Lockstep Consulting Pty Ltd
11 Minnesota Ave
Five Dock NSW 2046
Australia

The Committee Secretary
Senate Legal and Constitutional Committee
Department of the Senate
Parliament House
Canberra ACT 2600

22 February 2005

Re: Inquiry into the Privacy Act 1988

Dear Secretary

Lockstep Consulting Pty Limited is pleased to make this submission to your inquiry into the Privacy Act 1988.

Sincerely,

Stephen Wilson
Director

Summary

Australian citizens increasingly use the Internet to not just shop and operate their bank accounts, but to manage many aspects of their lives – including soon their electronic health records. Privacy threats in cyberspace are most often imagined by the public to involve "Big Brother" scenarios, like the tracking by governments of smartcard usage, and surreptitious linking of large and inscrutable databases. Yet Lockstep submits that privacy is more often invaded today by commercial and private interests exploiting weaknesses in cyber security. Examples include "phishing" scams, counterfeit websites, and identity theft. Our top priority in safeguarding privacy must be to ensure adequate levels of security around sensitive electronic services.

Our considered view, based on independent research and analysis, is that greater use of smartcards is urgently required to protect the privacy of Australians.

Smartcards do indeed present real privacy concerns, but if architected carefully and transparently, they become powerful Privacy Enhancing Technologies. In particular, smartcard technologies represent the only viable solution to phishing and, in the longer term, to spam as well.

It seems vital to us that governments not lag behind banks with regard to public policy on privacy and security. Most banks are responding to the threats of identity theft by making various forms of "two factor authentication" available. With medical records and other confidential government services being clearly more sensitive in nature than bank accounts, it is incumbent on law and policy makers to provide for at least the same level of security as do financial institutions. And yet, *of all the authentication solutions available today, only smartcard technologies can address phishing, counterfeit websites and spam*. To facilitate the rollout of this critical infrastructure, banks and governments could do more to work together, promoting the advantages of smartcards, and making available one another's respective card offerings for re-use by other service providers (with proper safeguards in place), thus providing consumers with as wide a range of choices as possible.

Turning to biometrics, privacy problems with these still immature technologies are, for the foreseeable future, far more likely to arise from their limitations than from their purported powers. The simple truth is that biometrics are not yet reliable enough for large scale rollout. And they fundamentally can never be as foolproof as people have been led to believe by cursory viewing of sci-fi movies. In this submission, by applying no more than a junior high school level of scientific analysis, we hope to expose some of the inescapable limitations of biometrics, and to thus help set more realistic expectations amongst law and policy makers as to the real abilities of biometrics.

Summary recommendations

1. Further research should be conducted into the threats posed to privacy by phishing and spam, especially in light of electronic health records and similarly sensitive personal information becoming increasingly available over the Internet.

2. Address the risks of data linkage and monitoring of individuals' online activities by fostering the use of multiple personal identifiers, and by adopting information technologies which actively support secure, segregated multiple identifiers.

3. Consider the possibility that government could play an active role in fostering the community's take-up of smartcards, by for example

expanding the role of the Medicare smartcard, and by encouraging computer manufacturers to make built-in smartcard readers more widely available in the Australian marketplace.

4. Further, explore ways in which all levels of government could work with financial institutions (being the other major players in smartcard rollouts in the medium term) to enable cooperative re-use of new smart debit and credit cards as secure "containers" for personal identifiers, to safeguard government and private sector electronic service delivery.

5. Take great care in the selection of any biometric authentication technology, especially in regard to their demonstrable rates of failure to enrol, False Positives and False Negatives.

6. Because biometrics cannot be revoked and re-issued in the event of identity theft, we strongly recommend that biometrics not be used as the sole means for authentication in any sensitive application.

7. Wherever possible, avoid the use of passive RFID transponders in e-passports, implantable microchips and so on. Rather, cryptographically active contact-less devices should be used, with the power to de-identify holders, mitigating unauthorised identification and linkages.

Declaration of interests of Lockstep Consulting

Lockstep Consulting Pty Ltd was established in early 2004 by Stephen Wilson, a leading international authority on identity management and information security. Lockstep provides independent advice, analysis and management consulting in security policy and strategy, authentication and privacy. Lockstep also undertakes independent research – some self-funded and some under contract – into applied technology, particularly smartcards and Public Key Infrastructure (PKI).

Current Lockstep clients include the Health Insurance Commission, the Australian Government Information Management Office (AGIMO), the Australia-ASEAN Development Cooperation Programme (AADCP), Telstra, and the US-based Organisation for the Advancement of Structured Information Standards (OASIS). Mr Wilson is a member of policy bodies including the Gatekeeper Policy Committee, the Australian IT Security Forum, the APEC eSecurity Task Group, and the national IT Testing Accreditation Advisory Committee. He was a founding member of the National Electronic Authentication Council (NEAC), and he sat on the previous Federal Privacy Commissioner's PKI Reference Group.

Our understanding of the Committee process

It is understood and agreed that submissions made to the Inquiry will in general be made public. Lockstep's Stephen Wilson is available to provide further information on any of the issues raised in this submission, and would welcome the opportunity for further dialogue.

The structure of this Submission

The Terms of Reference of the Committee Inquiry included:

> *(a) the overall effectiveness and appropriateness of the Privacy Act 1988 as a means by which to protect the privacy of Australians, with particular reference to:*
>
> *(i) international comparisons,*
>
> *(ii) the capacity of the current legislative regime to respond to new and emerging technologies which have implications for privacy, including:*
>
> ***(A) 'Smart Card' technology and the potential for this to be used to establish a national identification regime,***
>
> ***(B) biometric imaging data,***
>
> ***(C) genetic testing and the potential disclosure and discriminatory use of such information, and***
>
> ***(D) microchips which can be implanted in human beings (for example, as recently authorised by the United States Food and Drug Administration),*** *and*
>
> *(iii) any legislative changes that may help to provide more comprehensive protection or improve the current regime in any way;*
>
> *(b) the effectiveness of the Privacy Amendment (Private Sector) Act 2000 in extending the privacy scheme to the private sector, and any changes which may enhance its effectiveness; and*
>
> *(c) the resourcing of the [OFPC] and whether current levels of funding and the powers available to the Federal Privacy Commissioner enable her to properly fulfil her mandate* (emphasis added).

Lockstep is qualified to address the four technology issues (A) through (D). We trust that the Committee will consider our input about these particular issues. The remaining major sections of this paper address smartcards, biometrics, genetic data and implantable microchips in turn. Our major recommendations are summarised under the *Overview* above.

Smartcards and their potential use in a national identification regime

We assume from the way in which this matter is framed in the Terms of Reference that the Committee is taking advice on how to *prevent* the use of smartcards leading inadvertently to a national identification regime. There is some concern amongst the public that something intrinsic to smartcard technology could lead accidentally (or even surreptitiously) to an "Australia Card". This concern is certainly not groundless, but it does tend to overshadow the strong potential benefits of smartcards for privacy.

Pros and cons of smartcards with respect to privacy

As yet nobody seems quite sure if smartcards represent a Privacy Invasive Technology ("PIT") or a Privacy Enhancing Technology ("PET"). Being a programmable technology, the objective (if perhaps unhelpful) answer is that smartcards can be both.

There are indeed sound reasons to worry that smartcards can threaten privacy. Smartcards can store a relatively large amount of personal data, and that data can be written and/or read without the card holder being aware of it. Thus a smartcard might divulge detailed demographic data to a terminal and thence to backend IT systems whenever the card is used, allowing detailed profiles to be built up of one's buying habits, movements through ticketing and tolling systems, encounters with Human Services systems, and so on.

On the other hand, unlike regular magnetic stripe cards which can be "skimmed" if they fall into the wrong hands, smartcards offer robust protection of stored data against unauthorised access. To merely read , let alone modify, the data in a smartcard, it is necessary for an intelligent card reader to positively identify itself to the microchip on the card, and to satisfy various security criteria. In contrast, data encoded on magnetic stripes can be read off and copied using simple, readily available recording equipment.

A hierarchy of access controls can be programmed into modern smartcards, allowing personal data to be organised according to its ownership and level of sensitivity. Data can be optionally PIN-protected; access to different sets of data can be controlled with different PINs if so desired. One thing that makes smartcards "smart" is their ability to be programmed to make decisions about when and where they will exchange data with the outside world. If the correct conditions do not obtain – including the proper terminal equipment, the proper security protocols, and the presentation of a recognised PIN – then the smartcard will simply refuse to "talk".

These sophisticated capabilities can be used to protect card holder privacy in many different ways. In our opinion, of particular relevance to the Committee's inquiry are two unique abilities: management of multiple identifiers, and protection against website fraud such as phishing.

The benefits of multiple identifiers and multiple identities

Multiple identifiers are good for privacy. If individuals are permitted to use different identifiers when interacting with different information systems, then data matching is inherently more difficult. The National Privacy Principles themselves proscribe the re-use of government issued identifiers by the private

sector. And many independent commentators have argued for multiple identifiers; see for example the highly regarded recent US government sponsored study *Authentication Through the Lens of Privacy* [1].

It is common today for individuals to carry on their person a dozen or more personal identifiers.[1] Usually these are numbers, printed on the surface of plastic cards. If we think deeply about identifiers, it is possible to appreciate them as standing for separate *identities*. People take on different roles when they put their various identifiers to use, as the following examples show.

Multiple Identities Example #1: Business banking and Personal banking

I am the Company Secretary of Lockstep Consulting Pty Limited. Lockstep's bank is Westpac. I am a signatory to the Lockstep corporate account and I have custody of a Westpac key card issued for the company accounts. I also happen to hold a personal bank account with Westpac. When I bank on behalf of Lockstep, I exercise *a different identity* compared with when I bank on my own behalf, even if I am in the same branch or at the same ATM. For obvious reasons, it is unlikely that anyone would ever wish to merge these two identities into one.

Multiple Identities Example #2: Using the *HICAPS* healthcare payment system

"HICAPS" is a point-of-sale payments system with which patients with private health insurance can settle their accounts instantaneously at dentists, physiotherapists and so on. The receptionist swipes the patient's health insurance card and enters an item number for the service delivered, and the HICAPS system automatically transfers the insurance payout to the provider's bank account. The patient then pays the balance of the account. If the patient happens to pay by credit or debit card, then the receptionist will have to swipe that card separately. It strikes some as odd that the one patient has to have two cards swiped. Yet it is perfectly understandable that the HICAPS system in effect sees *two* identities: one being the insurance policy holder who has a certain entitlement, and the other being a bank account holder who is paying for the service. It would clearly be contractually and logistically complex to merge the two identities into one.

In general we should permit people to conduct their affairs using multiple identities and identifiers. And yet there is an implicit objective in many electronic authentication schemes to combine identifiers. Sometimes this is a worthy aim; when it comes to computer passwords, we often find ourselves with too many identifiers to manage efficiently and safely. In other areas however, it is less obvious why we should strive for a single identifier. For

[1] Common identifiers include a Medicare number, a state driver license number, one or more bank account and credit card numbers, a Tax File number, a telephone account number (and perhaps the telephone number itself), a private health insurance policy number, an employee number, one or more superannuation fund numbers, one or more e-mail addresses, a passport number, a utilities company billing number, a blood bank number, a frequent flyer number, and any number of association or club membership numbers.

instance, Public Key Infrastructure (PKI) is usually imagined to mean some sort of large scale identification scheme, and many PKI standards do indeed have their roots in these sorts of historical ambitions. Contemporary thinking about PKI however is starting to lean towards multiple credentials.

Dr Stephen Kent, joint chairperson of one of the authoritative Internet Engineering Task Force PKI standards committee, and author of the report *Authentication Through the Lens of Privacy* [1], recently offered the following observations and recommendations about identity:

> *For big [PKI], there is an implicit assumption that a single [credential] is all that a user should need. This assumes that one identity is sufficient for all applications, which contradicts experience. But, in support of personal privacy, authorization credentials that do not identify users are preferable.* [2]

And Lockstep's Stephen Wilson, writing for the Board of Directors of the Australian PKI Forum in late 2003, pointed out that:

> *The idea of multiple [credentials] was once alarming, but when embedded invisibly in convenient forms such as smartcards, they need not be any harder to use than conventional plastic cards. There is increasing awareness from the perspectives of privacy, usability and commerciality, that a single identity would not be useful in any case. The reality of physically different cards for banking, drivers licence, health insurance, professional memberships and building access is here to stay – irrespective of whether the cards are based on magnetic stripes or PKI.* [3]

Smartcards and the management of multiple identifiers

The first thing to be said about multiple identifiers is the almost trivial observation that the typical Australian in three or four years time can expect to carry multiple smartcards. These will embody diverse identities, just as plastic cards do today. Lockstep's own conservative prediction[2] of smartcard growth in Australia is shown below, based on the banking sector's transition from magnetic stripe to smartcard technology, plus just two publicly announced schemes – the Medicare smartcard [4] and the New Queensland Driver Licence [5].

---

[2] Our forecast is based on the following data. Medicare will roll out smartcards from early 2005; all six million Medicare customers are likely to carry smartcards by the end of the decade. Queensland's new licence will roll out from 2006, eventually reaching over 2.5 million people. The *Sydney Morning Herald* of 20 February 2004 reported that credit cards numbered 11.1 million in 2003; credit card companies require banks to migrate to smartcards from c. 2006. In 2004, the Global Platform group reported there were already half a million ANZ *First* smartcards on issue [6]. Overall, compound annual growth of 3% is assumed.

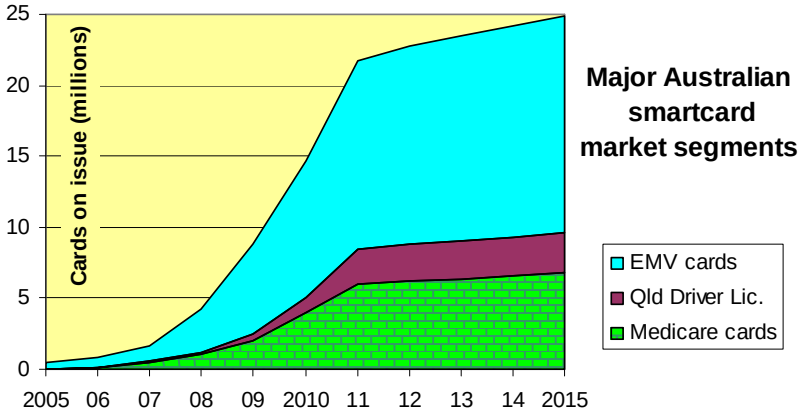

Figure 1: Lockstep's growth forecasts for domestic smartcards

Notable overseas smartcard rollouts include:

UK Chip and PIN cards	75,000,000
Taiwan health cards	22,000,000
US Visa smart credit cards	12,000,000
US government employee cards	4,000,000
Hong Kong SMARTICS identity smartcard	1,500,000

More significant than the mundane likelihood of there being multiple smartcards, the technology has the following crucial and unique abilities to safely manage multiple identities:

- *smartcards can carry multiple identifiers on a single card,*
- *they can un-link identifiers from natural identities*
- *they can keep multiple identifiers segregated from one another.*

It is beyond the scope of this paper – and probably the Committee's present inquiry as well – to go into technical detail on smartcard architectures. Nevertheless, Lockstep submits that it is important for the Committee to gain a high level understanding of how smartcards can handle personal identifiers, as will be explained in plain language in the following sections.

Smartcards are more than mere storage devices for personal data

It is commonly presumed (or represented) that smartcards basically act as storage devices, not markedly different in principle from conventional magnetic stripe cards, or even pieces of note paper, except for their capacity. And so, when health identifiers for instance are discussed and debated, it is assumed

that they will be stored in a variety of technically equivalent ways, and that individuals will have a broad choice of storage options.

For example, the recently released Health*Connect* Business Architecture says that "the unique Health*Connect* identifier issued to consumers is to be stored on a new Medicare smartcard" but that "the Medicare smartcard is one means of holding the consumer's Health*Connect* identifier" [7]. The Business Architecture stops short of elaborating precisely how the identifier might be stored. Elsewhere, the documentation is non-committal about the details, stating with deliberate imprecision that "the national consumer identification service is expected to include a client (consumer) master index linked to the consumer-held smartcard" [8].

It is understandable that the architects until now have wished to hedge their bets on technical details such as how the identifier should be stored. But on the other hand, if health identifiers are not managed properly from day one, and if the special capabilities of smartcards are not fully utilised, then unfortunate privacy compromises will be inadvertently built into the system.

Smartcards as Privacy Enhancing Technology: an illustration

Most modern multi-function smartcards can be used to manage and segregate distinct sets of private data as depicted below. ***Note that the proposed Queensland driver licence smartcard is used merely for illustration; the example systems shown are fictitious.*** Three sets of private data are shown, each retained inside its own secure memory area. The smartcard's computer processor can control access to each private data set in a number of ways, including requiring that a requesting system present the appropriate cryptographic key. Thus the driver licensing system and an additional e-health record system need not be able to see each others' data. Indeed, without access to a system's particular cryptographic key, its association with the smartcard is totally invisible to the outside world.

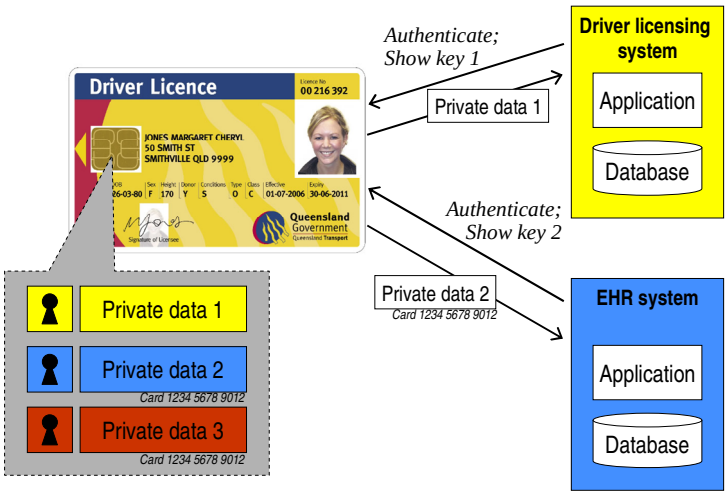

Figure 2: Schematic privacy architecture for multiple identifiers on single card

Minimising data linkages
Note also that cryptographic smartcards provide the extra capability of optionally tagging stored data with an indelible digital signature. This signature can be unique to the smartcard chip – as indicated in the diagram by the schematic card number "123456789012" – *without giving away any information at all about the card holder*. This provides powerful means for eliminating data linkages and anonymising transactions carried out using the smartcard. For example, if *Private Data 2* in the illustration was the health identifier of card holder Ms Jones, then no data written into the e-health record system can be linked by a third party to Ms Jones, unless they have access to her smartcard.

The fight against website fraud and phishing

Some of the greatest real-world threats to privacy today are a consequence of website fraud. Just as hackers can take control of Internet bank accounts, soon there will be escalating opportunities for them to hijack far more sensitive online resources, most notably electronic health records.

The major risks faced by all organisations offering services over the Internet – including governments – include "ghosting", where unwitting visitors are re-directed to phoney websites, and "phishing", where counterfeit e-mails dupe users into divulging personal information or undertaking bogus transactions.

In the cyber-crime arms race, standard software defences are breaking down. To protect the identities of websites, special security codes – technically known as SSL digital certificates – are loaded onto certified web servers, and checked against 'master codes' pre-programmed into browsers. Yet hackers can now surreptitiously swap these codes with fraudulent ones, and so, in the words of one reporter, the "reassuring little padlock icon is essentially worthless". A host of new personal security tools – one-time PIN generating tokens, secret look-up tables, biometrics and even SMS text messages – are being rolled out to protect users from identity theft. *Yet none of them protect the service provider itself from electronic impersonation by phishing and ghosting.*

Lockstep's research has led us to conclude that only smartcards (and the closely related USB key or "dongle") have the necessary functionality to prevent impersonation of web servers, at the same time as safeguarding users from identity theft. Put simply, smartcards can safely store the Internet security codes so they cannot be tampered with.

We are especially concerned that electronic health record systems are being contemplated where individuals will be able to access their personal information over the Internet with simple password security. The Health*Connect* Business Architecture for instance proposes that consumers be

able to log on with their Health*Connect* identifier and PIN, treating the identifier like it were a bank account number.[3] And the privacy implications of spam – which is bad enough in regular e-commerce – will be immeasurably worse when Health*Connect* consumers begin to receive unsolicited invitations from overseas drug sellers, fringe healthcare providers, disreputable insurance companies, and private healthcare groups touting for business.

Conclusions regarding smartcards

The most serious threats to privacy today are not related to the future "Big Brother" potential of monitoring card holders, as important as this possibility may be. Rather, far more serious privacy risks arise today from the increasing vulnerability of web-based electronic service delivery, and from dumbing down security for consumer access to their electronic health records. Smartcards offer the only truly effective solution to the scourges of website ghosting, phishing and, in the longer term, spam.

The main substantive objection to smartcards – that they might facilitate monitoring and profiling of card holders – is perhaps best dealt with by encouraging the use of multiple digital identities and multiple smartcards, and by deploying anonymity solutions that un-link identifiers from their holders.

Government could take a lead on this front in several ways, including:

- promoting and expanding the use of the new Medicare smartcard, especially to authenticate consumers' access to their Health*Connect* records over the Internet
- working with programmes such as the Queensland smartcard driver licence to secure online services
- investing in smartcard readers in all new government computer purchases, and
- providing incentives for computer manufacturers to accelerate the introduction of readers in their standard product offerings.

Government should not be seen to be lagging behind the financial sector in the adoption of sophisticated authentication technology, for surely our electronic health records deserve even greater security than do our bank accounts.

[3] "From the time of registration, consumers will have the ability to use a Health*Connect* Consumer Access Portal to view and contribute to their Health*Connect* record. This will initially be via a web browser system requiring the consumer's Health*Connect* identifier and PIN for access" [8].

Biometrics

The expression 'the devil is in the detail' holds true for biometrics like no other technology today. Biometrics appear profoundly simple in operation,[4] but the science, engineering and applied product design are all still in their infancy. Lockstep's most fundamental recommendation on biometrics to the Committee would be that a "reality check" is required as to their real world abilities, because for the foreseeable future, privacy problems from biometrics are far more likely to arise from their limitations than from their purported powers.

The public's inflated expectations of biometrics

Vendors frequently draw parallels between their technologies and the way humans perform 'biometric matching' when we recognise for example a friend's face or their voice. Yet such simple throw-away comparisons are misleading. Science still has only the most basic understanding of how the human brain works; "artificial intelligence" remains controversial in research circles, and wholly immature in the commercial world. Anyone who has used speech recognition software knows that it falls well short of 100% success in simply picking out words; the task of accurately recognising individual speakers by their accent and intonation may be expected to be harder still.

Enormous over-simplifications pervade most peoples' understanding of how biometrics work. We are frequently led to believe that it's as simple as scanning someone, looking up a database, and having their identity pop out.[5] But the reality of biometrics is more nuanced. Lockstep contends that not enough is generally known by laypersons about biometrics for robust policy determinations to be made. Large scale biometrics studies overseas have recognised that the field is beset by complex technical problems, a lack of standard ways to measure performance, and a bewildering range of competing methods (refer to [12] for a comprehensive survey of current concerns).

---

[4] Some biometric vendors show clips from science fiction films – like "Minority Report" and "Diamonds are Forever" – as part of their product demonstrations, as if they constitute actual *case studies*.

[5] Biometric applications which pick an individual out of a database of enrolled templates are termed *one-to-many positive identification* systems. For large user groups, one-to-many systems face major technical challenges; for instance, the US General Accounting Office has concluded that the "performance of facial, fingerprint, and iris recognition is unknown for systems as large as a biometric visa system" [13]. Successful one-to-many systems remain confined to secure data centres, military facilities, bank vaults etc. where the group of enrolled users is very small, and where relatively frequent false negatives are well tolerated. More common than one-to-many identification is *one-to-one verification* where a user is matched against a single template held in a portable device, like a smartcard or e-passport. Here the biometric is used in place of a PIN, and does not remove the need for the "something you have" factor, namely some physical token.

Further, there is a sort of in-built optimism about biometrics based on common ideas about the uniqueness of biological traits. It is lore for instance that fingerprints are unique. However, a number of criminal prosecutions based on fingerprint analysis have recently been over-turned on appeal, and this has led to a re-examination of the 'science' of fingerprinting. It turns out that very little scientific testing has ever been done of the accuracy of forensic fingerprinting. Even more fundamentally, the very question of uniqueness actually remains open: "although conventional wisdom since the nineteenth century has accepted the doctrine that no two fingerprints are alike, no one has really proven the proposition's validity" [14].

Even the 'gold standard' biometric – DNA[6] – is not what it might first appear. It is true that each of us (save for identical twins) carries a unique genetic sequence, *but the complete sequence is not what is measured in typical "DNA testing" for identity*. Instead, forensic DNA testing examines only particular subsets of one's genes. The founder of modern DNA testing, Sir Alec Jeffreys recently highlighted some of the practical issues, reported in the media as follows:

> *DNA testing is not an infallible proof of identity. While Jeffreys' original technique compared scores of markers to create an individual "fingerprint," modern commercial DNA profiling compares a number of genetic markers — often 5 or 10 — to calculate a likelihood that the sample belongs to a given individual. Jeffreys estimates the probability of two individuals' DNA profiles matching in the most commonly used tests at between one in a billion or one in a trillion, "which sounds very good indeed until you start thinking about large DNA databases." In a database of 2.5 million people, a one-in-a-billion probability becomes a one-in-400 chance of at least one match.* [15]

So we see that the field of biometrics is rather more murky than it might at first appear. Lockstep urges greater caution in the large scale application of biometrics, as a matter of principle, because the field is riddled with so much uncertainty. But rather than attempt to survey all of the issues in this brief submission, we seek to bring to the Committee's attention just two easily understood problems: the imperfect performance of biometrics, and the impossibility of recovering from biometric identity theft when it eventually occurs.

---

[6] In any event, the practical use of DNA for automatic, real time identification, akin to iris or fingerprint scanning, is many years away. State-of-the-art field laboratory analysis of DNA today (as is necessary for forensic analysis of disaster victims) takes several days using sophisticated apparatus costing hundreds of thousands of dollars. The technique will not be easily nor economically miniaturised.

The fallibility of biometrics in practice

The least technical criticism of biometrics concerns the fallibility of all measurement apparatus. Cameras, sensors and microphones – like human eyes and ears – are imperfect, and the ability of a biometric authentication system to distinguish between subtly different people is limited by the precision of the input devices. Even if the underlying biological traits of interest are truly unique, it does not follow that our machinery will be able to measure them faithfully.

One iris-based security manufacturer promotes its product with the impressive claim that the probability of two individuals' iris patterns matching is one in ten to the power of 78. These are literally astronomical odds; there are fewer atoms in the universe than 10^{78}. Yet does this figure necessarily tell us how sensitive the end-to-end biometric system really is? Consider the fact that there are ten billion stars in the Milky Way. If two people look up in the night sky and each pick a star at random, is the probability of a match one in ten billion? Of course not, because of the limits of our measurement apparatus, in this case the naked eye. Interference too affects the precision of any measurement; the odds of two people in a big city picking the same star might be no better than one in a hundred.[7]

"False Positives" and "False Negatives"

Biometric authentication entails a long chain of processing steps, all of which are imperfect. Each step introduces a small degree of uncertainty, as shown in the schematic below. Furthermore, uncertainty is inescapable even before the first processing step, because the body part being measured can never appear exactly the same. The angle and pressure of a finger on a scanner, the distance of a face from a camera, the tone and volume of the voice, the background noise and lighting, the cleanliness of a lens all change from day to day. A biometric system cannot afford to be too sensitive to subtle variations, or else it can fail to recognise its target; a biometric must tolerate variation in the input. Inevitably this means the system can sometimes confuse its target for someone else.

[7] Incidentally, independent testing shows a False Accept Rate of 0.00001% for iris [11]. This is 100,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000,000 times worse than implied by the figure of 1 in 10^{78}.

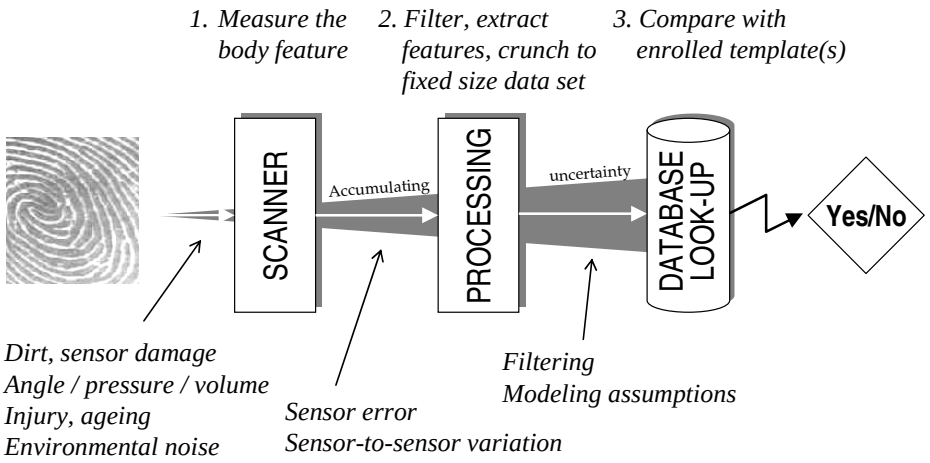

Figure 3: How uncertainty accumulates during biometric data processing

Therefore all biometric systems inevitably commit two types of error:

1. A "False Negative" is when the system fails to recognise someone who is legitimately enrolled. False Negatives arise if the system cannot cope with subtle changes to the person's features, the way they present themselves to the scanner, slight variations between scanners at different sites, and so on.

2. A "False Positive" is when the system confuses a stranger with someone else who is already enrolled. This may result from the system being rather too tolerant of variability from one day to another, or from site to site.

		Is the person legitimate?	
		Yes	No
Does the measured biometric match an enrolled template?	Yes	True Positive Admit an enrolled user	**False Positive** (aka "Type 2 Error") Admit a non-enrolled user, possibly an imposter
	No	**False Negative** (aka "Type 1 Error") Reject an enrolled user	True Negative Reject a non-enrolled user

Figure 4: The two basis types of error in any biometric

False Positives and False Negatives are inescapably linked. If we wish to make a given biometric system more *specific* – so that it is less likely to confuse

strangers with enrolled users – then it will inevitably become less *sensitive*, tending to wrongly reject legitimate enrolled users more often.

A design decision has to be made when implementing biometrics as to which type of error is less problematic. Where stopping impersonation is paramount, such as in a data centre or missile silo, a biometric system would be biased towards false negatives. Where user convenience is rated highly and where the consequences of fraud are not irreversible, as with Automatic Teller Machines, a biometric might be biased more towards false positives. For border control applications, the sensitivity-specificity trade-off is a very difficult problem, with significant downsides associated with both types of error – either immigration breaches, or long queues of restless passengers.

The following schematics illustrate how a highly specific biometric system tends to commit more False Negatives, while a highly sensitive system exhibits relatively more False Positives.

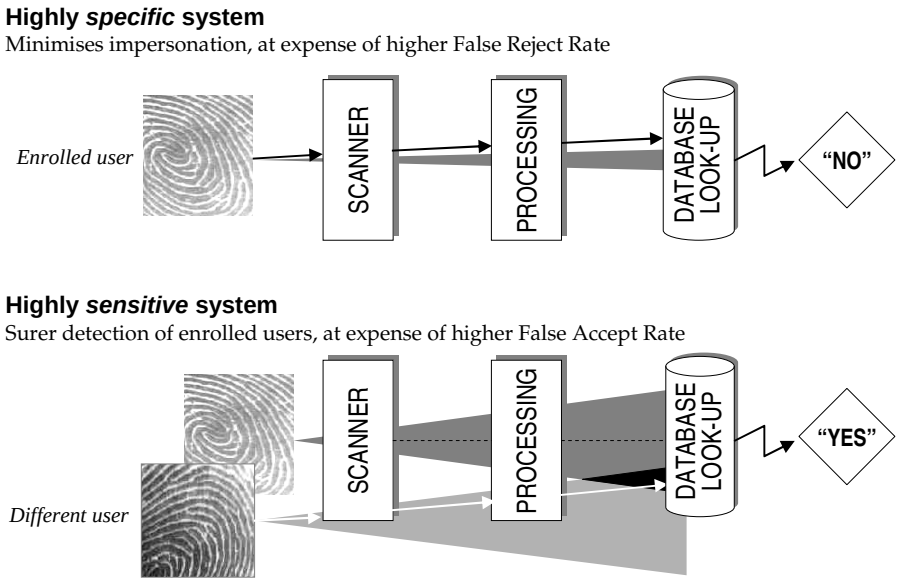

Figure 5: The "sensitivity-specificity" tradeoff illustrated

The trade-off is demonstrated in the next figure by actual data for three common biometrics, as tested by the British Government (graphic adapted from [11]).

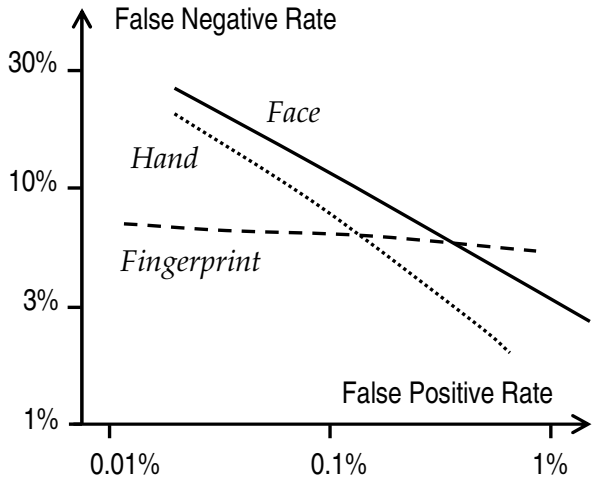

Figure 6: UK Government False Positive / False Negative test results

Failure To Enrol

Over and above the issues of False Positives and False Negatives is the unfortunate fact that not everyone will be able to enrol in a given biometric authentication system. At its extremes, this reality is obvious: individuals with missing fingers, or a severe speech impediment for example, may never be able to use certain biometrics.

However, failure to enrol has a deeper significance for more normal users. To maximise performance for both False Positives and False Negatives, a biometric method must impose requirements on the quality of its input data. A fingerprint scanner for instance will perform better on high definition images, where more fingerprint features can be reliably extracted. If a fingerprint detector sets a relatively stringent cut-off for the quality of the image, then it may not be possible to enrol people who happen to have inherently faint fingerprints, such as the elderly, or those with particular skin conditions.

More subtle still is the effect of modelling assumptions within biometric algorithms. In order to 'make sense' of biological traits, the algorithm has to have certain expectations built into it as to how the features of interest generally appear and how those features vary across the population; after all, it is the quantifiable *variation* in features which allows for different individuals to be told apart. Therefore, face and voice recognition algorithms in particular might be optimised for the statistical characteristics of certain racial groups or nationalities, making it difficult for people from other groups to be enrolled.

The impossibility of enrolling 100% of the population into any biometric security system has important implications for public policy. Clearly there can

be at least the perception of discrimination against certain minority groups, if factors like age, foreign accent, ethnicity, disabilities, and/or medical conditions impede the effectiveness of a biometric system. And careful consideration must be given to what 'fall-back' security provisions will be offered to those who fail to be enrolled. If there is a presumption that a biometric somehow provides superior security, then special measures may be necessary to provide equivalent security for the un-enrolled minority.

The inability to respond to Identity Theft

We hope that the foregoing discussion serves to show how biometrics in practice do not perform as well as often portrayed, and that they will inevitably commit various types of errors. Given this reality, *the most worrying aspect of biometrics is that they offer no fall-back in the event of identity theft.*

One could be forgiven for assuming that biometrics are immune to identity theft. At the biological level, the fundamental notion of course is that your biometric is not transferable, in contrast to passwords which can be copied or guessed, and physical tokens which can be stolen. Yet most biological traits can in fact be duplicated with sufficient fidelity to fool most biometric detectors.

In 2002, a German popular electronics magazine bench-tested a host of low price biometric systems, aimed at the small office / home office market. They found most of these devices almost comically easy to thwart. Some face and iris scanning products were subverted by photos of enrolled persons; some fingerprint readers were readily made to trigger on the latent print left behind by the last user [16]. We don't suggest that these results are typical of more expensive products, but they are illustrative of the in-principle vulnerability to identity theft. And more serious vulnerability assessments have been conducted, most notably the testing by Japanese researcher Dr Tsutomu Matsumoto, which found that 80% of fingerprint readers were susceptible to attack using replica "gummy" fingers made from gelatine [17][18].

Now the critical question is: *What are we to do in the event that an individual's biometric identity becomes compromised?* We know what do when any other authenticator is stolen, be it a password, a magnetic stripe card, or a smartcard: we simply revoke it and issue a new one. But as things stand today, *no biometric identifier can be cancelled and re-issued*. In the event of biometric identity theft, there would appear to be no alternative but to withdraw the affected user from the system.

Disaster recovery and contingency planning are the mainstays of good security practice. It is axiomatic that no system is ever 100% secure; all truly secure systems include a contingency plan to cope with the event of a critical breach,

even where the likelihood of such a breach is extremely low. Lockstep submits that all electronic service schemes proposing to deploy biometrics to members of the public should address identity theft contingencies plans as an absolute top priority.

Particular risks with storing fingerprint templates

The ease with which latent fingerprints can be retrieved from objects and reproduced in gelatine should be heeded by the designers of all schemes where fingerprint biometric templates are stored within a portable device. Affected proposals include the new US *Personal Identity Verification* smartcard for federal employees [19], and electronic passports the world over. Last year, Lockstep's Stephen Wilson wrote in *New Scientist* magazine of the dangers of storing fingerprint templates within portable electronic devices [20] (emphasis added):

[A] new cell phone incorporates a fingerprint reader as a safeguard against theft (New Scientist 24 July 2004, p 26). Several laptop computers, the odd mouse, even some experimental smartcards now have this "security" feature. Yet it is a worrying gimmick, ***closely equivalent to writing the PIN on the back of your credit card****.*

A majority of commercial fingerprint detectors can be fooled by replica prints. In 2002, Tsutomu Matsumoto devised the infamous "Gummy Bear Attack", in which a gelatin candy molded with latent fingerprints transferred from a drinking glass proved effective against 80 per cent of readers tested.

So if you lose your fancy phone [or smartcard], a clever thief will find your biometric security information very conveniently left behind all over the keypad. One wonders whether disposable latex gloves will become the next weapon in the war on identity theft?

The standard response to these issues is for "liveness" detection to be built into the fingerprint scanner, so it won't be fooled by replicas. Yet in commercial practice, robust liveness detection remains uncommon in fingerprint systems.

Conclusions regarding biometrics

Despite all the hype, no biometric system can preclude identity theft. Typical False Positive error rates of 0.1% are rather higher than the public might imagine, and indicate that biometrics in practice are not at all immune to impostors. Therefore, the most significant privacy risks posed by biometrics today relate to their shortcomings and not their purported powers:

— The consequences of identity theft will probably be exacerbated by a false sense of security engendered by exaggerated impressions of biometrics' accuracy, plus the fact that once compromised, a biometric identity cannot be readily re-issued.

— Policy makers must carefully guard against compromised individuals and those who, through no fault of their own, cannot be successfully enrolled into a biometric system, being relegated to second rate, fall-back security options.

Genetic data

Lockstep's primary concern in the area of genetic data is that the issue tends to overshadow more important everyday possibilities for exposing highly sensitive data about individuals' health.

Let us state up front that Lockstep shares the concerns of many commentators that the collection and retention of genetic material from large numbers of individuals poses serious privacy risks, especially with regard to function creep, and the long term security of associated identifying information (see for example [15]). Nevertheless, we believe that genetic data is relatively less critical than many other types of personal information.

Let us first review the definition of "health information" applied by the latest proposed National Health Privacy Code [21]:

> *(a) information or an opinion about:*
> *(i) the physical, mental or psychological health … of an individual; or*
> *(ii) a disability (at any time) of an individual; or*
> *(iii) an individual's expressed wishes about the future provision of health services to him or her; or*
> *(iv) a health service provided, or to be provided, to an individual … or*
> *(b) other personal information collected to provide … a health service; or*
> *(c) other personal information about an individual collected in connection with the donation … of organs or body substances; or*
> *(d) genetic information about an individual in a form* ***which is, or could be, predictive (at any time)*** *of the health of the individual or any other individual, including antecedents or descendants* (emphasis added).

During the National Health Privacy Working Group's public consultation process, Lockstep's Stephen Wilson submitted that there is additional information, not included in the proposed definition of "health information", which is in fact currently more sensitive than genetic data:

> *[There] is behavioural and attitudinal information that may be strongly predictive of a person's state of health (or of their belief as to their own state of health):*
> — *diet and eating habits*
> — *sports and exercise regime*
> — *usage of non-prescription medications*
> — *usage of herbal and non-traditional remedies.*
> *We suggest that* ***the predictive power of information about diet and exercise is stronger today than is almost any genetic information at hand****, and is therefore just as deserving (if not more so) of inclusion in the definition* [22].

The crucial practical issue here is that detailed lifestyle information, falling outside the current definition of "health information", is increasingly compiled via routine e-commerce transactions by any number of organisations, often inadvertently. Gym memberships, sporting activities, fast food orders, even the online purchases of herbal remedies like St John's Wort (a popular alternative treatment for depression), may all be monitored and tracked, apparently without attracting the stringent provisions applied to health information by the Privacy Act. *En masse*, this lifestyle information paints a vivid picture of an individual's perceived or actual health.

Lockstep suggests that this loophole in the definition of "health information" be addressed before too much further effort is expended on as yet more academic issues around genetic data.

Implantable microchips

The practical benefits of implanting microchips in humans would in some cases appear to be significant. Lockstep has not researched these developments in depth, but we understand that the primary application for now is tagging patients during hospital stays, with an electronic version of the traditional wrist band. By mitigating human error in the reading of patient details, it seems likely that medical misadventure can be reduced. If the RFID technology can somehow be constrained to the hospital environment, then the benefits would seem to outweigh the privacy risks, which primarily relate to eavesdropping.

If indeed human-implantable microchips are only envisaged in hospital applications, then we can presume that the per-chip cost is not as important a consideration as it is for the disposable RFID tags used in manufacturing and retail distribution. If so, then in principle it should be possible to deploy rather more sophisticated microchip technologies than the ordinary passive RFID transponder, which does little more than act as a long range bar code.

Lockstep therefore submits that cryptographically active wireless devices – equivalent to contact-less smartcards – should be preferred in implantable microchips, and that they be programmed with anonymity features, such as the chip-specific digital signature technique mentioned above under *Minimising data linkages*. We suggest that the same principles also apply to the new electronic passports.

References

[1]. *Who Goes There? Authentication Through the Lens of Privacy* Stephen Kent and Lynette Millett, editors, Committee on Authentication Technologies and Their Privacy Implications, US National Research Council of the National Academies, National Academies Press, 2003.

[2]. *Challenges to PKI Deployment* Stephen Kent, Asia PKI Forum, Shanghai, July 2004; http://asia-pkiforum.org/july_shanghai/2004July/(4)Challenge.ppt

[3]. *Position Statement on PKI of the Australian Security Industry*, Stephen Wilson, Australian IT Security Forum (AITSF) V3.0, November 2003 www.aitsf.aeema.asn.au/resources/doc/documents_10.pdf

[4]. Medicare smartcard homepage www.hic.gov.au/yourhealth/our_services/medicare_smartcard.htm

[5]. New Queensland Driver Licence homepage www.transport.qld.gov.au/new_driver_licence.

[6]. *Global Platform Implementation Overview* November 2004; www.globalplatform.org/uploads/GlobalPlatform_NOV_2004.pdf.

[7]. *HealthConnect Business Architecture V 1.9 – Specification of HealthConnect Business Requirements*, Commonwealth Department of Health and Ageing, November 2004.

[8]. *HealthConnect Business Architecture V 1.9*, Commonwealth Department of Health and Ageing, November 2004.

[9]. *SSL defeated in IE and Konqueror*, Thomas Greene, *The Register*, Aug 2002; www.theregister.co.uk/2002/08/12/ssl_defeated_in_ie.

[10]. *Securing NHS Care Records* e-Health Insider, 16 December 2004 www.e-health-insider.com/comment_and_analysis/index.cfm?ID=41.

[11]. *Biometric Product Testing Final Report*, Report X92A/4009309, Biometric Test Programme, UK Government Communications Electronics Security Group, 2001 www.cesg.gov.uk/site/ast/biometrics/media/BiometricTestReportpt1.pdf

[12]. *Biometric Security Concerns*, UK Government Biometrics Working Group, September 2003.

[13]. *Using Biometrics for Border Security*, Report GAO-03-174, United States General Accounting Office, November, 2002; available at www.gao.gov.

[14]. *The Myth of Fingerprints: A forensic science stands trial*, Simon Cole, Lingua Franca 10(8) pp 54-62, 2000; http://fp.bio.utk.edu/evo-eco/resources-this_semester/Cole-fingerprints.pdf.

[15]. *DNA fingerprinting sparks fresh worries; discoverer says genetic databases could be misused*, Jill Lawless, The Associated Press, 8 Sept 2004; www.cbsnews.com/stories/2004/09/08/tech/main641998.shtml.

[16]. *Body Check: Biometric Access Protection Devices and their Programs Put to the Test*, Lisa Thalheim, Jan Krissler and Peter-Michael Ziegler, c't magazine, November 2002; www.heise.de/ct/english/02/11/114.

[17]. *Impact of Artificial Gummy Fingers on Fingerprint Systems*, T. Matsumoto et al, Proceedings of SPIE, 4677, Optical Security and Counterfeit Deterrence Techniques IV, 2002.

[18]. *Fun with Fingerprint Readers*, Bruce Schneier, Cryptogram, May 2002; www.schneier.com/crypto-gram-0205.html#5.

[19]. *Personal Identity Verification (PIV) for Federal Employees and Contractors*, FIPS PUB 201 V1.0 Public Draft, US National Institute of Standards and Technology, 20 Dec 2004; http://csrc.nist.gov/publications/drafts/draft-FIPS_201-110804-public1.pdf.

[20]. *Telltale prints*, Stephen Wilson, New Scientist, 14 August 2004; www.newscientist.com/article.ns?id=mg18324604.200.

[21]. *Proposed National Health Privacy Code*, National Health Privacy Working Group, Australian Health Ministers' Advisory Council, August 2003; www7.health.gov.au/pubs/pdf/code.pdf.

[22]. *Submission on the Discussion Paper "National Health Privacy Code (Draft)"*, Stephen Wilson, SecureNet Limited, 2002; www7.health.gov.au/pubs/pdf/46.pdf.