Blind in the Panopticon

The decreasing marginal utility of access to information as an aspect of cyber policy

The uncertain price of information

There is no clear way, in the present, of deciding what information will be relevant in the future. Although axiomatic, the problem of valuing information based on future utility is at the heart of paradoxes around privacy, intellectual property and secrecy. In general, the difficulty in setting a value (not simply a *price*) for information results in a desire to stockpile-and-hoard all data, resulting in a wide variety of *diminishing returns* on an initial investment in intelligence gathering.

In the field of business, *invention* is one of the key drivers of wealth creation. It has multiple functions, most clearly to change the structure of society as new infrastructures like railways, aeroplanes or digital systems allow new economic and social norms to form around them. For agents of invention, *patent* provides a relatively secure return on the investment in toil and gold required to make something new. Society benefits from patent twice: *first* by encouraging innovation, and *second* by making sure that knowledge is not lost to trade secrecy. The social benefits of knowing how something works *generally* outweigh the social costs of occasional monopolies. In technology, the situation has become increasingly complex as incredibly valuable patents have changed the business strategy around patent - 19 years is much longer in software than it is in mechanical engineering, for example, but generally speaking the equilibrium is as stated.

The *valuation problem* in patents is as follows. A patent has a fixed cost, often estimated at $30,000, largely composed of legal fees. Some contingency is also required to use or defend the patent. An *invention* has uncertain returns, by virtue of being new. Therefore, a gamble is being taken at every patent filing: will the economic return on protecting this invention exceed the cost. As with all gambling, those with deep pockets are favoured over the smaller players because they can tolerate far more loss on the way to a big win, and therefore a *conglomeration pressure* is created. Larger and larger enterprises are created to tolerate the costs of research and patent, and to maximize the capacity to exploit and defend these territories. The recent litigation between Apple and Samsung over key smartphone interface components is not-unlike two scribes engaged in a duel over the precise orthography of English in the first scriptorium. Turning innovation into property, which disputes can then be fought over, may not always accelerate progress or produce substantial social benefit. Patent offices lack the financial muscle to make first class decisions in esoteric areas, leaving the (much more socially expensive) court system to corral the necessary expertise, burning resources in expensive litigation.

The cost of knowing in national security

A substantial subset of the national security apparatus, and the more ordinary policing system, gathers or stores information. These datasets are acquired on a cost/benefit basis, with both financial and social costs considered. A changing technological landscape has dramatically changed the acquisition costs of knowledge, but not necessarily the social cost of this knowledge being gathered.

Consider a notional state in which homosexuality is traditionally illegal, was illegal in the pre-digital period, and is currently illegal. It is likely possible, without undue effort, to

analyse network traffic to estimate the probability that any given digital citizen is breaking the laws on homosexual activity. Let us consider four ways this could be ascertained:

1. Prima facie illegal activity like viewing homosexual pornography
2. Use of predominantly gay social networking tools and sites like Grindr
3. Discussions of a frank and personal nature including incriminating admissions
4. Social network analysis tracking from known-gay persons to likely-gay persons

A state which pursued such a course in an efficient and methodical fashion is going to change their society in one of three ways:

1. Some 1% to 3% of their citizens will be incriminated, or
2. Gay activity will be pushed back off the internet and into other, more secret forms
3. Populations will deploy security software and techniques to protect themselves from the state online.

None of these outcomes were possible before the internet age.

Let us consider outcome (1). The notional authors of our policy presumably want this outcome, but now face a problem: processing 1%+ of your population, cross-cutting across all age groups and positions in society, through the criminal justice system. The economic costs alone are enormous, perhaps even substantially impacting GDP from loss of productive members of society, and a more nebulous *chilling effect.* The ability to effectively enforce law has greatly increased, but the actual impact of universal enforcement turns out to be at least in our terms of reference, generally counterproductive. Increased ability to enforce turns out to require a different approach.

Now let us consider this from an *information valuation* perspective. Faced with an individual citizen of interest, a police official may want to know “is this person gay?” and, indeed, may be conducting an investigation to that effect. In such an investigation, the ability to effectively pry into internet traffic may quickly sort the issue out, at least to a first approximation, modulo countermeasures (2; hiding) and (3; security software) above.

However, the cost of *all* such information is much, much higher than the utility in an individual case: the entire structure of a society has to change in order to be able to prosecute 1% to 3% of a population. The harm caused to the integrity of government also has to be considered: if the ability to do uniform enforcement of the law exists, but is not acted upon except in special cases, the question can be raised “are they serious about this or not?”

In short, the potentially useful single data point, when multiplied across an entire society, can turn into an unprecedented social cost. This is a simple issue of scarcity, akin to the tragedy of the commons: prosecution capacity is limited, surveillance capacity outstrips prosecution capacity, and a problem is created. Enforcement is now either arbitrary or differential, and either option removes the appearance of uniform enforcement of the law.

Past a certain point, each new surveillance case of homosexuality has a negative marginal utility.

Knowing one of something can be valuable. Knowing all of something is, in many cases, actively counterproductive. This is an effect beyond the law of diminishing returns: it’s the *decreasing value of omniscience.*

Simply put, all that is happening is cases are piling up beyond ability to prosecute, or society is transforming into a police state to do the enforcement which goes in step with new surveillance capabilities. The social equilibrium of the pre-network surveillance days cannot be maintained if network surveillance greatly increases criminal detection powers, but not enforcement with them.

This seemingly paradoxical relationship between increased detection capability and reduced social benefit is the correct analytical framework for thinking through the information acquisition potential of the network.

Now let us consider data retention. Suppose that our notional state caches all internet traffic for six months or a year, as is commonly mandated in modern states. We must then consider the potential for not simply instantaneous detection, but pervasive tracking over time, weaving networks of incrimination which take weeks or months to unfold, correlating the fine structures of the society's online life with illegal activity in the real world.

The cost of acquiring the information goes down and down. But the social cost of having it, even in a regime which thinks it wants to stamp out homosexual behaviour, goes up and up.

Cyber-security and cyber-crime

Let's briefly consider cyber-security and cyber-crime from this perspective. It's widely acknowledged that cyber-everything is an evolutionary arms-race which by nature favours the attackers. Defences are built up in response to perceived threat, and the much-sought-after zero day exploit is the imperceptible threat, the blow you never know is coming. Defence in depth and stealthy countermeasures reduce the risk of harm, but just as there is no flying tank (the air favours evasion, not durability) so the natural weave of the cyber-landscape favours attackers.

In this environment, hoovering up known-threats is much more credible than detecting and nullifying new, unknown, dangerous threats. People attacking out-of-date webserver software and turning machines into spam servers are a social nuisance, certainly, but there is no national security implication. Credit card fraud online punishes credit card companies and mid-tier retailers, but again is not socially destabilizing.

Amid all the noise, how to detect probes into critical infrastructure or national security targets?

Facing the price of magic bullets

We love standardization and monocultures. The *cheap acquisition cost* of certain kinds of information in a networked environment conceals the very substantial *costs of knowing* given certain common operating conditions.

Stamping out, for example, music and movie piracy could cost us the current relatively resilient and robust, internationally standard internet experience. National borders, and equivalent structures to ports and airport security could emerge in response to the desire for uniform enforcement of pre-digital copyright standards, with huge net-negative social costs. We may indeed be able to detect all the copyright infringement going on, but the price of the enforcement apparatus may vastly exceed the social benefit of uniform enforcement.

How many genuinely dangerous cyberwar actors are there, really?

Separating the wheat (cyberwar) from the chaff (cybercrime) requires a willingness to clearly define cyberwar, and to understand the classes of actors which might perpetrate it.

But let's return to the problem of assessing the *correct value of information*. How large is the social benefit or cost of failing to accurately estimate the correct risk of cyber? How does that change 5, 10, 15, 20, 25, 30 years in the future? Trident has a 30 year expected service life, and the B52 has planes 50 years old in service. 30 years backwards takes us to 1980, when only the very boldest science fiction foresaw the likely trajectory of digital networks. 30 years forward takes us into a time when quantum computing and similar will likely be routine, and millions of times more computer power and network bandwidth will be available to us, as dust-like network nodes float around and monitor our environments.

We're building a foundation for sliding up an exponential cliff of innovation while maintaining the basic functional structures of our societies. Nothing less.

Distinguishing what parts of this landscape are militarily relevant is critical to understanding how to the correct the valuation of relative capacities in the cyber domain.

The "magic bullet" capability in cyber is incredibly expensive to acquire. It likely requires deep hardware-level access to the majority of computing devices used by our enemies, and friends, and whole-of-technical-society collusion to this effect. China has some capability to generate that kind of capability, doing a substantial part of the world's whole manufacturing capacity. America might have a similar lead in software, owning two of the world's dominant operating systems (four if we count mobile devices.) For the UK to make an equal play at the level of *embedded capability* would require us to build up a domestic electronics and software industry which exported globally, and to engage deeply with those commercial processes. Short of this, all capacity is going to be partial.

But what do we actually need to know? The information with a positive marginal utility is

1. Hard to define, except after the fact.
2. Immune to analysis, even given the totality of information in the situation.
3. Concentrated in heavily-defended clearly-targeted secret networks.

It's the third of these facts which works to our advantage. While it's clear that there's a "brass ring" of *total information awareness*, the odds are that the marginal utility of the majority of the information retrieved and extracted from such systems is, in fact, negative. Much hay, few needles. On the other hand, the clearly marked "hard targets" contain more or less everything of direct military relevance anyway.

This may sound paradoxical, but the simple fact of the matter is that anything which is not guarded is not likely to be worth knowing.

We can winnow wheat from chaff simply by tending to ignore anything which is not defended. The risk from low-level hackers and cyber-criminals will force genuinely interestig information into easily identified bunkers, which can then be broken when possible.

This, of course, buys us little protection from non-state cyber-actors, hackers and cyber-terrorists and the like, but *attackers advantage* is not substantially diluted by pre-emptive attempts to detect danger. The zero day exploit in the hands of a newly radicalised hacker is not something a *good offence* can protect us from, and therefore we can only harden high

value targets like power grids and hospital IT systems: building our own bunkers in turn. We cannot comprehensively *get them before they get us* because who *they* are changes depending on government policy and the mood of unstable individuals with high technical competence. Attempts to prevent cyber-crime and cyber-terror before the fact are unlikely to be universally effective, and therefore do not obviate the need for staunchly robust critical systems and defence-in-depth of our assets.

Given that magic bullet capability is too expensive, and omniscience overwhelms our ability to act on the information gathered, what then is a rational strategy for integrating our new capabilities, and this new landscape, based on *realistic valuation of information*?

Path-oriented vs. Map-oriented knowing

Whether we like it or not, there is too much to know. The dreams of reason and the illusion of omniscience which go along with extended technical capabilities do little to actually make a reasonable map of how to integrate new technologies into existing organizational and bureaucratic structures. The people are constant! Technological changes do not revolutionize our ability to think, at least not yet.

Mis-estimating the value of information is at the heart of our cyber-challenges. We need to be consistent and clear in our understanding of *what access to information is worth* on the information gathering side, and *what control of information assets is worth* on the operational side. In most cases, we're going to find that only a very small part of the information operations landscape actually measures up well relative to other priorities, and that the costs associated with protecting those critical systems (in our case) are huge, and the costs of building reliable compromises of potentially-hostile systems are equally huge. Building maps of everything which is going on is vastly too expensive relative to mapping the critical pathways, but this approach requires clarity about what is critical and what is not.

Everything outside of those hardened networks is in the domain of *civilian inconvenience*. Who really cares if embassy web sites are down, as long as classified communications systems are up and secure? While there is enormous potential for economic damage if business systems, including non-critical infrastructure SCADA systems are attacked, given that there is no realistic hope of total interdiction on civilian networks, this is once again a case of target-hardening and whack-a-mole. The core assets of society are vulnerable, and will be so until civilian operating systems are secure: possibly never. But a cyber-attack is unlikely to be able to damage those assets beyond hope of restoration from backups, given offline storage, write-once media and the tendency for people to fix things after a crisis. The actual assets of relevance from a military perspective are as clearly flagged as they ever were: the other guy's stuff, in its hardened bunkers.

Morale and society

At this point in history, the common perception, reinforced by media of all kinds, is that life is pretty good in America and Europe (for slightly different reasons.) The internet is extremely disruptive to regimes which attempt to maintain a closed society with traditional values, and the social transformations wrought in the 1960s and 1970s as liberalization and social forces like the birth control pill took effect have come later (and in much harsher forms) to the Middle East and North Africa. It is likely that static, rigid, hostile societies will continue to have more to fear from their own populations adopting some variation of our social values than they have to fear from our intentional policy decisions. What is the real

cost of social control in an environment of pervasive monitoring: becoming the kind of society which keeps track of everything, and drowns in its own red tape.

Uniform enforcement is a close parallel to the dream of *total information awareness.* Information absolutism. The desire of totalitarian societies (and, to some degree, this now includes more extreme forms of Islam) to use technology to repress their own populations closely parallels the US DOD's persistent desire for omniscience. In both cases, a severe underestimation of the strain that effective ICT systems put on the rest of the organization has resulted in a certain kind of organizational paralysis: one seeks to control everything, the other seeks to know everything. A middle way is needed, in which clear distinctions are made about what we *need* to know, what we *want* to know, and what is irrelevant. Increasingly it will be necessary to expand, rapidly and clearly, the category of the irrelevant: to turn an institutional blind-eye to almost everything except a certain set of critical systems, knowledge and individuals who extreme focus can be placed on. It is this ability to partition the relevant from the irrelevant which is the key to operating in a cheap information environment, and understanding that the *value of information is often negative* is the most important lesson of all in plotting a strategy for dealing with this cheap information future that we find ourselves in.

Conclusion

There is an old saw among computer professionals, which goes:

data < information < knowledge < wisdom

Thus far we lack a coherent language to describe the limits of our ability to metabolize data up this chain into useful distinctions which allow us to act.

There is a complex and precise problem at the heart of our modern social problems related to the rapid advances in ICT: *we do not know how to value the information we now have access to.*

That is not a problem which can be solved in a military context, or likely in an economic/ business one. We are waiting for society to adapt to change, to make clear our new models, modes and objectives, and to provide the actors who are actually capable of making not simply intelligent but wise decisions about the cyber environment.

There's a delicate balance here between the environment left by nuclear, biological and chemical weapons and the basic social fabric. The age of total war has left us with a very distorted landscape of conflict, relative to human historical norms, and our human nature is poorly suited to understanding such total environments. Rather, a continual pressure and back-pressure, with shifting tides of advantage is the historic norm.

Cyber, because it is high tech, is being framed in a manner not-dissimilar to other total war weapons while, in fact, giving little real advantage in a total war context. Perhaps it is possible to shut down the power grid with cyber-weapons, but it is *almost certainly* possible to shut it down with hostile special forces. Reconceptualizing cyber along the lines of espionage, rather than placing it with high-tech weapons systems (nuclear, biological and chemical) greatly helps frame the real landscape of advantage here. Reflect on ENIGMA (and its breaking) and RADAR and its concealment. While the advantages derived from these breakthroughs was profound, it was as a *sharper point to an otherwise extremely hearty spear.*

Cyber’s value outside of a total war environment is still difficult to ascertain. The current push-and-pull around STUXNET, FLAME and possible Iranian retaliation for the same is perhaps our first opportunity to learn the limits of this space in operational terms. Being prepared to prosper on the second pass may be more useful than being the first actor.