



An Exploration of Identity

Ian Grigg

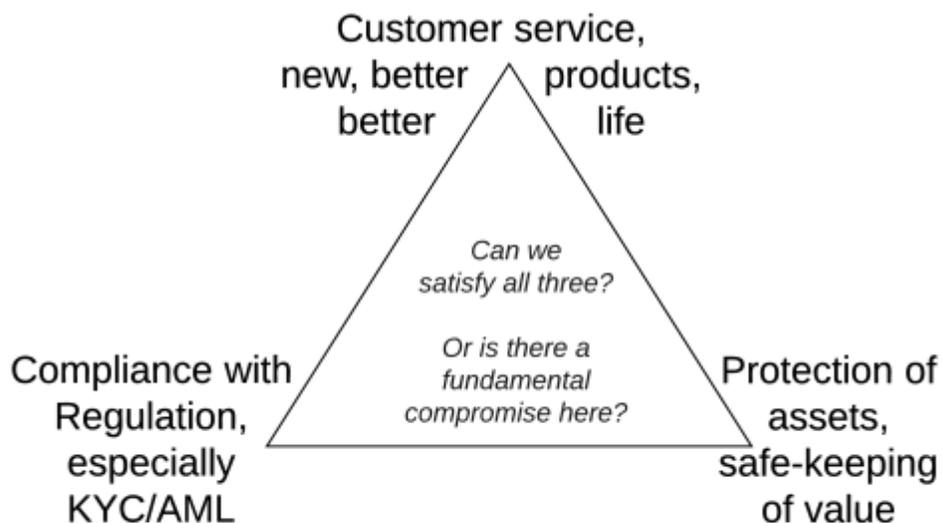
March 3, 2017



Three motivators for identity

It seems as if there are three general motivators for the commercial notion of identity: compliance, security and customer service. Knowing which person is whom is a compliance issue. Knowing that a person authorised use of his own funds is a security issue. Knowing that a person was able to move her money easily without obstructions is a customer service issue. Sometimes these three motivators are compatible and sometimes they are in conflict.

The Financial Identity Trilemma Syndrome



FITS – the Financial Identity Trilemma Syndrome

If you're suffering from FITS, it might be because of compliance. As discussed at length in the White Paper on Identity, McKinsey has called out costs of compliance as growing at 20% year on year. Meanwhile, the compliance issue has boxed in and dumbed down the security, and reduced the quality of the customer service. This is not just an unpleasantness for customers, it's a danger sign – if customers desert because of bad service, or banks shed to reduce risk of fines, then banks shrink, and in the current



environment of delicate balance sheets, rising compliance costs and a recessionary economy, reduction in bank customer base is a life-threatening issue.

If we can redress the balance - put customer service back up on the list - without compromising the other two, then not only might it be worth the customers' while to invest in a new system, it might actually save a few banks.

How might we do that? From the perspective of our current understanding of identity, it's somewhat clear that today's systems are as much the problem as they are the solution. It then behoves to reexamine the process from scratch.

Let's start from a clean sheet of paper and ask – how do we make decisions?

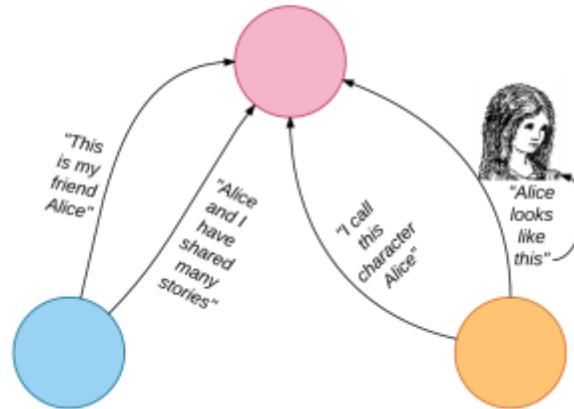
Context means everything

We humans can make quick assessments, unconscious ones, from context.

If for example, all of our communications with a group of people are in a closed network of friends or known coworkers, we can make assumptions about what we can share with that group, which we would not make for a public forum. Or, outside work, if we found ourselves in our local bar, we might assume that all the people in the bar are likely as honest and reliable as the average person in our city – in one city we might leave our phone and keys on the table while visiting the conveniences, but not in another.

In the above examples, when deciding whether to trust me to not steal your phone, you don't need to know my identity at all. What you need to know is, where am I? Which set of norms do I subscribe to?

Fig 3. Identity is an edge protocol



Taking this across to the context of banks and transactions and so forth, it may well suffice to know that I'm a trader with that famous power-broker Broker A. That might be all you need. OK, it's not quite all, you'd also want to know which desk I'm on, what my limits are, and what I'm authorised to trade. But what you don't need to know is my name.

When we limit ourselves to corporate banking – context - then we're interested in the two identities - the trader and the company, but what we're really interested in is not the identities per se, but the relationships and interplay between the identities. Let's think about this as an edge protocol (see Figure 3 above from the earlier piece on Edge Protocols) and apply it to trading: everyone says Broker A is a great broker, Broker A says that I am a trader there, and that should be enough.

If we do business with the Broker A at all, we're basically riding on that corporation's trust level, and shouldn't need to worry that much on which of the many people in the company we're talking to - they all should be good, else why deal with that company?

If it was just about edges, then we could just collect them up, analyse them and we'd be liquid. Connect an edge Hoover to a relationship AI and we'd be done. Utility, outsourcing, profits, here we come!

But there are a few impediments to this process. Let's look at three, being risk, reliability and liability.

The Facts of Others

Ideally, we'd like to take one bank's decision over a person and copy that, just like data, across to other banks. But risk analysis precludes that -- one bank's risk is not the same as another bank's risk. We therefore need to stop short of outsourcing decisions, and in today's world, we are limited to out-sourcing the gathering of facts – relationships or edges.

The Hunt for Facts

Let's put the spotlight on the fact. A typical fact would be that Bob signs over this statement:

"The passport held by Alice has her name and a good photo of her" /Bob/

This is evidence, to some value, that Alice is Alice, and if relying party Carol needs that evidence, she might be happy to be able to rely on Bob's statement.



Or she might not. This could go wrong in many ways but let's say we've filtered out the non-useful facts and what we are left with is the golden nuggets of trade. We're still left with:

The Nouns...

What's a passport, anyway?

What's in a name? A rose by any other...

What's a good photo? Fierce, pretty or bland?



Who is Bob?

Why does he care?

What incentive does he have to tell the truth?

What incentive does he have to do a good job?

Is the fact within reliable?

Was it reliable then, is it reliable now, will it be reliable tomorrow?

Chief amongst our concerns is that the edges as collected might be unreliable. I.e., what happens if my name is not Alice? Or I work next door to Broker A and just snuck in to use their terminals one night? Or, any one of a thousand scenarios - the conclusion here is that while the fact might still be a fact, in that Bob said those words, it might not be an accurate or reliable representation of the real world. If such happens, then the golden analysis above turns to fool's gold.

The Source of Our Unreliability

The reasons that any given fact might be unreliable are legion – I once wrote down 99 of them. We could dutifully promise to take more care, but this hasn't worked well in the past. We need better than marketing and empty campaign promises to make this work. Luckily we've got some more clarity on why a fact isn't reliable and how to fix it. Four techniques will create a foundation for the facts:

1/ Skin in the game – every agent needs to be not only positively incentivized to work in this relationship building, but also positively corrected when things go wrong. There needs to be what the engineers call a negative feedback loop – one of correcting capability.

2/ Quality control – if the above correction is dramatic, we need a way to show that the agent has done a good job. Statements are in words and they can be both wide of the mark and misinterpreted. To address this, we can set up-front minimum quality standards that are clear and applicable to both the makers of facts and to the users of the facts, or “relying parties,” and operate them diligently.

3/ Redundancy in sources – to get a single complete accurate fact is very expensive, but to get many small facts converging on the same approximate large truth is cheap.

4/ Authorities in sources – some facts are ‘owned’ by certain parties, and if we can get them on board in a secure context then we can improve the quality, at least in their facts.



These should be familiar, and will create a base of reliability but we need more.

Liability of the provider sets the quality of the facts

To be useful to Carol, we need the facts to be not only reliable, but exportable. That is, suitable for other people to rely on the facts in their own risk assessment. We can do this in the first two basic ways as listed above:

Firstly, by setting liabilities for poor work, especially but not only by not following the standard of the second way.

Secondly, by setting up front and operating to a minimum quality standard that is clear and applicable to both the makers of facts, and to the “relying parties” of the facts.

Imposing liabilities for poor work needs to be done carefully because there are two general possibilities, being

- the work and care that is done in creating the fact, which has one value, and
- the damage that can result from relying on the fact, which damage has another value.

These two values are sometimes wildly different.

In general, it is harder to assess liabilities to the damage that can result *in advance* because it is implausible to predict to what use the facts are put. This is what the cryptographers call the problem of Grandma’s house: if I sign that Mallory is a good guy, and Grandma relies on my statement, but the result is that Grandma loses her house, who’s to blame? One school has it that she fairly relied on me, so I have to pay her a house. Another school has it that because I only reviewed Mallory’s passport, it is a process or administrative implication and no more, and back to passport-review school I go.

Which is it?

To cut the Gordian knot, we typically place such problems before *a resolver of disputes*, a person who can decide which of the interpretations apply. This accepts that we are entirely uncertain how a given dispute will pan out, both I and Grandma, but we know that it will resolve one way or another. So, and this uncertainty is the crux of the argument, I will probably do a better job than merited because my liabilities may go sky-



high, and Grandma won't put her house to the gamble of one claim, because her assets may go to rock bottom!

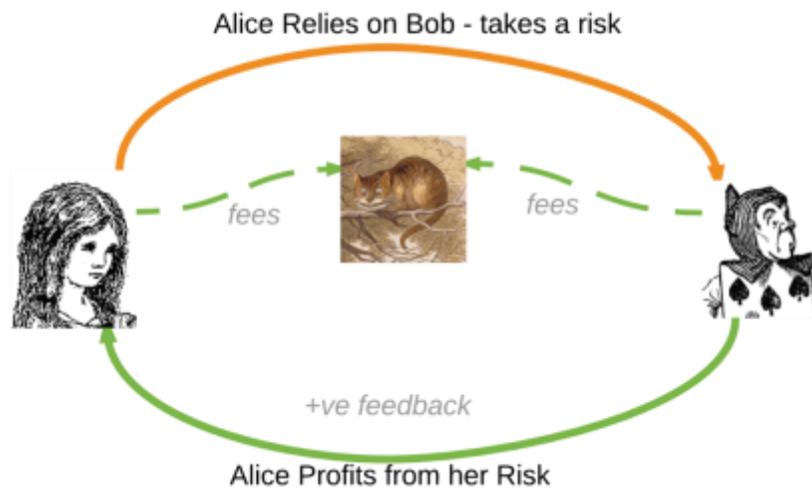
But, if my liabilities could go sky-high, why would I ever get involved? It is for this reason that I need to be protected by a standard approach – one that is well thought out, agreed, documented and *auditable*. Especially, that last step is what will convince an Arbitrator that I have done the job I was asked to do.

A good liabilities framework then is initially limited to the correctness of the facts. However, in order to get any traction on relying on those facts, we at least need to improve the quality of the facts such that they are reliable – they meet a minimum bar to allow others to rely on them.

Hence, while the liability solution is initially necessary to address the liabilities incurred when one person relies on a fact produced by another, it has another side-effect - it improves the quality by encouraging the provider of the fact to take especial care up front, *as if* they are liable to another person not as yet identified for risks not as yet quantified. For this reason, we need to protect the provider of the facts with a standard to follow, so that they are not on the hook for impossibly high liabilities from a simple process operation.

The alternate route lacks accountability

In the typical alternate approach, the provider of facts asserts zero liability, as that is the only business solution to unpredictable liabilities. But, this is what we call a positive feedback loop, in which the provider gets paid for good and bad results, equally. In a positive feedback loop, activity grows and grows until the machine destroys itself. As there is no correction when the machine goes off the rails, a lack of liability also means a lack of accountability, and in that scenario, there is an unfortunate consequence: the quality of the data shrinks to nothing. In effect, the zero-liability solution causes a *race to the bottom*, and the provider prints unreliable statements without limit.



We need to solve the liability issue not only because of the direct liabilities themselves but also because the system needs an appropriate level of quality, and to deliver that, the provider of fact needs an appropriate incentive to reach that level of quality.

In short, we need a feedback mechanism that convinces the provider of facts it is worth taking real care as opposed to advertising care.

An example of adverse liability consequences can be seen in the Certification Authority (CA) business. The provider of facts, the CA, typically says two things about a corporation, being that R3 for example

- holds the private key that signs the public key that is identified, and,
- is the holder of a domain name, e.g., *r3.com*.

These two facts are memorialised (if not explained) in the certificate.

But the provider of facts, the CA, disclaims all liability. So, in consequence of that disclaimer, if a bank were to do a Corda transaction that relied upon the certificate to (for hypothetical example) download the latest Corda release, and the bank got man-in-the-middle attacked by the Closet Internet Anarchists which injected a zero-day and then led to a Bangladeshi-style sweep of all the bank's money via the Philippines....

Then...*who stands up? Who carries the liability?*

Not the CA – its contract (buried deep) says "zero liability" and it means it – for various and many reasons we cannot take the CA to court. The CA is entirely covered, no matter what goes wrong. And, therefore, eventually, inevitably, by all the laws and



history of governance and economics and erosion and evolution and all such scientific notions, the CA takes low levels of care to check the data being asserted into the certificate. The CA has entered *the race to the bottom*, and we're not the winners of this race.

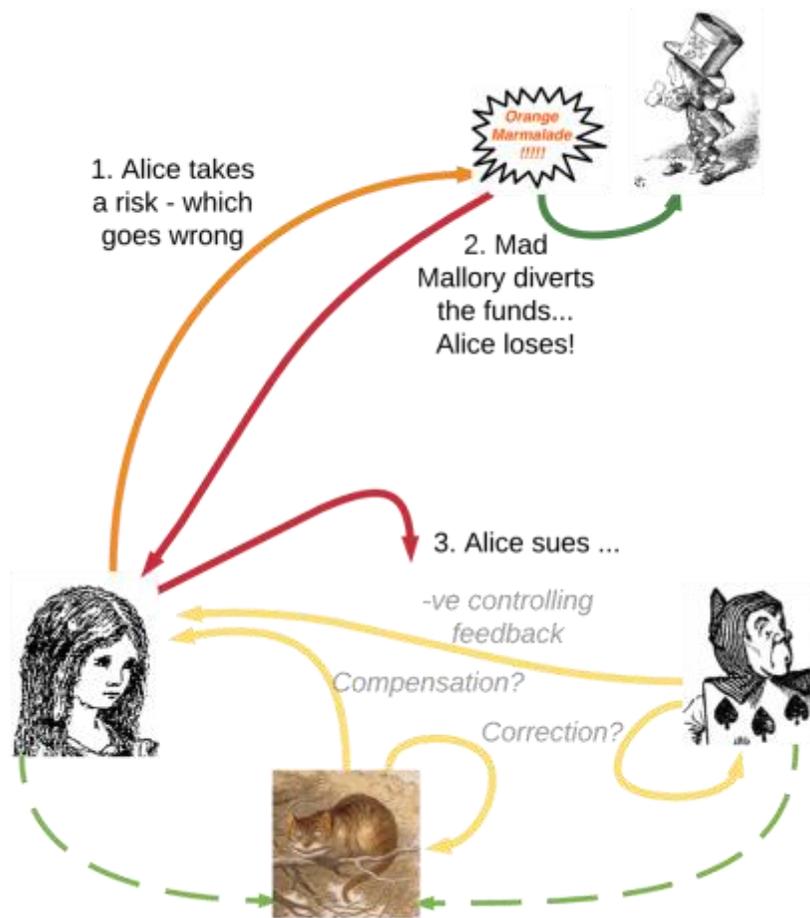
The CA can only have an appropriate level of quality if there is a feedback loop that acts to move the quality up or down as appropriate to the wider uses made by customers. If the liability is dumbed down in some fashion - if some group of players is insulated from the effects of their actions - then the quality sinks down over time and the entire system becomes useless. This then eventually results in a disaster such as fraud or hacking or phishing, which then triggers either a bureaucratic overdrive to add more and more pointless feel-good rules, or a regulatory kickback. Or both.

Closing the loop

It isn't about the bank, or Alice, or the regulator, or the CA or nature or desire or humanity.

It's about *the system*. We need a feedback loop that controls the quality of information, one that the engineers call a negative feedback loop.

The nature of this information / this data is that disasters are unpredictable - there is no way to add a dial *ex ante* that allows the setting of liability to some level. We all take on some risk, each of us, individually and society-wide, what we need is a mechanism for controlling the risks when they blow out. Hence, we need a dispute resolution mechanism that can examine the disaster after the fact, take on the question of wider liabilities within the context of a standard, and return a human answer when the data fails.



We need both Alice and Bob to not fall into the trap of moral hazard – hoping that some mysterious other covers them for all eventualities. We need both of them to take some care, and be prepared to stand up when the care wasn't sufficient. We also need to wrap this up into an efficient package, such that they are incentivised to participate.

If Alice and Bob and everyone else has reached consensus on participation, then the edges will grow and flow. And with enough edge liquidity, the task of decisions over a node becomes tractable, even across the globe, across languages, across jurisdictions.

Then we can get back to the business of creating profitable, shared trade. Deals that are backed by a fabric of identity, built of the edges of human relationships.