Misplaced **Trust**

Distributed ledgers have a role to play in identity verifications, but there are potential pitfalls, write Professor *Michael Mainelli* and *Vinay Gupta*

Distributed ledgers are increasingly touted as the answer to the identity problems plaguing finance and government. They may well be part of the answer, but more important is recognising they are only tools to help us tell the identity 'stories' we need to make finance and society work.

People talk about 'finding themselves' or 'losing themselves'. We talk about friends who have 'changed out of all recognition', but the restaurant happily believes they are the same person as their valid credit card. For people that do find themselves the result may be a higher quality of life and a less turbulent mind, less self-doubt or similar benefits, but they never quite seem to be able to explain what it is that they found, much less how a third party might validate it. You probably don't know who you are. And if you do, you probably can't tell me about it in a way that communicates it to me.

There are a lot of identity projects underway trying to provide new hardware, software, systems, and organisational approaches to identity. Identity is expensive for financial services where know-your-customer, anti-money-laundering, and ultimate-beneficial-owner requirements cost a lot. Yet the answer to "what is identity?" is obvious. Governments, police, financial institutions, credit rating firms, social networks, copyright lawyers, probate lawyers, websites and blogs all know it. Obviously identity is …

- your physical body;
- a government identifier such as a tax or social security number;
- status as incapacitated, a married couple or triple, a minor who can drink, enlist, vote or drive;
- your bank, mortgage, insurance, or credit card account;
- the network of people you know and the nature of those relationships;
- your community reputation;
- your capability to do something, a skill;
- the sum of your decisions made and your propensity to make new ones;
- proof of memory, such as a username/password combination;
- a cryptographic pair of public and private keys;
- a continuous and persistent sense of self or selfhood, "I";
- a metaphysical construct, a "true self" or soul;
- proof of intention, what you wanted to happen after you died;
- a story woven around a persistent image of a human being.

Temporary answers like "I am Pat's

mother" can be factually true but miss the core of a person's identity. The sum of all the roles we play seems to fall short of describing the totality of our identity. There is something in the whole which is beyond description. And we're supposed to write software about this???!

The story may not even be created by us. 'Linux Linus' doesn't refer just to an engineer; 'Linus says' is not a mid-40s engineer speaking, but a structural role. Surveillance also creates identities. Google, Facebook, and Twitter tell advertisers they know us better than we know ourselves. The Snowden revelations show that national security agencies have access to most or all of our social media, web surfing habits, financial and health records and more. But do they know us?

This leads back to self-knowledge. Is information about us which we don't know part of our identity or someone else's story? Is unconscious knowledge or behaviour part of our identity or part of someone else's analysis? A classic example might be a card player's 'tell' or gait analysis – an unconscious behaviour that gives your intentions or mood away. Identity is full of half-understood, half-baked theories yet enormously powerful organisations use identity every day to answer our questions – "can I have a mortgage, please?"

The convergence of interests among financial institutions, national security agencies, and advertisers is becoming one of the most unholy pacts made in our time. To get out of this hole we have to find new ways of paying for services and new ways of identifying ourselves. The most exciting changes are those provided by mutual distributed ledgers which seem to be providing new identity systems where the users own the data, where the data is universally consistent, and where the data cannot be destroyed.

Let's recognise two poles in the identity debate: humans as things versus humans as changing, narrative beings. Biometrics has an appealing simplicity. The temptation is to nail people to the meat. DNA is a seductively perfect biometric, but DNA gets everywhere that cat hair gets, and further. And when the stakes are high enough it's worth breaking in for your nest egg.

In practice biometrics has problems of false positives and false negatives and irrevocable release of information. But even were these overcome, we want different identities for different communities. Suppose my DNA says I'm female but I've changed gender? Suppose I am a closeted gay member of a church which is rather hard on gay people. If I use the same biometric, e.g. facial recognition, for both my church and my gay dating sites, I'm open to blackmail and exposure. People need to consider what stories the world really needs to know about us, what specific agencies really need to know about us, and do our stories for certain purposes 'add up'.

Mutual distributed ledgers are designed to be pervasive, persistent, and permanent. They can store the data to support our stories for ages – "look at my life history, I'm a responsible financial individual." The person most interested in the data is the person it's about. That person needs to be convinced it is their data and that they control who sees their data (for how long, how many times, for what purpose).

Solutions must be person-centric in order to have a stable point which will survive the extremely rapid rates of change which are so much a part of our twenty-first century. Put the person in the middle of the story, and work out from there. In a world of constant form-filling, the person also needs to be convinced that any new identity system will gain wide coverage and persist, that it will be worth taking the time and effort to join a new system. We might suggest that successful identity systems based on mutual distributed ledgers will give people control and distinguish actions from profiles.

On a distributed ledger I can present you with proof of what I have done without any third party having to vouch for me. Distributed ledgers provide objective permanent record storage without creating a data provision intermediary like a driving license office or MasterCard. Plus we get things like anonymity, pseudonymity, and public key infrastructure for figuring out who we're talking to without having to meet them in person.

The model of identity encoded in distributed ledgers is cryptographic keypairs: you have a secret (key) which you can prove to other people without ever revealing it. So we can prove (mathematically) that a person who knows a secret (X) is the person who (for example) signed both document A and document B. We cannot tell anything about this person – not a name, not an age, not a face – without additional information. If we link your real name to a secret key, and then put all the actions taken with this secret key into a distributed ledger, you have essentially zero privacy. But if we have a wallet full of keys that you control to boxes with your data, and these keys can be time-limited and box-limited, then you have highly selective privacy.

What goes in the distributed ledger isn't the data or even the hash of the data. Instead, we store a signed statement from a seriously credible source, or the necessary information required for a zero knowledge proof. We might often store a public key and a set of statements-and-signatures on a public key. This is a vision of the distributed ledger: something that's a little less about storing the entirety of our being on a permanent record, but more focussed on disclosing the minimum required for people to be able to do business (or other critical functions) together.

Maybe information is partitioned: your school grades, but not your name, all separated with a zero knowledge proof. If you have to prove that those are, in fact, your grades – well,

that's what that public key is for. You demonstrate possession of the relevant private key – you know the secret – without revealing the key. Every piece of personal information about you is stored with a different public key: you have a big fat keyring, and each key reveals only a single fact.

Such systems do exist. Both authors are working on systems that do give ownership of the data to the person. In the extreme, a state authority asks an institution for a copy of all the anti-money-laundering information used to make a decision to accept a person's deposit and they can say, 'we don't own that data, go ask the person. "But they're not cooperating." "That's your problem."

People will leak their keys and there will have to be schemes to kill and replace keys while maintaining the old identity. These schemes are 'key revocation certificates' and can be managed much like a backup copy of your private key. Once your revocation certificate is published to the ledger, your key is dead. People will continue to lose their keys, and there will have to be schemes which allow recovery (e.g., key split 100 ways among your social circle, with 75 pieces required to reassemble it). Perhaps most important, we need simple metaphors so people don't give away 'master' keys lightly, know what a thrice-time use key is, and understand how the system works as easily as they understand locks on doors. It won't be easy. Research and experimentation on the user interface for wallets of keys (the big fat keyrings) is now one of the most difficult problem areas.

Who are you – a complex construct inside of the head of a naked ape that looks at the night sky in wonderment or the sum total of your mortgage payments? Actions are life decisions, like signing an employment or mortgage contract, taking out citizenship, or making a will. Actions are also smaller decisions, like recording biometrics or locations at points in time. Action-based identities are very, very rigid things, and represent a credential nearly as hard as a State ID, and often more useful.

Profiles are a set of opinions that somebody has about me. These might include things like medical records, which are stories about people. We have little idea what is in our medical profile, are rarely qualified to validate it, and might not even know it exists. And then it might be filled with utter stuff and nonsense. Your profile is your business, even if it is about me. My profile about myself is a special case: a set of things I say are true about myself ('I like cold beer') which (still) may or may not be true.

A dangerously hot area is merging profiles, e.g. credit records crossed with social media to give even more accurate forecasts of loan risk. The problem is that each profile contains within it a set of assumptions (a hidden frame, a perspective) which is incompatible with much of the other data in a combined profile.

The distinction between profile and action is important in ensuring that identity decisions are kept close to the data. A combination of profile and action is powerful. If an institution 'co-signs' data we might call this a 'privileged profile'. This is a profile stating that I (as the person this profile information is about) currently consider it to be correct, and sign along with an institution. The combination of privileged profile and action records is pretty much unbeatable: "they say I got a law degree, and here's my evidence that I've paid for one."

The purpose of identity systems is to provide data to make a decision. Give, or not give, access, services, credit … All identity-based decisions involve risk. Where ever there is risk, then someone can provide indemnity, i.e. insurance against a bad identity or a bad decision based on identity. The person who pays a mortgage for 18 years is a tangible person with thousands of pounds proof of financial trust. The cumulative impact of multiple historical payments constitutes a compelling fact, but the next payment may not be made for any variety of reasons from fraud to loss-of-job to death. These are all games of odds. If the permanence and flexibility of distributed ledger identity systems leads to wide take-up, there will be vast new markets in the systems themselves, the institutions that provide 'privileged profiles', and those that provide indemnities or insurance against mistakes.

What this leaves us with is a tentative path forwards. Put some things into distributed ledgers, particularly things like the public keys of major organisations, and certificate revocations for people whose keys have been stolen or whatever. Use additional cryptography to provide things like zero knowledge proofs of age without revealing identity. Compartmentalise information about our various roles in ways which make it hard to join profiles without our consent.

The result of this 'wise as a serpent' approach might be a mixed ecology of identity solutions. But this slight confusion is worthwhile. We clearly need new identity infrastructure, but we also need time to experiment, scale experiments, and understand where we are in the ever-changing terrain which so much marks our current lives. With mutual distributed ledgers, what is not disclosed can be disclosed in future, but what is published now is published forever, encrypted or not.

Let's step forward on identity, carefully. ***BT***

About the authors

Professor Michael Mainelli is Executive Chairman of Z/Yen Group and Principal Advisor to Long Finance. His latest book, The Price of Fish: A New Approach to Wicked Economics and Better Decisions, *written with Ian Harris, won the 2012 Independent Publisher Book Awards Finance, Investment & Economics Gold Prize.*

Vinay Gupta is an architect with Consensys Systems and the Ethereum Foundation.