvent digital rights management (technologies designed to prevent copying). These activities have been partially criminalized by the Digital Millennium Copyright Act of 1998. Third, the investment in research and development needed to create many patentable commodities is very large relative to the investment needed to create most works of art. The pharmaceutical industry, for example, spends large sums on R&D and it seems likely that this would not occur without the ability to patent discoveries. So the case for restrictive and strongly enforced patent law is more compelling than that for copyright law. As with many other aspects of property rights, the trade-offs here are complex, and they remain the subject of active consideration.

24.4 Exercises

1. Consider an airport that is trying to sell exclusive rights to operate a wireless access network in its terminals. Depending on what proportion of all travelers use the network, it can become congested and result in a low-quality experience for everyone.

[5] http://www.uspto.gov/main/patents.htm

Specifically, suppose for simplicity that there are N travelers at any given time during normal hours in the airport, and if an x fraction of them attempt to use the network concurrently, then the payoff to each of them will be $\frac{1}{2} - x$. (We can view this payoff as the amount they'd be willing to pay for the service.)

(a) When the airport sells the rights to operate the network to a third-party access provider, this provider will make back what it pays for the rights by charging travelers in the airport a fee to connect to the networks. How much should the airport expect to be able to sell the rights for, how much will the third-party access provider in turn charge to travelers, and what will be the resulting sum of payoffs to all travelers? Explain.

(b) Suppose instead that the airport were to let people use the service for free. What would be the sum of payoffs to all travelers in this case? Explain.

2. Consider the set-up from Exercise 1, but now let's change the scenario a bit to suppose that there are two populations among the travelers through the airport, with members of one population valuing the wireless access service more than members of the other.

In particular, suppose that travelers of Type 1 receive a payoff of $\frac{1}{2} - x$ when an x fraction of all travelers use the service. (Here, the x fraction is based on the total usage by travelers of both types, since travelers of both types contribute to the shared congestion.) Travelers of Type 2 receive twice this payoff, $1 - 2x$, when an x fraction of all travelers use the service. Notice that both payoffs become 0 when x reaches $\frac{1}{2}$, since at that point the high congestion makes the service useless for everyone.

As in Exercise 1, the airport is going to sell the right to operate the network to an access provider, who will then charge a single price to all travelers (regardless of which type they are).

(a) Suppose that the airport and the access provider know that half of all travelers are of Type 1 and half are of Type 2. How much can the airport expect to sell the rights for, and how much will the access provider charge?

(b) Now let's consider a variation: suppose instead that only 5% of all travelers are of Type 2, and the rest are of Type 1. Again, both the airport and the access provider know this. Now how much can the airport expect to sell the rights for, and how much will the access provider charge?