# Economics and Security: Scribe Notes 4/12

Sidhanth Mohanty

**First case study -- a market of lemons and limes**

The thought experiment about a market for lemons came up while studying the used car market. A simplification to the used car market is to assume there are two kinds of cars, ones which are reliable, which we refer to as "limes" and those which are not reliable, which we call "lemons". Additionally, let's assume that roughly half the cars are lemons and rougly half the cars are limes.

Every buyer has a valuation of lemons and limes. In particular, a buyer is willing to pay up to $x for a lemon and up to $y for a lime. To make matters concrete, let's think of x as 5000 and y as 9000. Similarly, every seller is willing to accept at least $x' to sell a lemon and at least $y' to sell a lime. And once again, to keep things concrete, think of x as 5000 and y as 9000.

Let's assume that a buyer does not know whether a given car is a lemon or a lime. So given a random car, the buyer offers $7000 since that is the expected amount they are willing to pay (.5$5000 + .5$9000). If the car is a lemon, the seller who is willing to sell the car for $5000 immediately accepts the bid, whereas a seller selling limes rejects the bid and takes their car out of the market. The effect of every buyer bidding $7000 is that the distribution of lemons and limes in the car market gets skewed and lemons become an overwhelming majority. And when this happens, the bids slowly approach $5000 and the market arrives in a situation where exclusively lemons are sold, and it becomes impossible to sell a lime for its actual value.

It is an undesirable situation for both sellers of limes as well as for buyers who are highly interested in buying a lime.

In classical economics, there are several theorems articulating the phenomenon that markets converge to equilibria where every commodity is sold at a price equal to its value. Nevertheless the above thought experiment illustrates a counterexample to this phenomena. This raises the question as to what the discrepancy between the above discussion related to the market for lemons, and classical economic wisdom.

The key to this discrepancy lies in the fact that the classical economics results work in a setting where every agent in the market -- every buyer and every seller -- has perfect information about all the involved commodities, an assumption which is false in the market for lemons (in particular, while every seller knows if they are selling limes or lemons, the buyer does not).

**Application to security**

In the previous example, the main lack of information was in the buyers not knowing the reliability of the car. In the context of security, we can treat the unknown information as the security of the product. Both the buyer and the seller are unclear on the exact security of the product as there could be vulnerabilities that haven't been discovered by anyone yet. However, as the creator of the product, it is reasonable to assume that the seller has a significantly better idea of the safeguards as well as security flaws that a given product has.

A potential danger analogous to the market of lemons that happens is if software with poor security sells for the same amount of money as software where developer invested lots of time and effort into heightening the security, it removes incentive for developers to invest that time in future product, resulting in a situation where the market, the world's computer and the Internet are filled with insecure software.

One proposed solution to this potential danger (both in the economic setting and in the security setting) is to have a trusted third party, such as an organization comprised of experts that evaluates the quality of the product. For example, a third party evaluating used cars could classify each car entering the market as "lime" or "lemon", and in the best case this has the potential to change the market to one with perfect information where conventional economic wisdom says the equilibrium behavior should be a desirable outcome. Even if this ideal outcome of perfect information is not achieved (i.e. there could still be security vulnerabilities that a third party misses), it certainly has the effect of decreasing the amount of uncertainty that a buyer/customer has about a car/piece of software. An example of this is: in the "market for lemons" case study, when a buyer encountered a car, their distribution on whether the car was a lime or lemon was uniform -- irrespective of the true value of the car, in the buyer's judgment with probability 0.5 the value was $9000 and with probability 0.5 the value was $5000. However, given ratings by a third party a buyer might see the value of a lemon as between $5000 and $5500 and the value of a lime as between $8500 and $9500, which is much closer to the truth, and the effects this would have on the market for limes is not nearly as drastic as in the zero information setting discussed earlier.

## Miscellaneous economics concepts

**Signalling.** Another interesting idea from economics is that of "signalling". A canonical example is in the majestic tails of peacocks. Though it serves no functional purpose, it consumes a lot of energy and hence can be a signal of good health when looking for mats. In particular, healthy peacocks are more likely to have excess energy to invest in functions beyond what is necessary for survival and they signal that via a colorful and beautiful tail.

**Externalities.** An externality refers to any adverse effect on the world that is caused by the production of commodities, but which neither the buyer nor the seller pays an explicit price for. A classic example of an externality is that of pollution and its side effects in terms of atmospheric global warming caused by the production of electricity via burning fossil fuels. It is an overall negative consequence for the world but the price isn't paid by any particular individuals or organizations involved in the production or consumption of electricty. An analog in the security world is that of software vulnerabilities having a negative consequence for people using systems which have them. For example, a large company can offer a software service such as an email hosting service to customers, but due to security vulnerabilities customers reliant on this email hosting service have their data breached by hackers, and face negative consequences. However, these negative consequences of security vulnerabilities are not borne by the email service provider.

A standard way to deal with externalities from economics is to explicitly identify externalities and have an agency such as the government place taxes for suppliers and consumers causing externalities. Once again, a concrete example here is that of taxing energy companies for the pollution they cause. Another idea is that of the "least cost avoider principle". Given a situation where there is an externality caused, the least cost avoider principle purports that the solution is to have the party that would incur minimum cost in correcting the externality to bear it. For example, in the case of a factory causing pollution affecting residents that live nearby, two possible options are for the factory to install equipment that reduces pollution and another option is for people living near the factory to move further away. And the least cost avoider principle suggests that a legal system should enforce that the option one should take for correcting the externality is to choose the one which costs minimum.

An example in security for the least cost avoiding principle is that of who is liable if there is a bank fraud. Different policies can result in varying behavior and incentive structure. If the laws say that the banks are liable (as they are in America) then that would result in banks investing more in software infrastructure for protecting against security vulnerabilities. On the other hand, if the liability falls on the customer (as it does in Europe), this would result in customers upgrading their own security measures, such as using more effective passwords, and installing antivirus software.

**Case study: security of Windows vs Mac.** There is a paper that reports that Windows machines have their security compromised four times more often than Mac machines. In order to understand why such a phenomenon occurs, consider the following thought experiment. There are two players, an attacker and a defender. An attacker can develop either a Windows malware or a Mac malware. And the defender can setup a Windows defense or a Mac defense. Let's assume that if the defender creates a defense for a malware intended for one OS, it certainly stops the attack. However, if the defense is created for the

OS different from the one the attacker chose to attack, every machine running the OS the attacker targeted gets compromised. Now let's assume we are in a setting where 90% of the machines are Windows machines and 10% of the machines are Mac machines. Then the equilibrium behavior would involve the attacker choosing to make Windows malware with probability 90% and Mac malware with probability 10%, and the defender would make a Windows defense with probability 90% and Mac defense with probability 10%. When this happens, the expected number of Windows and Mac machines that get compromised in the long run would be equal. This means this simple model does not explain the phenomenon observed in the paper. Which means an alternate model that takes into account the success probability of an attack when a defender chooses to defend against Mac or Windows malware is necessary, which suggests that Macs might indeed be more secure than Windows machines.