Applications of Economics

A diverse collection.

Economics as a Unifier of Law: A True Story

While teaching a course on intellectual property I came across an article that Lou Kaplow, a prominent (and very able) law and economics scholar at Harvard, published in 1984.[1] What especially interested me about the article was that I wrote it. In 1981.

Neither Lou nor I engaged in plagiarism, with or without the aid of a time machine. I first saw his article more than two decades after I wrote mine; so far as I know he had never seen my article when he wrote his and may not have seen it yet. His was on patent law, mine on criminal law. Yet they were, in their essence, the same article. Each hinged on a single simple idea, simple enough that I can explain it in a page or two, the same idea in both articles.

The conventional view of patent law is that it rewards inventors with a temporary monopoly in order to give them an incentive to make and reveal inventions. Lou was looking at the question of how long the term of the monopoly should be. The longer the patent term the more valuable the patent, which raises the question of how large the reward for making an invention should be.

There is an obvious answer, obvious at least to an economist. Set the reward equal to the social value of the invention. That way it will be in the interest of inventors to make any invention that costs less than it is worth. Applying that rule in practice faces a host of difficulties but the theoretical answer seems straightforward.

It is also, as Lou pointed out, wrong. The reason it is wrong is that giving the reward is costly. For reasons familiar in economic theory, the benefit a monopoly provides to the monopolist is less than the cost it imposes on his customers, the difference being what economists refer to as deadweight cost.

Imagine that there is an invention whose social value we can somehow measure as ten million dollars. Further imagine that we have calculated that ten years of monopoly will earn the inventor a million dollars a year, adding up to a reward of exactly ten million. Should we give it to him?

No. Reducing the term of patent protection from ten years to nine reduces his reward from ten million dollars to nine million. If the cost of making the invention is less than nine million dollars he will still make it, we will still get the benefit, and we will have a year less of deadweight cost.

If the cost is between nine and ten million the reward would be less than the cost so the invention will not get made. That is a cost but a cost of less than a million dollars, since we (consumers and inventor together) will lose a ten million dollar benefit but save a cost of between nine and ten million. To find the optimal length of protection we would need more information, a probability distribution for the cost telling us how likely it is that any reduction in the reward will result in the invention not being made and a way of calculating how large the deadweight cost is for any length of protection. But even without that information, it is easy to see that the optimal term of protection can be less than ten years and only a little harder to see that it almost has to be. If the term of protection is 10 years - X, both the chance that the shorter term will result in not getting the invention and the net cost of failing to get the invention are proportional to X, making the combined

[1] "The Patent—Antitrust Intersection: A Reappraisal, 97 Harv. L. Rev. 1913 (1984).

effect proportional to X squared, what an older generation of scientists referred to as of the second order of smalls. The savings in deadweight loss is proportional to X, since that is how much less time we bear it. So if X is small enough, the gain has to be larger than the loss.

Lou concluded that the conventional answer, optimal reward equal to value of invention, was wrong. As long as giving a reward costs something, the optimal reward is less than that. That was the central point of Lou's article and it was correct — obviously correct, once stated.

My article was on how to calculate the optimal penalty for a criminal offense.[2] In that case too there was an obvious answer, obvious at least to an economist, and the logic of the answer was the same. Set the penalty (more precisely, the cost to the criminal of the combination of penalty if convicted and chance of conviction — if criminals are risk neutral, the product of the two) equal to the damage done by the offense. The only offenses it is then in the interest of the criminal to commit are those where the gain to the offender is greater than the loss to the victim, in which case deterring the offense would make us, on net, worse off.

That obvious answer is wrong for precisely the same reason as before. Catching and punishing criminals, like rewarding inventors, is costly. If an offense costs the victim $100 and benefits the criminal by $99, it imposes a net cost of $1. But if raising the punishment by enough to deter that offense costs $10 in extra enforcement and punishment, costs of paying cops and running prisons, we are better off not doing it. The level of punishment that minimizes net costs is the level at which any further increase would cost as much in extra enforcement and punishment costs as it would gain in net damage deterred.

There are differences in detail between my case and his, in particular the fact that the cost of deterrence is sometimes negative, since if you deter an offense you do not have to punish it. Anyone sufficiently interested can find the details in the relevant chapter of my *Law's Order* and, in a more mathematical form, in a virtual footnote to that chapter. But the logic of the two articles is identical, as is the logic of the two errors, one in patent theory and one in criminal theory, that they critique.

The two articles are evidence of how economics unifies the law, makes the same analysis, the same ideas, the same logic apply in apparently unrelated legal fields.

Executive Compensation and the Economics of Insurance

Since one of the things legal rules do is to allocate risk, one of the subjects routinely covered in the economic analysis of law is risk allocation, conventionally presented in terms of the economics of insurance. Designing optimal rules in this context is hard because there are three different objectives and no reason to expect the same rule to be best for all of them.

The first objective is risk spreading. For familiar reasons, many people prefer a certain income of (say) $50,000 a year to a coin flip between $10,000 and $90,000, even though the average outcomes are the same. Insurance provides a way of converting the riskier outcome into the less risky.

The second objective is optimum incentives for controlling risk. Often, although not always, the person who starts out bearing the risk is the one in the best position to take precautions against it.

---

[2] "Reflections on Optimal Punishment or Should the Rich Pay Higher Fines?," Research in Law and Economics, (1981).

Insurance is then the problem, not the solution. Once I have insured my house or factory for most of its value, my incentive to take costly precautions to make it less likely to burn down is low. If an insurance company is so imprudent as to insure something for more than its full value, the chance of a fire may become very large indeed. This is the problem referred to in the literature as moral hazard. "Moral" in this context has a meaning, now largely obsolete, close to "psychological," as in "moral support."[3] The hazard is due to the incentives of the actors not the physical characteristics of the situation.

The third objective has to do with the fact that the choices we make signal our private information, sometimes incorrectly. That I want to buy lots of life insurance today is evidence that I know something the insurer does not about my chances of living to tomorrow, which is a reason for him to offer the insurance to me only at a high price. The same argument applies, although less strongly, to anyone who wants to buy insurance against any risk that he has better information about than the seller. Wanting to buy is evidence that the risk is higher than average, a fact which the seller will take into account in pricing the insurance.

That is a problem if you want to insure against an average or below average risk, since the seller will misinterpret your desire to do so as evidence that the risk is larger than it appears and charge accordingly. A classic example is the market for used cars. Since sellers know more about their cars than buyers, the fact that I want to sell my car signals that it is likely to be a lemon so I get offered a lemon price. That makes it even less likely that I will choose to sell it if it is not a lemon. Lemons end up selling for a lemon price, cream puffs often don't sell even if they are worth more to a potential buyer than to their present owner. The problem is known in the insurance literature as adverse selection.

The same set of problems applies to executive compensation. One way of better aligning the interests of executives with those of stockholders is to require the executives to own stock in their company; that reduces not only moral hazard but adverse selection as well. An executive whose private information implied that hiring him would be bad for the company, perhaps because he planned to keep the job only until a better offer, expected shortly, came through or because he planned to supplement his income at the expense of the stockholders, would have less reason to take the job if it paid in part in stock which he was not permitted to sell until some time after leaving the company.

Requiring executives to invest a large fraction of their wealth in the stock of the company they worked for might be a good way of aligning their incentives with the interest of the stockholders but it would be a bad way of spreading risk. If the company did well, both the executive's salary and his stock portfolio would go up. If it did badly, both would go down. Better for him to spread the risk by owning stock in some company, practically any company, other than the one he works for.

Here again economics unifies the analysis of apparently unrelated subjects: Insurance, the used car market, and the compensation of executives.

---

[3] The same obsolete meaning shows up in the quote attributed to Napoleon (obviously in French) that in war "The moral is to the physical as three to one."

Should Contracts Made Under Duress be Enforceable?

A mugger in a dark alley offers you a bargain: A hundred dollars for your life. Since you are not carrying that much cash, you ask if he will take a check. When you get home, you stop payment; contracts made under duress are not enforceable. Expecting this he either refuses your check or accepts it, reneges on his side of the bargain, and cashes it before the news of your death reaches your bank.

You prefer paying a hundred dollars to being killed. He prefers receiving a hundred dollars to killing you. It occurs to you that even a contract made under duress produces benefits for both parties and so should be enforceable. What's the problem?

The problem is that making such contracts enforceable makes mugging more profitable, which makes it more likely that you will be offered similar bargains in other dark alleys in the future. Whether contracts made under duress ought to be enforceable depends, economically speaking, on elasticity at two margins. If making the contract enforceable results in a large increase in how often duress occurs but only a small decrease in the damage done each time, almost certainly the case in my mugging example, we are better off if such contracts are unenforceable.

A few hundred years ago, prisoners of war, at least officers, were routinely asked to give their word not to try to escape and then permitted to wander around unguarded. Sometimes the prisoner was even permitted to go home, having promised not to rejoin his army until he had been exchanged for a prisoner of equal rank from the other side. A prisoner who violated his parole had proved that he was not a gentleman and treated accordingly by people on his own side. The parole system made war somewhat less costly for both sides, which presumably increased the amount of war, but it seems unlikely that the effect was very great. It substantially decreased the cost born by captive and captor. Its eventual breakdown during the Napoleonic wars probably made the world a worse place.

For a similar tradeoff in a different context, consider the question of whether Augusto Pinochet ought to have been tried for crimes he was accused of committing while dictator of Chile. Enforcing the terms of an agreement that immunizes an ex-dictator from prosecution makes it less expensive for dictators to commit crimes while in power. But refusing to enforce the terms of such agreements makes it more expensive for dictators to give up power. Pinochet is one of the rare examples of a dictator who voluntarily relinquished power to an elected government. If his reward was a jail cell, the next dictator might not make that mistake.[4]

In all of these cases, the agreement was made under duress: the threat of killing you, of keeping a prisoner of war locked up, of keeping dictatorial control over a country. The same is true of a peace treaty; threatening to drop bombs on someone until he agrees to your terms is about as clear a case of duress as one can find. Yet most of us feel as though it is a good thing for treaties to be kept; in a world where a treaty is only a piece of paper, it is hard to end a war short of annihilation of the loser.

Science fiction readers may want to consider the reaction of the aliens in *Footfall*, by Niven and Pournelle, to individual humans who surrender and then violate their parole.

---

[4] For an account of the actual cases, see Wikipedia. Pinochet spent some time under house arrest but was never jailed.

Accidental Incentives: Speculation and Reputational Enforcement

If a copper mine shuts down in Chile the price of copper goes up, giving other producers an incentive to produce more copper, consumers an incentive to consume less. The objective of the individual producer or consumer is to improve his own welfare not the functioning of the economy but he is led, as Adam Smith long ago pointed out, as by an invisible hand to achieve a desirable objective that is no part of his intent.

From the individual's point of view the effect on the market is an accident, an undesirable one; producing more makes prices go down, which is not what the producer wants to happen. But seen in a broader sense the market is a system of feedbacks, signals, that give the individual actors the right message, make it in their interest to produce more copper or consume less when and only when doing so improves the overall outcome. The gain to the individual actor is a measure of the social gain from his action.[5]

There are, however, cases which appear the same but are not, cases where the link between the benefit his action produces and the reward to the actor is in some sense an accident.

One example is speculation. A successful speculator buys things when they are cheap, sells them when they are expensive, both benefitting himself and smoothing out price movements. The latter effect can be a very large benefit to other people. A speculator who sees a food shortage coming well in advance and takes the opportunity to buy grain early, driving up its price, gives less well-informed people an incentive to use less grain, plant more of other food crops, import more food, export less, take actions that alleviate what might otherwise be a serious famine.

But in the case of speculation, the profit is not a measure of the benefit produced. To see that, consider a case where the speculator learns of the shortage only a short time before everyone else would have learned, short enough so that the increase in price due to his activity has no significant effect on other people's behavior. He can still make a lot of money, not because he has produced valuable information but because goods belong to him instead of someone else when their price goes up. One implication, pointed out long ago in a classic article by Jack Hirshleifer, is the possibility of inefficient speculation. It might pay a speculator to spend a million dollars acquiring information about future price movements, information whose social value is zero, with his profit coming at the expense of whomever would have held the goods when their price went up if he had not bought them first.

The point is illustrated by a famous law case, Laidlaw v. Organ. A tobacco trader in New Orleans got word of the treaty that ended the war of 1812 and took the opportunity to buy a large quantity of tobacco from a seller who had not yet heard the news at the low price that had resulted from the British blockade. When the seller discovered that the war was over he attempted to renege on the contract. The court held that the contract was binding. A modern example would be paying for high speed data links to a commodities market in order to be able to engage in profitable arbitrage a second earlier than anyone else.

The same pattern occurs in another important context: reputational enforcement. I buy a jacket from a department store that guarantees to refund the purchase price if I am not satisfied. When I return it they refuse to give me back my money. It is not worth the cost of suing them but it is worth telling my friends, and anyone else who will listen, how badly I have been treated. The result

[5] For a more detailed defense of this claim see my *Price Theory* or *Hidden Order*.

is that other people stop buying from the store. That is a reason for stores to live up to their promises even if they are not likely to be sued for breaking them.

Reputational enforcement plays a large role in encouraging commercial honesty but, just as in the case of speculation, the incentive that makes it work is not linked to the usefulness of the behavior. People stop buying from the store not to punish it for mistreating me but to protect themselves.

To see why this matters, suppose it is not immediately obvious which party to a dispute is at fault. In order for interested third parties to punish the right person they have to know who it is. In most cases, they have little incentive to investigate the claims of each since they have the easier alternative of no longer doing business with either. Anticipating that I realize that, having been cheated, I am better off saying nothing. At which point the mechanism for keeping firms honest stops working.

It follows that reputational enforcement works only in contexts where it is easy for interested third parties to discover who was at fault; elsewhere I have argued that one function of arbitration is to lower the cost to third parties of doing so. My point here is a more general one: It is worth distinguishing between those feedback mechanisms that work because the incentive to act measures the value of acting and those that work, as it were, by accident.

Why Bureaucrats don't Maximize their Budgets

William Niskanen, in a book published many years ago, proposed a simple model of government bureaucracy. The more money a bureaucrat controls the more important he is, so bureaucrats want to maximize their budgets. The legislature knows how much any level of output from a bureau, whatever good or service it produces, is worth to it but not what it would cost. So the rational bureau misrepresents its production function in a way designed to trick the legislature into giving it the largest possible budget. It does so by finding the largest level of output that it can produce at a cost just below the value of that level of output to the legislature and then exaggerating the cost of any lower level of output by enough to make it higher than its value.

When I first read the argument it struck me that it contained a fundamental mistake. So far as I can remember, I never published that conclusion, so am doing so here.

Consider two bureaucrats. Abe has a ten million dollar budget and is required to purchase $9,900,000 worth of paper to be sent to the IRS for printing tax forms on, leaving him $100,000 for himself, his secretary, and rent for his office. Bernie has a one million dollar budget and is required to do nothing at all. Which would you rather be?

Generalizing the example, I suggest that the size of the budget in Niskanen's model ought to be replaced with the surplus, the difference between the size of the budget and the lowest cost at which the output the bureau has agreed to can be produced. That difference represents resources that bureaucrats can divert to their own purposes.

Niskanen's argument was criticized by Dunleavy but, so far as I can tell from a summary of his work, he did not offer my alternative. I have found no references to anyone else doing so.

Imperfect Information and Suboptimal Design

In designing the sensor of a digital camera or a cell phone there is a tradeoff between pixel count and sensitivity. The more pixels the sensor is divided into, the smaller each pixel is. More pixels give higher resolution. Larger pixels are more sensitive, which means less noise, better low light

performance. The number of pixels is routinely included in a camera or cell phone's characteristics. The sensitivity is a subtler issue there; is no simple number that sums up image quality. A good review will describe the quality of the image and provide examples, but most consumers are unlikely to make use of that sort of detailed information.

If consumers are aware of pixel count and value it, are not aware of sensitivity, the manufacturer has an incentive to give the sensor more than the optimal number of pixels. He is trading off two desirable features only one of which is easily observed by the consumer, hence will put too much weight on that one. The result is a twenty megapixel camera, higher resolution than most photographers most of the time have any use for, that takes lower quality pictures than a camera with the same size sensor divided into fewer and larger pixels.

There are other examples of the same pattern. A car's mileage can be increased by making it lighter at the cost of making it less protection in a crash. Dishwashers use substantially less water now than they used to — also, in our experience, do a worse job and take longer to do it. The cause in that case is federal regulation, which sets a 3.5 gallon limit for a dishwasher's normal cycle, but the logic is the same — water consumption is a quantifiable cost, so it is tempting to restrict it. How clean the dishes come out, how much rinsing is needed before they go in, are harder to measure and regulate.

The same analysis applies to power usage. A dishwasher we once had saved power by having no heated drying cycle, relying on the heat from the wash to dry the dishes — and frequently didn't. A review can report on how dry the dishes come out, but it is less immediately obvious information than whether the dishwasher qualifies for the (government backed) energy star badge.

When buying a product that has tradeoffs among two or more valuable characteristics, avoid the temptation to give too much weight to a characteristic that has a simple objective measure. A product that does unusually well on that may do so at the expense of other desirable characteristics that are less legible. The camera or cell phone with the highest pixel count, the dish washer that uses the least power or water, is probably not the best.

The approach applies more broadly, to any context where there is a tradeoff between two desirable features and a reason to favor one of them. Consider my favorite recipe for apple crisp. The ingredients for the topping are a cup of flour, a cup of rolled oats, a cup of brown sugar (and a half cup of butter). Having identical amounts of the three main ingredients makes the recipe easier to remember; it would be harder if it was a cup, seven-eighths of a cup, and a cup and a quarter. That's a reason to make the quantities come out the same even if the topping would be as good, even if it would be a little better, with slightly more or less of one of them. It follows that if you think the dish comes out too sweet (arguably it does) you should feel free to cut the brown sugar to three-quarters of a cup without worrying about substituting your judgement for the expertise of the cookbook (or web page) author.

The U.K. used to have a system where a child at ten or eleven took an exam called the eleven plus; a good score meant that he would be channeled into the academic track, grammar school and possibly college. My family spent most of a year in Cambridge, England, in the early fifties, when that system was still in general use;[6] my sister was in the year just before the test. Her entire school year was spent studying how to do well on that exam.

[6] The Eleven plus still exists but taking it is voluntary and only some schools use it.

She was not the only one.

> Children in school were drilled in the 11-plus until it was "coming out of their ears". Families had to play the system, little booklets were available from local newsagents that showed how to pass the exam and contained many past papers with all the answers provided, which the children then learned by rote. (Wikipedia)

The test was supposed to measure how smart you were. It actually measured how well you had been trained to take the test. The problem has been generalized as Goodhart's law: When a measure becomes a target, it ceases to be a good measure.[7]

Online Applications of Economics

Instead of a Graduation Speech

The Function of College: An Orientation Speech

[7] This phrasing is due to anthropologist Marilyn Strathern. Goodhart's version was "Any observed statistical regularity will tend to collapse once pressure is placed upon it for control purposes."