The Market for Silver Bullets

Ian Grigg
Systemics, Inc.

2nd March 2008

Abstract: What is security?

As an economic "good" security is now recognised as being one for which our knowledge is poor. As with safety goods, events of utility tend to be destructive, yet unlike safety goods, the performance of the good is very hard to test. The roles of participants are complicated by the inclusion of agressive attackers, and buyers and sellers that interchange.

This essay hypothesises that security is a good with insufficient information, and rejects the assumption that security fits in the market for goods with asymmetric information. Security can be viewed as a market where neither buyer nor seller has sufficient information to be able to make a rational buying decision. Drawing heavily from Michael Spence's "Job Market Signaling," these characteristics lead to the arisal of a market in *silver bullets* as participants *herd* in search of *best practices*, a common set of goods that arises more to reduce the costs of externalities rather than achieve benefits in security itself.

Introduction

In an investigation into security, Adam Shostack posed the question, *what are good signals in the market for security* [1] [2]? In addressing this apparently clear question we find ourselves drawn to the question of *what is security?* One avenue of potential investigation is to ask what the science of economics can provide in answer to this question. In economics terms, security could be a "good" as it is demanded and traded for value. This essay seeks to cast security as a good, and attempts to classify what sort of good it is?

In so doing, the results may be of interest to economists as well as security professionals. Investigation along the lines of asymmetric goods would suggest that the market for security would be an inefficient one and evidence so far has not disconfirmed this. This essay contrasts the security good against Akerlof's market for lemons, Rothschild and Stiglitz's market for insurance and Spence's market for jobs [3] [4] [5]. We are led to the observation that security is a good without sufficient information on *either side of the buy/sell divide*. Such a good would seem to defy rationality, as without information, how do purchasers select, and how do sellers present to market? Spence's model for education and jobs best provides a starting point with the introduction of *signals*, and I extend that model with the additional exogenous incentives found in the market for security.

Notwithstanding any surmised inefficiency in the market for security, it is a huge market. Billions of dollars are spent annually on information security, and hundreds of billions, if we extend to for example the markets for defence goods or "national security". For example, the price for a missile

defence system was "nearly $8 billion [6]." As nobody has ever fired an ICBM in anger, it remains a challenge of some depth to decide whether and how much to spend on such a security tool: Zero, the asked price of $8 billion, or the price of the late President Ronald Reagan's fabled "Star Wars" programme? Any potential to increase efficiency in such a large market should be viewed optimistically, and the first step to this would be to establish a theory as to its structure and forces.

A note on Structure. This essay can be seen as the application of three essays in economics to the field of security. The essays are the seminal set identified in the award of the Nobel Prize in Economics for 2001 to Professors Akerlof, Spence and Stiglitz. For brevity, where the essays are drawn on directly, the names of the authors will suffice as references.

I proceed as follows. First, I discuss some characteristics of Security as "goods", primarily the appearance of the active attacker, which helps us to claim that the buyer lacks information. This might suggest the market for lemons. Second, I look at the Asymmetry Hypothesis, and reject it, as the seller also lacks information. Third, I arrive at the only place left, the market in insufficient information. This is then developed with the characteristics of security, to propose a market in silver bullets. Finally, I attempt to draw some lessons, albeit few and tentative.

I. Some Characteristics of Security Goods

In this section, I show by several views that the buyer lacks information. This is uncontroversial, and the reader may skim quickly or skip to the next section.

A Simple Mathematical Model

A security good is purchased by a party or actor that is facing a threat by an attacker. A threat is a costly event that has some small probability of occurring. The probabilities can vary over time and over circumstances. An alternate, complimentary good would be *insurance* which would compensate for costs after the fact (insurance is not considered further in this essay).

The seller of the good makes a claim of some as yet uncertain pedigree that it can defeat or defer the threat, either wholly or partly. As such, the good may *reduce the costs incurred* in the eventuation of the threat, and it may *reduce the probability of the event*.

Some Equations. We can model this with some simple mathematical equations, primarily so that we can question the predictions of the model below (to which the reader is encouraged to skip). The threat can be seen as a series of events each of which incurs an average cost C with a probability of p, giving an expected loss of ($C\ p$). A security good claims to reduce the cost of the threat by S according to some metric, and reduces the probability by s. Given a particular event, benefit b is:

$$\begin{aligned} b &= (C\ p) - C\ (1 - S)\ p\ (1 - s) \\ &= (C\ p)\ (1 - (1 - S)\ (1 - s) \\ &= (C\ p)\ (S + s - s\ S) \end{aligned}$$

Or, total benefit B summed over many events.

$$B = \sum b_i$$

A profit accrues to the buyer if she can pay some price less than benefit B. Note that her profit rises fast with *either* cost reduction (S) or probability reduction (s) [7]. This would be seem to be an easy calculation. Companies are well equiped to deal with calculations of much greater complexity (e.g., "Capital Asset Pricing Model" or CAPM) and consumers intiutively deal at this level (do I really need hurricane shutters? when was the last hurricane? do they work anyway?). Especially, if either of S or s can be shown by a supplier then profit should follow for a purchaser.

Questioning the Predictions? Why then do we not observe such calculations in the market for these goods? Likewise, why no effective insurance? The difficulty with security appears to arise from the intractability of determining the paramaters -- the sufficient set of events, and the good's reduction paramaters *S* and *s*. This is in part as we shall see due to several factors: a lack of information, the existance of an active attacker who disdains any statistical boxing, and exogenous factors that override endogenous calculations.

A Practical Example: Burglar Alarms

> *"Today we were unlucky, but remember we only have to be lucky once. You will have to be lucky always."*
> PIRA statement to Thatcher govt after Brighton bomb, 1984 [8].

Consider a burglar alarm. This is a good that can be tested on a basic level according to the supplied instructions. In this sense, it is similar to a protective gate, in that both deliver a good test based on feedback and routine. One opens and goes click on closing; the other sets and resets, with beeping.

The Inadequacy of Testing. But there the comparison ends. Once we have shown the burglar alarm is activated, we still have no effective way of determining that it achieves its nominal goal of reducing burglaries or the cost of them. The threat that is being addressed cannot be easily simulated. Pretending to be a burglar is not an efficient option through mismatch of knowledge, so simulated tests are ruled out. In contrast our gate is easily tested, we can kick a gate but not an alarm.

A more realistic test is equally problematic. Hiring a burglar is no easy task, as even if one were to achieve this questionable task, it is hard to ensure that the burglar is unbiased, competent and representative. Would he tell you the truth? Is he adequate to the task? The results of any such test would not be strongly indicative of a fair unbiased event of the same characteristics.

Active Participation by the Attacker. It gets worse. In the business of security, the attacker is not a party to our testing procedure. Even though the attacker plays within the game, he is an *active party*; he is not an unbiased, rules-based agent that permits himself to follow statistical patterns. He deliberately attempts to pervert our security, and intends to cause the costs that we are hoping to avoid. As such, he is unimpressed with our efforts and seeks the gaps in them. Any test by one burglar will not find all those gaps, so even a real burglary is not a good predictor of any other event of distinct characteristics.

Indeed our attacker is incentivised to not play by any rules that have been set by the purchasing party. Standardised tests derive from and feedback into standardised models of security, and attackers evolve and migrate their threats to work outside those models. Underwriters' Laboratory is therefore not a good model for security with active attackers.

The active participation in the security process by the attacker might indicate that *game theory* would be a useful direction. The buyer / owner of the good employs it, and the attacker responds. Yet there is often a limited scope for active responses to attacks. That is, many security goods are employed in an environment where there is limited ability to hit back with a *tit for tat* or similar strategy [9].

Summary of Characteristics

> *"You're proposing to build a box with a light on top of it. The light is supposed to go off when you carry the box into a room that has a Unicorn in it. How do you show that it works?"*
> Gene Spafford [10].

Let's summarise. We have these characteristics in the market for security:

- a test of the product by a simulated threat is:
  - expensive and/or destructive, and/or
  - the results cannot be relied upon to predict defence against a real threat;

- The attacker is an active actor in the process:
  - profits by success, and thus is more economically aligned
  - bypasses defences and avoids statistical boxing
  - can be relied upon to be dishonest
- a test of the product by a real threat is:
  - difficult to arrange,
  - insufficiently reproduceable to facilitate statistical analysis,
  - could be destructive, and
  - the results cannot be relied upon to predict defence against any *other* real threat.
- A real event is both
  - destructive and costly, and
  - designed to minimise response, neutralising aggressive defence tactics, and
  - versatile, seeks to migrate.
- The result is a negative sum game, in that
  - as long as the attacker's profit exceeds their costs, it is worth doing,
  - the costs to the victim will generally exceed the profits to the attacker,
  - there may be no strong insurance product.

Buyers are facing quite a challenge. For the good described above, I therefore postulate that there is a lack of information to inform good decisions by buyers. Let us call this hypothesis 1:

> *H1. In the market for security goods, buyers are not informed sufficiently to make rational decisions.*

Indeed, *H1* is fundamental to the very problem at hand, rather than being a simple failure.

Are Buyers in the Dark?

How common is this scenario? How frequent is it that goods purporting to deliver security are installed according to those claims, but the purchaser lacks any confidence or objective way of testing that claim? Anecdotal evidence suggests that much of the market for security goods in the information technology sector is marked by such uncertainty:

> *“ "I go to security conferences where we all sit around puzzling about what kind of metrics to use for measuring the results of security programs," says Adam Stone, an analyst who specializes in security management for the financial services industry. "The metrics we have right now -- the ones we use for assessing vulnerability and measuring the effectiveness of our investments -- are all based on subjective judgments. They're fundamentally flawed [11]." ”*

Circumstantially, Simson Garfinkel suggests a concentration on "security tools when greater benefits can be achieved ... [using] controls and process improvements." It is easier to change tools than to change ways which would indicate a choice of simplicity over security [12]. Ross Anderson concurs that purchasing is distracted from security [13]:

> *“ ... managers often buy products and services which they know to be suboptimal or even defective, but which are from big name suppliers. This is known to minimize the likelihood of getting fired when things go wrong. Corporate lawyers don't condemn this as fraud, but praise it as "due diligence". ”*

II. The Case for Asymmetry

For the rest of this essay *H1* is assumed. This suggests a presumption that sellers therefore have sufficient information, and we then have an asymmetry in information. Recent literature in information security accepts the case for asymmetry, or, the assumption has yet to be disconfirmed [14] [15] [16]. Let us call that the Asymmetric Hypothesis:

> *HA. The market for security goods is a market for lemons.*

In this section, we look at *HA*, testing security within the market for lemons, and reject the hypothesis.

"The Market for 'Lemons'"

In "The Market for 'Lemons,' " George A. Akerlof introduced with the question of why there is a "large price difference between new cars and those which have just left the showroom." His analysis was based on four kinds of car, being, new and used cars, intersecting with good and bad cars, the last being lemons (a slang term for bad cars). A buyer of a new car does not know whether his new car is good or bad, but, Akerlof writes:

> *"After owning a specific car, however, for a length of time, the car owner can form a good idea of the quality of this machine; i.e., the owner assigns a new probability to the event that his car is a lemon. This estimate is more accurate than the original estimate. An asymmetry in available information has developed: for the sellers now have more knowledge about the quality of a car than the buyers. But good cars and bad cars must still sell at the same price -- since it is impossible for a buyer to tell the difference between a good car and a bad car. It is apparent that a used car cannot have the same valuation as a new car -- if it did have the same valuation, it would clearly be advantageous to trade a lemon at the price of new car, and buy another new car, at a higher probability q of being good and a lower probability of being bad. Thus the owner of a good machine must be locked in. Not only is it true that he cannot receive the true value of his car, but he cannot even obtain the expected value of a new car."*

A Market for 'Limes'? For concrete counter-example, it is now time to turn to a prediction of Akerlof's: the market for insurance. In "Markets: an Essay on the Economics of Imperfect Information," Michael Rothschild and Joseph Stiglitz further examined this competitive market as one "in which the characteristics of the commodities exchanged are not fully known to at least one of the parties." Rothschild & Stiglitz claimed to show that

> *"not only may a competitive equilibrium not exist, but when equilibria exist, they may have strange properties. In the insurance market,... sales offers ... do not specify a price at which customers can buy all the insurance that they want, but instead consist of a price and a quantity -- a particular amount of insurance that the individual can buy at that price. Furthermore, if individuals were willing or able to reveal their information, everybody could be made better off. By their very being, high-risk individuals cause an externality: the low-risk individuals are worse off than they would be in the absence of the high-risk individuals."*

Echoes of '*Lemons*' indeed, but they make "the bald assumption that individuals know their accident probabilities, while companies do not." That is, buyers and sellers are reversed in their asymmetry, and thus, this market is the mirror image of Akerlof's model. I will exploit this reversal of fortunes in the *market for limes* later.

How to Slice a Lemon. If, as Akerlof quiped of his paper's original treatment, no equilibrium need exist, "then no goods could be traded .... [and] economics would be different [17]." Of course, used cars do indeed sell. Akerlof predicts the arisal of institutions: guarantees, brand-names, chains, licensing practices, and certification to meet the needs of buyers and sellers. For example,

Underwriters' Laboratory, a US institution owned by insurers, tests many mechanical and electrical products. Rothschild & Siglitz predict distinct institutions: products that encourage buyers to reveal their advantage and grouping to create self-selection of the low risks.

Does the Seller have Superior Information?

> *"1. Good enough is good enough. 2. Good enough always beats perfect. 3. The really hard part is determining what is good enough."*
> Ravi Sandhu [18].

Having laid out the basis of '*Lemons*', we now return to *HA*, that the sellers of security goods have an advantage in information. Recent literature may have assumed that identifying a shortage of buyer information permits the conclusion that the market is one of asymmetric information, yet this needs to be tested. What might this advantage for the seller be? Ravi Sandhu answers honestly but with evident frustration [19]:

> *"We are completely clueless about what is good enough. That is the rub. Business people cannot tell us because they don't understand security, and security people cannot tell us because they don't understand business."*

This *cluelessness* appears confirmed by marketing claims of high form and low substance. Much documentation is produced to give apparent mass to the marketing proposition. Yet actually reading the documentation leads the eyesore analyst to contracts and disclaimers that suggest the reverse conclusion, that they seek to neutralise all promises of fitness for purpose: "This product performs as documented and may not be suitable for your purpose [20]." Or, as Scott Berinato put it [21]:

> *"It's just the logical extension of marketing's dominance over IT in the first place. Long ago, in an era called the dotcom boom, marketing finally neutered information security. Vendors promised "solutions" to Kool-Aid-drinking marketing veeps. Those veeps in turn promised to alchemize revenue out of consumers' private information. Go, said the CEO. Buy these technologies, collect this data and we shall dominate and our stock prices will soar."*

If marketing dominates, what need is there for extensive knowledge of security? The literature is not forceful on this point. Anderson suggests [22]:

> *"The real driving forces behind security system design usually have nothing to do with such altruistic goals [as security]. They are much more likely to be the desire to grab a monopoly, to charge different prices to different users for essentially the same service, and to dump risk. Often this is perfectly rational."*

Rejecting Asymmetric Information

It is a failure of logic to suggest that the buyer's lack of information means that the seller has that information. If sellers are indeed peddling goods labelled as security goods for purposes more subtle and strategic, there is no necessary need for them to understand any more of security than buyers. Indeed, there appears no good way to test whether suppliers have any advantage in security information (hence the difficult question of *signalling* that raised this essay).

Further, with a view to the picture drawn above by Berinato, who are the buyers and who are the sellers? In Akerlof's used car market this was clear. The owner of the car was the seller, and the person with the cash was the buyer. Even though a seller could become a buyer, these acts were sufficiently separated so as to not confuse. Similarly in the market for insurance, Rothschild & Stiglitz's "bald assumption" clearly separated the parties.

I suggest then that roles played by actors around such goods are traditionally fixed in markets in asymmetric information. In security there are a wide and confusing range of alternatives, some substitutable and some maybe not. Further, there is a rich and complex supply chain; one who is a

buyer today is likely a seller tomorrow. When a manager proposes a big company's product to the board, is she buying or selling, and what is her reward in either analysis? The key assumption of fixed roles, simple products and straight forward incentives limit the easy applicability of the theory of asymmetric information, at least to the extent of giving the buyer a free pass in liability.

I therefore reject the Asymmetric Hypothesis, and suggest that the seller has no advantage over the buyer. Further, the seller lacks sufficient information to make a credible selling decision. Call this Hypothesis 2:

> *H2. In the market for security goods, the seller lacks sufficient information to make a credible selling decision.*

III. Markets in Insufficient Information

Then, security is a good for which *there is insufficient information available on the parts of both the buyer and the seller*. Security is not a good with asymmetric information, in the sense of used cars or insurance. It is more a good exhibiting knowledge that is too expensive for either party to obtain. The efforts of either party could see them knowing more than the other, but even reasonable efforts leave parties without sufficient information to make a rational decision.

"Job Market Signaling"

To investigate this space, I now turn to the third of the triumvurate of essays celebrated by the 2001 Nobel Prize, Michael Spence's essay “Job Market Signaling.” This essay on information poverty is rich in insight, and I take the unusual step of fully summarising the essay [23]. The author's own words are best [24]:

> *“ In most job markets the employer is not sure of the productive capabilities of an individual at the time he hires him. Nor will this information necessarily become available to the employer immediately after hiring. The job may take time to learn. Often specific training is required. And there may be a contract period within which no recontracting is allowed. The fact that it takes time to learn an individual's productive capabilities means that hiring is an investment decision. The fact that these capabilities are not known beforehand makes the decision one under uncertainty.*
>
> *To hire someone, then, is frequently to purchase a lottery. ”*

This market raises a particular problem for honest negotiation: "if the incentives for veracity in reporting anything by means of a conventional signaling code are weak, then one must look for other means by which information transfers take place." How then do two parties, employee and employer, neither of whom have sufficient information to rationally trade, move forward and conclude negotiations?

Spence viewed formal education certifications (or qualifications or degrees) as *signals*, being proxies for the more normal claims made in efficient markets [25]. Presuming some verifiability, I can improve my prospects in the job market by acquiring some educational qualifications [26]. Yet, these cost, in time and in money, and in principle, anyone can participate. Then, "it is not difficult to see that a signal will not effectively distinguish one applicant from another, unless the costs of signaling are negatively correlated with productive capability."

Hence we arise at one of Spence's key insights: it is not the education provided that makes the market for education work, but the selection that results. By separating students into two groups, one that finds it hard to do a course, and one that finds it easy, this is a sufficient *signal* for the employer to be taken as the proxy for productivity. That is, education does not (only?) improve productivity. Its (more?) important purpose is to present a barrier against low productivity; the optimisation problem for the employee is solved by securing the education that, *for the most ease*, generates the best wages, and is solved for the employer by selecting those who most easily pass the purported

productivity test [27].

The foregoing is simply setting up for the next step:

> *" At this point it is perhaps clear that there is informational feedback to the employer over time. As new market information comes in to the employer through hiring and subsequent observation of productive capabilities as they relate to signals, the employer's conditional probabilistic beliefs are adjusted, and a new round starts. The wage schedule facing the new entrants in the market generally differs from that facing the previous group. "*

Spence then spends the bulk of his essay in analysing the stability of the model generated by the above feedback loop.

> *" An equilibrium is defined in the context of a feedback loop, in which employer expectations lead to offered wages to various levels of education, which in turn lead to investment in education by individuals. After hiring, the discovery of the actual relationships between education and productivity in the sample leads to revised expectations or beliefs. Here the cycle starts again. An equilibrium is best thought of as a set of beliefs that are confirmed or at least not contradicted by the new data at the end of the loop just described. Such beliefs will tend to persist over time as new entrants into the market flow through.*
>
> *Multiple equilibria are a distinct possibility. Some may be Pareto inferior to others. Private and social returns to education diverge. Sometimes everyone loses as a result of the existence of signaling. In other situations some gain, while others lose. Systematic overinvestment in education is a distinct possibility because of the element of arbitrariness in the equilibrium configuration of the market. ... everyone is reacting rationally to the market situation. "*

How can we summarise this set of claims? The markets in question are characterised by an inefficiency of information flow between buyer and seller that is then finessed by means of *signals*. Buyer beliefs over the efficacy of the signals are key to a feedback loop, and stability can be reached if the signals are not disconfirmed. Yet, the mechanism is not sufficient to move towards an optimum, nor even to ensure a positive return. Indeed, the only thing that seems likely is that it reaches an equilbiria, and if this results in a postive return to welfare, the reasons are likely external to the model.

It is perhaps apparent by now that the presence of signals would be *bona fide* evidence of an inefficient market, and that signals themselves are, by fortune, inefficient stand-ins for metrics that are presumably efficient yet unavailable. Thus, the question that opened this essay is addressed: there are no good signals, only uncertain signals.

How plausible is Spence's model? I see many real world examples: I have always wondered why the British civil service class read the classics at Oxbridge; we might also observe that the institution of higher pay for men, less for women, has exhibited remarkable stability over time, which Spence directly analyses and confirms. And, more apropos, security-minded readers might mentally substitute *security* wherever *education* is found above and experience a sense of *deja vu*.

The Matrix of Markets in Imperfect Information

Having introduced three distinct inefficient markets, above, it should now help to place these markets in a comparitive matrix. Following Akerlof, when the seller has an information advantage, this is referred to as the *market for lemons*. If the buyer has an information advantage, this is Rothschild & Stiglitz's market in insurance. As this market shares many characteristics with that of lemons, I term it the *market for limes*. As we now have two binary choices, we can map this in a two by two matrix (a favoured tool of economists) giving us two more spaces. See Figure 1.

The Market for Goods, as described by Information and by Party	**Buyer Knows**	**Buyer Lacks**
Seller Knows	Efficient Goods	Lemons (used cars)
Seller Lacks	Limes (Insurance)	**Silver Bullets (Security)**

Figure 1. Security is a Symmetrically Insufficient Market

Of course, one space is efficient goods, where both parties have information. The remaining fourth space is that market where both sides lack sufficient information. I term this space the *market for silver bullets*. A silver bullet is a term of art in the world of software engineering for a product or process that is presented as efficacious without any logical or rational means to back up that claim [28]. Silver bullets are goods traded in markets in *insufficient information*. Along with the diagonal of lemons and limes (*asymmetric information*), they form the markets in *imperfect information*.

Security as a Market in Silver Bullets

I note in Section I that the buyer has no good test with which to confirm the veracity of the seller, and so cannot economically determine *ex ante* that the good meets needs. The seller lacks that information as well, as he has no advantage with the attacker [29]. Relating to the theories of imperfect information, it strikes that the security good does not reveal *to either party* a sufficiency of information needed to make a rational decision.

Following Spence, we can intersect *H1* and *H2* as Hypothesis #3,

> *H3. In the market for silver bullets, neither buyers nor sellers of the good are informed sufficiently to make rational decisions.*

Security goods are such goods of insufficient information. For the rest of this essay, we assume these hypotheses. Given that, we can also make the following claim:

> *H4. Decisions are therefore made primarily on the basis of exogenous factors that do not strongly (causally) relate to the nominal security goal of the good.*

These "exogenous factors" would be the signals of Spence's models.

IV. Rational Trading of Silver Bullets

It is our claim that decisions are made in the market for silver bullets regardless of the nominal needs of security. This section explores how actors buy and sell, and why this is rational.

Extraordinary Costs and Fingerpointing

> *“ New security measures, such as mobile two-factor authentication, will become the norm for businesses and institutions of all sizes sooner rather than later, said Michelle*

> *Warren, a Toronto-based IT analyst with Evans Research. "Strong security is all about avoiding litigation, bad PR [publicity] and fines for non-compliance. It all goes to costs." "*
> Paul Lima [30]. (my emphasis).

Let us turn now to an examination of breaches of security. What happens when security goods are employed in anger? Why is the information not shared in these cases? What can it tell us about security choices that arise?

Ordinary Costs. When a costly event arises due to a realised threat, there are the costs as suggested above: value is transferred from victim to attacker, as well as destruction of value. Let us call these *ordinary costs*, being the direct costs due to the event, that can be modelled in a simple world of victim and attacker. For the most part, I ignore the part of the attacker.

Extraordinary Costs. There are also additional costs which occur, for example, from exposure and adverse publicity [31] [32] [33]. In a competitive market there are many stakeholders: secondary victims (customers of the victim), competing actors (unaffected competitors), future victims (vulnerable competitors) and information channels (news, researchers, suppliers). If we accept many players, and assume that information leaks more or less freely, the costs of adverse publicity can easily outweigh the costs of the direct event itself. Let us call these costs *extraordinary costs* (economists might better term them *externalities* or exogenous costs).

Fallout. Observers in the market know that competitors are also subject in general to the same risks, and are thus subject to future events, and the consequent ordinary costs. Thus, competitors also face the same adverse publicity from the immediate event [34]. That is, competitors do not face the direct costs of the event, but many or all may well be exposed to extraordinary costs of *fallout* from the immediate security breach event that effects one victim. Due to *H3*, fallout costs are indiscriminate and even unaffected competitors incur extraordinary costs.

Fingerpointing. As victims always have an incentive to reduce costs, the competitors can seek to claim that they themselves are not at risk. They can do this by differentiating themselves from the victim. I term this tactic *fingerpointing* to describe the deliberate shifting of attention away from ones own vulnerabilities to another's, in order to reduce own fallout costs.

I do not examine here the merit or otherwise of fingerpointing, only its effects. As secondary victims engage in fingerpointing at the hapless original victim, the costs for the latter are likely raised. At best, this is a zero sum game. For the fingerpointers it reduces their own costs even at the cost of one of their number.

If the extraordinary costs of an event (e.g., the bad publicity from both the event and the fallout) are greater than the direct costs of the event (e.g., the stolen money) then the strategy of defending victims likely switches to reducing the higher extraordinary costs. Preference is given to absorbing direct costs, and security purchasing strategy switches from preserving owner's value to preserving the safety of the purchaser. And thus, notwithstanding the absence of information, decision making in security goods is led away from the nominal goal of security towards the externality of the reduction in fallout costs.

> *H5. The primary motivation for decisions on security goods is purchaser's safety from exposure to fingerpointing by other stakeholders.*

> *H6. The direct costs of security breaches are absorbed by the primary victim without further consideration.*

The public exposure of breaches following the Choicepoint Affair is unprecedented. It has been suggested that the breaches have always been there. Assuming that claim, and given the excessive and ongoing efforts of companies to avoid disclosure, even when their own value is at risk, we find support for the claims that *extraordinary costs in disclosure and adverse publicity are in excess of the direct and ordinary costs of the breach* and therefore *the startegy of defence is to reduce extraordinary costs, and absorb direct costs.*

Herding and the Emergence of a Market for Silver Bullets

Adverse publicity explains the need for hushing up security breaches. It may also explain the popularity of secrecy in security models and reputations of impenetrability, which may shed some light on the popularity of excessive cryptography. A mythology of secrecy and impenetrability, based on likewise impenetrable descriptions of mathematics, cryptoprotocols and patented features, may make it no harder to breach, but it does make it easier to hide the breach when it does occur.

Yet fingerpointing is a first order reaction, and not only is it widely recognised as unprofessional in sophisticated markets, the economics of the Prisoner's Dilemma suggest a less efficient market. What then can participants do to limit the costs of fingerpointing? What happens when ones risk reduction strategy rebounds and also raises costs?

Consider a market with 3 goods *A*, *B*, *C* and 3 potential victims *a*, *b*, *c*. A market of participants that are engaged in distracting attention from their own risk is initially unstable. If each employs one good (e.g., the namesake good), then each remains at risk of fallout, as I suggest that no good predicts in advance that it alleviates all risks. Thus, even if good *A* protects, this is unknown, and employing victim *a* remains, at the minimum, with significant costs due to the risk of fallout.

If *a* employs also *B*, there is still the risk inherent in fallout from *C*. Other things being equal, only if victim *a* employs also *B* and *C* is the risk of fallout alleviated. There is then no longer the means for *b* and *c* to offset their fallout costs at *a*'s expense, as the same goods are in place. Differentiation by *b* and *c* is strictly weaker. (Other things not being equal, *a* could band with *b* and attempt to ridicule *c* into dropping good *C*.)

If the sum cost of a set of security goods is less than the costs of fallout, there exists a natural equilibrium where all participants employ the same set of goods, however they are arrived at. Equilibrium obtains when all participants employ a standard set of goods. If a participant lacks a good, there is an incentive of reduced fallout to add it, until all goods in a set are present. Adding a missing good both reduces own fallout costs and reduces others' potential rewards for fingerpointing.

Once the set is selected, participants face negative incentives to differentiate. Further, if any one participant introduces a new security good, each other participant has a choice of following along, or of attacking the good. This raises the costs of each new measure as it has to face the combined scrutiny of other participants, and all participants have to invest.

Once the set is achieved, the group-wide investment cost dominates. This would oppose an unknown incentive gained from improve security with a new technique. A likely result is for the new technique to be ridiculed by all, rather than all accepting increased costs. This high barrier is highly rational.

As with Spence's market for education, there are many equilibria in the market for security goods. Likewise the set of equilibria is robust for different numbers of players and different numbers of goods; it takes a good that promises sufficient benefit to cover the investment costs of all parties as well as the coordination costs in order to breach an equilibrium. Unsurprisingly, once a good is entered into the set, there is little incentive to remove it even when widely criticised.

Hence changes to the selection of security measures happen slowly, and when they do happen, they tend to ripple fast across the community. This characteristic of security in communities is known to economists as *herding*. The goods in the set are *silver bullets*. The set of all such goods in a given community is known as *best practices*. The inclusion of a silver bullet is based on the random chance of being used by a participant in beginning rounds, and surviving the possibility of being isolated and withdrawn. Although this earliest process may be motivated by security thinking, that motivation shrinks dramatically once the set is chosen; deviations are costly. Changes are more likely to relate to indirect effects of the costs and benefits of maintaining the community's set than be to be connected to a good's nominal mission of security.

In summary, where indirect costs dominate direct costs, participants are encourage to reduce the risk of differentiation. Silver bullets are chosen and trialled randomly. In order to reduce the costs or fingerpointing, a set of best practices emerges, and all participants conform to the set.

Less Security Results

> *"Best practices look at what everyone else is doing, crunch numbers-and come up with what everyone else is doing. Using the same method, one would conclude that best practices for nutrition mandates a diet high in fat, cholesterol and sugar, with the average male being 35 pounds overweight."*
> Ben Rothke [35].

I have shown that any deviation from best practices is costly, *including* towards a presumed direction of greater security. Only the most profitable of security measures will produce enough benefit to overcome the cost of breaching the equilibrium. As the cost of breaching the equilibrium is proportional to the number of community members, the larger the community, the greater the opportunities for security are foregone, and the more the vulnerability.

Herding is a Nash equilibrium as well as being rational behaviour [36]; if one player chooses a new silver bullet, other players do not have a better strategy than sticking to the set of best practices, and even the player that changes is strictly worse off as they invite extraordinary costs in the event of a breach. This approach lowers the more significant extraordinary costs and accepts direct costs as unavoidable and to be absorbed.

V. Insights into Insufficiency

Modelling security as an economics good within the market for insufficient information is a useful step to designing methods to reduce the inefficiency in the market for security. This section looks at what we have learnt from the above analysis and what we can suggest to improve the market for security. There are in essence these strategies:

- *a better best practices* — Break the current equilibrium and shift to a more secure equilibrium.
- *Break Free of the Loop* — Discard best practices entirely and move security back onto the agenda.
- *migrate to Lemons or Limes* — Improve information and give one side or other an advantage.

A Better Best Practices

Selecting a better best practices may result in a theoretical improvement, but as a strategy it is clearly flawed. Firstly, given our assumptions, we have little confidence in our ability to pick a new equilibrium. Secondly, our underlying security needs drift faster than our ability to select and implement a new equilibrium. For example, banking regulators in the USA started insisting on "two-factor" security in late 2004 [37]. Within one year, and *before the industry had moved or even agreed* the breach was spotted: *Man-in-the-browser* attacks [38].

To take a turn through John Boyd's OODA loop, the attacker manoeuvres inside our best practices loop [39]. Thus, the presence of the active attacker provides for a strictly worse result than Spence imagined: in the education market, the correlation between signal and metric drifted without sufficient feedback; in the security market, the correlation diverges deliberately to a net negative result. The market for silver bullets naturally converges on a net loss of welfare. Therefore, recommending an improvement to best practices is to signal your surrender to the enemy, confirm your intentions to fight again on his battleground, and have no better strategy than to prepare for his next victory.

Breaking Free of the Feedback Loop

It is clearly superior to attack the forces of equilibrium, and to break the loop entirely. As the feedback loop breaks, and equilibrium absents itself, differentiation would result. Competition and experimentation would result in improved information, setting the stage for migration to the market for lemons.

Replace the signals. In principle, we could replace the signals with the underlying metrics. In education, the elusive metric was productivity. In Security, this is an open area of research, so we can say little conclusively except to underscore the importance of research in security metrics. This approach would move each agent's set of goods away from the set of best practices towards focussing on more precise metrics. Their precision would more clearly relate to the individual circumstances of each agent, and thus force more local alignment and greater sector diversification.

Rebalance the costs. If the extraordinary costs can be reduced, especially to below the ordinary costs, this would assist to bring the ordinary costs back into primary focus. By rebalancing the costs, attention of security buyers could be naturally switched back to reducing the primary damage.

- Reduce fingerpointing. This could be done by institutional methods (associations, professionalism, codes of conduct, etc), and would seem to be a suitable avenue for security associations and regulators to investigate. Fingerpointing was traditionally reduced by means of secrecy of breaches, but this approach is out of favour due to its overall reduction of information. It is a singular observation that to the extent that *breach disclosure* (SB1386 and friends, discussed below) increases fingerpointing, it strengthens the Nash equilibrium and may make matters worse; this places a strong challenge before the professionalism of the industry.
- Make direct costs more direct. Many of the ordinary costs are lowered by various tricks. The goal is to direct these costs to the party that can best deal with the cost, by way of modifying the security. Often, the larger party creates the system but transfers the direct costs of breach onto end-users. Canonical case studies here would include online banking and phishing, or credit card fraud fees paid by merchants. As a thought experiment, insist on an audited estimate of ordinary costs, and that amount could be paid out in cash to the end-victims, or a third party fund where victims are not found. A further tool to focus attention is to mandate multiples such as double liability.
- Sellers to share in the pain. Calls for software liability are unpopular, but there is no reason why it need be mandated. Compare Uninteruptible Power Supplies (UPSs) with anti-virus software; the former comes with up to $100,000 of coverage, written on the box, whereas the latter dislaims all. What precisely could be done here is open to innovation. Would for example a security supplier be present at a press conference to disclose an SB1386 breach?

De-certify use of best practices. A further direction is to reduce or reverse the tendency to promote best practices. This could be best done by institutions (associations or regulators) that celebrate daring experiments and differentiation rather than conformance and fear. Institutions would serve better by avoiding best practices, facilitating the open sharing of information, and insisting that members of the flock find their own paths.

Sunlight is the Best Disinfectant. By removing secrecy from practice, the feedback loop is exposed to all, allowing rational examination and deconstruction. Secrecy is the breeding ground of security's equilibrium of insecurity. The process in brief works like this. Each player keeps their security model secret (for a variety of rationales). Then, regulators breach the secrecy of each player and learn the model. They aggregate this knowledge into common wisdom and then feed that wisdom back to each player as common frameworks. best practices is thus forced through the regulatory channel wherever secrecy permits the channel to emerge and dominate as a centralised filter of knowledge [40]. The more a sector is regulated, and the more secrecy is favoured, the more the equilibrium is locked into the beliefs of the central power. Breaking down the secrecy reduces the power of the regulator to impose best practices equilibria, because all can see, challenge and diverge from the models of others, and the beliefs of the regulator.

Migrating to Lemons

The matrix of information in Figure 1 claims that security goods are found in a square of insufficient information; the goods could benefit by reaching sufficient information for either the buyer or seller. A goal then is to improve the information available on the good, for either or both buyer and seller. Then, the good is encouraged to migrate into the domain of asymmetric information (lemons or limes).

Sharing Information. Better information can be had by sharing and building. Schechter and Smith suggest [41]:

> *"Sharing of information is also key to keeping marginal risk high. If the body of knowledge of each member of the defense grows with the number of targets attacked, so will the marginal risk of attack. If organizations do not share information, the body of knowledge of each one will be constant and will not affect marginal risk."*

A fruitful direction for new information might be projects such as Honeypots and Honeynets [42].

Changing the Payoffs. Sharing, as above, suffers from a Prisoner's Dilemma: it raises costs for the sharer, and the benefits are not accrued to the sharer. There may well be a higher payoff if all victims share their experiences, yet those that keep mum will benefit and not lose more from sharing. As all potential sharers are joined in an equilibrium of secrecy, perhaps in fear of fingerpointing, little sharing of security information is seen, and this is rational.

An insight of the asymmetric and game theory literature is to change the payoffs so as to incentivise the more informed player to reveal information. Until recently, security breaches were indeed generally and widely hushed up. The Prisoner's Dilemma may have been "solved" by a California law, SB1386, brought into effect mid 2002. The law forces companies to notify victims of losses of identity information and has had the effect of springing a mandated leak in the secrecy of breaches. At first a steady trickle of smaller breaches garnered minor press attention. Then Choicepoint burst into the public consciousness in February of 2005, due to several factors. This major breach caused not only a major dip in the company's share price (c.f., fingerpointing), but also a wave of similar revelations within the month [43]. The threat of civil action by Attornies General and by victims was sufficient to spread the effect of California's law to the entire USA and beyond.

Such public exposure of breaches is unprecedented. Either we have just observed a breaches "shock" *or* the breaches are normal, but the disclosure is abnormal. Anecdotal evidence in the security industry supports the latter theory; therefore I suggest that the payoffs in the Prisoners' Dilemma of security disclosure have been adjusted by SB1386.

Frequency of Events. Spence directly warns that the "relative infrequency of appearance in the market which defines the class signaling phenomena under scrutiny here, is not characteristic of many markets, like those for consumer durables..." I can then suggest that the infrequency of events in the security market — breaches — works against it, such events are not frequent enough to provide a base for decision. Following Spence, when the threat is low enough to give few or no events, a silver bullet will not disconfirm its efficacy. In contrast, when a threat is continuous enough to generate statistically meaningful results, it becomes routine, thus migrating out of the domain of information insufficiency: silver bullets are shown to be what they are.

This could be posed as the *Barings-Visa paradox*. Barings had one breach, and failed. To that moment, it was totally secure. Visa has millions of breaches every year, pays some small percentage, and is comparitively solid.

Incentivising the Breach. In markets with insufficient information, there is no information advantage on the other side, but there is a third party, the attacker. Can we incentivise the attacker into revealing information? Consider paying for successful attacks, as has been employed under the guise of *crack challenges* and *bug bounties*, and more lately *exploit markets* [44].

Perversely, I suggest that organisations not "try too hard" to push attacks so far away that they can no longer be seen. Incentivising managers to bury events and probabilities for fear of repercussions will fundamentally reduce information and will also raise a *false sense of security*. Seen in this light, laws to criminalise activities such as cracking, reverse engineering and academic research into security products will probably work to (further) reduce information in security. Following the model of information insufficiency, these laws will worsen security by further embedding the equilibrium.

Asymmetric Insights

I placed security as a good in markets with insufficient information, yet some observations on asymmetric information markets are merited. Firstly, once migrated, security goods could benefit from the wide range of possibilities suggested by the asymmetric literature: branding, standards, testing associations, regulation, sharing of threats, contracts, self-selection. Secondly, no good is so cleanly placed as to be exclusively explained in one model. Other lessons may already apply. Thirdly, we can now see that many of the proposals in the security literature are better explained within the scope of Spence's model, not that of asymmetric information. Which is to say that we are already part the way there. Finally, the perfect goal is to move goods to the efficient market space, and *not to trap the goods in an asymmetric space* by reserving the information for one side or the other.

VI. Conclusions

Summary

Security goods, when they exhibit poor testability and the presence or perception of active and aggressive third party attackers, place themselves in a very difficult space. By lack of approachable metrics of quality, buyers lack sufficient information to support a purchasing decision. Likewise, sellers are stymied by the attacker's refusal to hold to theoretical and statistical models. It is likely that at least in some security markets the seller also lacks sufficient information. Asymmetric prescriptions of information sharing will be inadequate as there is none to share, and will likely raise institutional, signalling and screening costs that make matters worse.

I hypothesise security is a good in the market of insufficient information (*H3*), where buyers and sellers alike lack the information to support decisions. Following on from Spence's market in education, the critical beliefs of buyers are confirmed by the presence of *signals* that are not close to security (*H4*). Those that are not disconfirmed by new data arriving in cycles generate an equilibrium which exorts sellers to provide (at cost) those signals. I refer to goods in this market as *silver bullets*.

In an initial market, there are no signals. Participants find that any purchase of security goods sets them at odds with their peers; a natural equilibrium arises where all participants use the same set to avoid the extraordinary costs of *fingerpointing*. The result can be termed *best practices*, as a set of silver bullets. Once equilibrium obtains, the set cannot shift without excessive costs. Minor innovations are blocked, and major innovations must deliver rewards to all participants at once. The presence of the active attacker ensures that best practices converges to a loss in welfare over time, making it strictly worse than Spence's model.

Predictions and Criticisms

The hypotheses presented should be tested. The possibility of such an economic good certainly exists, and Spence's work has stood the test of time. In addition to the comments in this essay, confirming observations would include:

- the presence of signals, the lack of metrics,
- the frequent arisal of ineffective and costly institutions,
- the difficulty of calculating ROI for security projects,

- the arisal of new security threats such as phishing, and
- the inability to stem the losses nor formulate a compelling picture.

Testing the hypotheses could be a challenge, perhaps predicted by the minimal developments within the market for education and jobs in the 30 years since the publication of Spence's work.

Where's the Lemon? To disprove *H3* for a particular good, simply surface the information that would place it as a lemon or lime, and show that if that information were placed in the hands of the other side, the market would then reach efficiency [45].

Loading the silver bullet. Research into decision making in security goods trading would be fruitful. What are the forces that lead to the purchase being needed? *H5* suggests that avoidance of fingerpointing is important. What are the forces that effect the decision that is made? How are the costs distributed? As per *H6*, does the buyer end up with costs, and the seller is nowhere to be found? Is herding present, and does it effect decisions? What forces work within herding [46]?

Vampire? What vampire? In closing, it has to recalled that there is only limited traction in any model or test that does not involve all the parties. Especially, the active attacker is absent in most economics, including that of imperfect information. How much point is there in debating endlessly the strained relationship between buyers and sellers, and ignoring the critical part played by the attacker? The imperfect information literature broadly assumes markets with two actors, whereas the security market has three. How much changes with an economics of three-party "trading" ?

References

[1] Adam Shostack, "Avoiding Liability: An Alternative Route to More Secure Products," working paper (*FC++*), 2005. Also see earlier notes on Emergent Chaos blog, " Ratty Signals" 01 Jan 2005.
[2] Ian Grigg, " Security Signalling - the market for Lemmings" Financial Cryptography blog, 02 Jan 2005.
[3] George A. Akerlof, "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism," *Quarterly Journal of Economics*, v84 (1970) pp488.
[4] Michael Rothschild and Joseph Stiglitz, "Equilibrium in Competitive Insurance Markets: an Essay on the Economics of Imperfect Information," *Quarterly Journal of Economics*, v90 (1976) pp629.
[5] Michael Spence, "Job Market Signaling," *Quarterly Journal of Economics*, v87 (1973) pp355.
[6] James Hacket, "Missile defense trajectory," The Washington Times, 10th October, 2005. dead link.
[7] We should also discount each event's profit over time *t*. If this were done, it would skew results against the defence, as the purchase price of the good is paid up-front yet the benefit is discounted into the future. As a straightforward result, risky projects with high discount factors will not extract as much benefit from security goods as low risk projects, which might explain the concentration of security sales in stable industries such as banking and government.
[8] BBC, "The IRA campaigns in England," website
[9] Internet security exhibits this when ones attacker is protected behind hacked servers acting as proxies. Guerilla warfare similarly emphasizes the minimisation of potential for response, as with the avoidance of set-piece battles.
[10] Gene Spafford, " Collected analogies," Ed. Mahesh Tripunitara.
[11] Lawrence A. Gordon and Robert Richardson, "InfoSec Economics," *Security Pipeline*, 15th April 2004. CCRA.
[12] Simson Garfinkel, *Design Principles and Patterns for Computer Systems That Are Simultaneously Secure and Usable*, Thesis submitted to MIT, 2005.
[13] Ross Anderson " Why Information Security is Hard - An Economic Perspective," Proceedings of the 17th *Annual Conference on Computer Security Applications* , p.358, December 10-14, 2001
[14] Ross Anderson, *op cit*.
[15] Douglas Barnes, " Deworming the Internet" *Texas Law Review*, Vol. 83, No. 1.
[16] Rainer Böhme, " Vulnerability Markets - What is the economic value of a zero-day exploit? " Proceedings of *22nd Chaos Communication Congress*, Berlin, Germany, December 2005
[17] George A. Akerlof, "Writing the "The Market for 'Lemons'": A Personal and Interpretive Essay," Nobel Prize website. 2003.
[18] Ravi Sandhu, "Good-Enough Security," *IEEE Internet Computing*, Vol.5, No.3, Jan-Feb 2003, pp66-68.
[19] Ravi Sandhu, *ibid*.
[20] Jane K. Winn, " Couriers without Luggage: Negotiable Instruments and Digital Signatures," *49 South Carolina*

Law Review 739 (1998)
[21] Scott Berinato, "Waterloo," Alarmed column, *CSO Online*
[22] Ross Anderson, *op cit.*
[23] Apropos of inefficiencies in the market for economics research.
[24] And, it is perhaps interesting to interpret the liberal use of direct quotations as a signal.
[25] A definition of signals is hard, as Spence did not provide one, perhaps deliberately. The reader is left the puzzle, with some clues.
[26] I might then signal a Bachelor's in Computer Science and an MBA.
[27] It was apparently easy for me to do Comp Sci and Business, yet not so easy to do Economics or Information Security?
[28] Frederick P. Brooks, " No Silver Bullet," 1987, in the 2nd Ed (1995) of Frederick P. Brooks, *The Mythical Man Month*, 1975.
[29] *Prima facea* the seller is at an advantage, as he knows what his good can do. Yet he only *needs to know* enough to sell the good, he does not especially need to know how secure it is. The art and science of selling being well advanced, we are well beyond the point where goods must be sold solely on their merits.
[30] Paul Lima, "ID double-check systems stymie on-line thieves," April 2005, *Globe and Mail*, dead link.
[31] Cavusoglu, "Effect of Internet Security Breach Announcements on Market Value of Breached Firms and Internet Security Developers," *Economics of Information Security* Camp, L. Jean; Lewis, Stephen (Eds.) Springer.
[32] Katherine Campbell, Lawrence A. Gordon, Martin P. Loeb and Lei Zhou, " The economic cost of publicly announced information security breaches: empirical evidence from the stock market," Accounting and Information Assurance, Robert H. Smith School of Business, University of Maryland, 2003. Springer.
[33] Arjen Lenstra and Eric Verheul, " Selecting Cryptographic Key Sizes," *Journal of Cryptology*, 1999, warns of adverse media attention as being a risk from the mere discovery of utilising small key sizes.
[34] This is routinely seen on the stock market. When one company in a sector surprises the industry with adverse news, the entire sector is downgraded, at least until deeper analysis is undertaken.
[35] Ben Rothke " Vince Lombardi: Role Model for CIOs," *EWeek*, 15th December 2003.
[36] John Nash, Nash equilibrium as described in Wikipedia.
[37] FDIC, "Putting an End to Account-Hijacking Identity Theft," website 14th December 2004.
[38] Philipp Gühring, "Concepts against Man-in-the-Browser Attacks," working paper, (*FC++*) June 2006.
[39] John Boyd's *Observe, Orient, Decide, Act* model is described in various sparsely documented presentations to US Department of Defense, e.g., "Patterns of Conflict" 1986, and "Organic Design for Command and Control" 1987.
[40] I use the term *regulator* broadly. For example, the same feedback loop can be observed through compulsory audits.
[41] Stuart E. Schechter and Michael D. Smith " How Much Security is Enough to Stop a Thief?," *Financial Cryptography 2003* LNCS Springer-Verlag.
[42] The Honeynet Project & Research Alliance, " Know your Enemy: Tracking Botnets," March 2005.
[43] Bank of America, LexisNexis, three Universities, Las Vegas' Department of Motor Vehicles, DSW Shoe Warehouse, and Westlaw.
[44] Mozilla Foundation, " Security Bug Bounty Program," 2nd August 2004.
[45] For example, if a survey of security suppliers tested for the most useful knowledge from their competitors revealed a great desire to share others' threats databases, this would challenge the hypothesis. If competitors desired their opponents' customer lists this would tend to support the hypothesis. Alternatively if it were possible to establish independently that some markets were subject to herding, and others not, then the presence and popularity of a security good in each of those markets might be indicative.
[46] Cartel economics are well studied by Hamel & Prahalad, for example. Once a favourite emerges, it tends to gain traction and reward the early investors. This predicts a strong incentive for sellers to invest heavily in up-front perception and marketing, rather than wait for results from the herd. Are there centers of power where those that have the most leverage win? If so, this would predict frequent arisal in new consortia, and rapid uptake of members, as vendors move to establish their leverage whenever a new center is formed.