

BeefLedger Working Paper Series No. 1

Provenance Value: Supply chains, blockchains and BeefLedger - issues and design considerations

Extended Version of Paper presented at Supply Chain on Blockchain Conference, Brisbane, 15 July 2019.

Warwick Powell (Chairman, BeefLedger Limited) (This is a work-in-progress draft; please do not reference without permission)

Communications with author: wp@beefledger.io

If everybody always lies to you, the consequence is not that you believe the lies, but rather that nobody believes anything any longer -
Hannah Arendt

What is needed is an electronic payment system based on cryptography instead of trust - **Satoshi Nakamoto**

Introduction

This paper sketches out some elements of a design frame that has guided the development of BeefLedger as a specific application of blockchain technologies to a particular supply chain context. I do this first by reflecting on four propositions about supply chains and blockchain technologies in general. Our approaches and lessons are generalisable to the extent that they provide conceptual and practical tools to help frame the design and deployment of technologies in the name of better social, ecological and economic outcomes in a broad range of contexts. I then illustrate with some example application cases, in which I introduce the idea of **provenance value**.

Propositions

Here are the four propositions.

1. Supply chains in complex financialised economies largely **do not presuppose trust to function**. This is because transactions - the flow of goods and services on the one hand, and the flow of funds on the other - typically involve ‘strangers’¹ and institutional functionaries. Rather, **the condition of existence of functional supply chains is transactional dependability in conditions of zero trust**. A zero trust environment is characterised by the presence of *fundamental uncertainty* as opposed to simply the presence of calculable *risk*². Much discussion³ claims that there is a significant social cost to establishing trust for commerce, via mechanisms and institutions such as “the rule of law”, and various associated arbitration and court systems and such like. I would suggest that, on the contrary, these are better understood not as the costs of establishing trust but as the costs of dealing with the consequential uncertainty of outcomes of non-trust and the costs exacted upon particular parties due to counter-party non-performance or change of circumstances. The existence of legally enforceable contracts, governing supply chain transactions, that include punishments in the event of non-performance, is symptomatic of a zero trust environment.⁴ If parties trusted each other, there would be no reason for such contracts. Blockchains in supply chains will not hasten the demise of legally enforceable contracts.
2. The deployment of blockchain technologies in supply chain ecosystems is **not about instituting trust**. This is unnecessary. Rather, it is about increasing **dependability in conditions of zero trust** (at best), and distrust / mistrust at worst. In this context, the extent of **dependability refers to the presence or otherwise of information asymmetry** between and amongst participants as the basis of actor decision-making and conduct and the likelihood of actor behaviours and behavioural outcomes to meet required conditions. Should

¹ Peter Murphy has argued that a hallmark of successful economies is that they involve the ability for strangers to engage in commerce. See Murphy, P., (2015) “The Stranger Society: The Case of Economic and Social Development in the Tropics”, *Budhi* 19.2 & 19.3 pp. 107-134.

² On this important distinction, see Keynes, J.M. (1921), *A Treatise on Probability, The Collected Writings of John Maynard Keynes*, Vol. VIII, London.

³ For example, Davidson, S., et al (2018) “The Cost of Trust: A Pilot Study”.

⁴ There are some similarities in this argument to that advanced by Williamson, O.E., (1993) “Calculativeness, trust and economic organization”, in *Journal of Law and Economics* XXXVI (April) pp. 453-86. For Williamson, trust is reserved for the economically trivial, non-commercial and personal relations. For him trust is a passion rather than a modality of economic action.

trust emerge at all via these dynamics and interactions, it is a by-product.⁵ Additionally, information asymmetry can be understood as an embedded property of supply chains involving multiple parties of varying configurations of power relations, which is a precondition for **capricious action**. Industry structure is a key institutional and market condition of existence of real supply chains, wherein capriciousness is always present-in-potential.

3. Informational dependability **is not about the truth** in any epistemological sense. Rather, it is about the establishment of a body of **common knowledge** upon which actors can go about their business wherein the downside risks of action taken on the basis of asymmetric information relations is mitigated. We can dispel with the myth that deploying blockchain technologies is about delivering an epistemological truth. This is an important point, because this also deals with a range of associated conflationations that have resulted in a misdiagnosis of blockchain's opportunities for better supply chain outcomes (the hype) and misplaced criticisms (the anti-hype).
4. **Transparency is not a precondition of trust**. It is the reverse. Transparency is a condition of existence of dependability in conditions of zero trust. Where there is trust, transparency is no longer necessary. When one's word is as good as one's bond, there is no requirement for transparency beyond the promise. Transparency is demanded because there is no trust. For one party to "just say so" is not enough for the other to accept the claim or the promise. Where there is, at best, zero trust, actors demand a level of transparency that is not necessary where there is trust. Blockchains are dependable in conditions of zero trust amongst a network of strangers for the very reason that each stranger is both watching everyone else and is also being watched by everyone. (Again, if trust emerges then that is a by-product.) This is the Panopticon effect of blockchain consensus algorithms.⁶

These are possibly provocative claims, particularly when much discourse on blockchains and supply chains revolves around ideas of re-instituting trust, and delivering transparency and truth. I would humbly suggest that blockchain's potential lies not in this unworldly Nirvana, which sets a

⁵ We adopt the approach to trust as developed by Solomon, R., (2000) "Trusting" in in Wrathall and Malpas eds, *Heidegger, Coping and Cognitive Science. Essays in Honour of Hubert L. Dreyfus Volume 2*, which treats trust as an ephemeral condition produced through human interactions. This is contrasted to relationship between unfamiliar humans and between humans and machines, which are not trust-based. Not unrelated is the idea that trustworthiness *is a result of* safety and dependability, as articulated by Roth, A., (2015) *Who Gets What - And Why* p. 116. Market design, therefore, must be about safety and dependability for participants. Trustworthiness is a residual consequence. Sabel, C., (1990) *Studied Trust: Building New Forms of Collaboration in a Volatile Economy*, introduces the idea of "learning by monitoring" as an explanation of successful collaboration. In this case, trust is the consequence of learning embodied in ongoing monitoring.

⁶ The Panopticon is a prison design proposed by Jeremy Bentham. The circular structure positions a single guard in the central tower, with an outer ring of cells facing the guard tower. Prisoners in the cells cannot be sure they are being watched from the guard's tower; but given the prospect that they are being watched, prisoners will tend to regulate their own behaviour as if they are being watched. The word panopticon derives from the Greek *Panoptes*, meaning all-seeing.

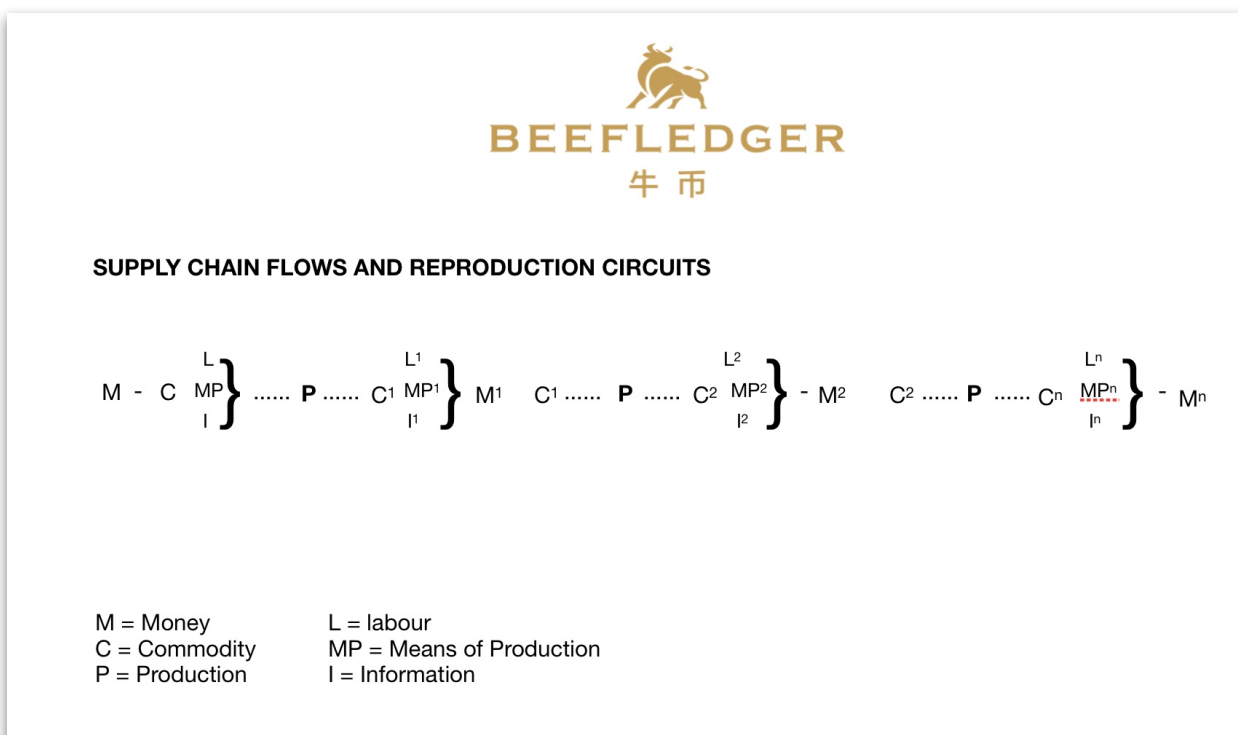
very high bar indeed, but in the less ambitious possibilities of improving supply chain outcomes without the need to presuppose trust or truth.

Supply Chains

Supply chains involve economic agents interacting across time and space to transform and create things, ultimately for consumption. A division of labour is a precondition of existence of a supply chain. These agents are connected at a material level in terms of the movement of things from one to the other as the processes of transformation take place. Agents are also connected in a non-material sense via the movement of monies in exchange for the movement of goods. The movement of money is governed by the flow of information about the things, because money doesn't flow unless certain informational conditions are satisfied. The financial relations between agents in supply chains are characterised as credit-and-debt relationships. In financialised supply chains, this usually involves a financial third party to mediate payments.⁷

The following depicts the supply chain flows in a financialised system:

FIGURE 1: FINANCIALISED SUPPLY CHAIN SCHEMATIC



⁷ See Nakamoto, S., (2008) *Bitcoin: A Peer-to-Peer Electronic Cash System*, in which he discusses the requirement of financial third parties, and posits the bitcoin blockchain as an alternative means by which financial transactions can be executed on a peer-to-peer basis. In strictly financialised terms, supply chains can be conceived as a series of discrete balance sheets linked to each other via relations of debt and credit.

... in which money (M) is exchanged for a commodity (C) which is then transformed through a production process (P) to become a new commodity (C¹). This new commodity C¹ is then exchanged for money (M¹), and so on. The scheme draws from Marx's *Capital* Volume 2, though I have added explicit notation to deal with the role of information in supply chain flows.

A supply chain has two purposes, so to speak. The first is reproduction. In modern financial economies, supply chains and their constituent agents are in the first instance driven to reproduce themselves. In accounting terms, this means that each agent must recover its costs. Costs should involve depreciation and therefore replacement cost of capital consumed during the processes of production. The second purpose is accumulating profit above and beyond its reproduction costs.

As much as there are debates about the *source* of surplus value, for our purposes I propose a simple parsimonious “social accounting” approach that views the financial surplus in terms of the difference between (a) the total aggregate revenue of the supply chain measured in terms of the price the end consumer pays for a given unit of product, and (b) the costs of reproduction of the chain as a whole in producing that product. The financial flows within a supply chain effectively distribute the aggregate revenue to the various actors, wherein above-cost surpluses are incrementally captured along the way. Flows are neither smooth nor continuous; rather, they are episodic and lumpy. (For this exposition I do not discuss the possibilities of end consumers being manufacturers of data, which can be purchased by others.)

We therefore are addressing a circulation system that has three interconnected flows. These are:

1. The flow of products;
2. The flow of money; and
3. The flow of information about the products.

In this schema, it is necessary to note the role of information about the products in the transaction M-C itself. In this formulation, we could say that the hyphen denotes the flow of information.

The counter-position of product to monetary flows is mediated by the flow of information. This is information *about* the product. In our particular supply chain, this includes information about cattle breeds, sex, age, feed, weight and locations. It also includes information about meat quality, weight, and storage and transport conditions. Verification that this information complies with procurement requirements is the condition precedent of the exchange of products for money.

Note that we have said information *about* the product, which is distinct from the product itself. The word “about” here is crucial as will become clear later.

Because information about the product is the key to the exchange, it is also the locus of **capricious** conduct. As the definition (see box) indicates, capriciousness is characterized by sudden, unaccountable and unpredictable actions. Information censorship, mis-information and arbitrary construal of the meaning of information and its relationship to the products themselves are all “up for grabs” when actors are in positions to subvert the validity of data as a basis of common action - that is, the exchange. Asymmetrical relationships between different actors involved in a transaction, coupled with siloed data creation and storage systems, is the condition of existence of informationally driven caprice. In supply chain terms, such caprice could include delayed receipt of goods, delayed payments, refusal to pay in whole or refusal to accept the products. It could also include refusal to accept responsibility for losses, for example for spoilage of goods. In more benign terms, caprice could be described as ‘friction’.

capricious | kəˈprɪʃəs |

adjective

given to sudden and unaccountable changes of mood or behaviour: *a capricious and often brutal administration.*

• changing according to no discernible rules; unpredictable: *a capricious climate.*

DERIVATIVES

capriciously | kəˈprɪʃəsli | adverb

capriciousness | kəˈprɪʃəsnəs | noun

ORIGIN

early 17th century: from French *capricieux*, from Italian (see *capriccioso*).

The prospect of payments delay or partiality in payments has definite economic and financial implications on upstream suppliers. From a whole of chain perspective, there is also the prospect that such delays effectively allocate risk sub-optimally often to those least well positioned to handle these variabilities in cashflow. Cost of cash flow finance is high, and yet, such debt is non-productive. These costs are imposed on specific nodes whose risk profiles are usually treated in isolation of the whole-of-chain flows.

Because capriciousness is sudden and unpredictable, it is a form of fundamental uncertainty rather than a calculable risk. When information asymmetry enables capricious conduct, supply chain sub-optimality is an almost inevitable outcome. In other words, what is good for one actor is sub-optimal on the whole. Blockchains have the potential to address these uncertainties and reduce risks of capricious action.

Blockchain Attributes


I now turn attention to what it is that blockchain's are actually good at. In fundamental terms:

1. They are good at ensuring that the application data state is valid - that is, the state got there by virtue of some process that followed some explicit rules;
2. The process itself is transparent to the participating network of actors, so people can see what the rules are and how states transitioned from one to another;
3. The states as recorded are irreversible (you can't wind back time); and
4. The data is censorship and fraud resistant, meaning altering or blocking data cannot be done singularly and capriciously.⁸

These properties are achieved by a consensus mechanism underpinning the updating of data states, in which the mechanism's purpose enables the secure updating of a state in accordance with specific state change rules, where the right to perform the state update is distributed amongst some set of economic agents.

It is for these reasons that many have been drawn to the possibilities of blockchain data systems to deliver an immutable record of a supply chain. Here, ambitions around traceability are postulated as a core attribute of blockchains. Indeed, a data structure based on Merkle Trees and block hashes enables the “tracing” of any state via its link to previous states and so on (see Figure 2). This is a data traceability truism. It is useful for a host of reasons, predominantly enabling a *proof of data state validity*, and therefore a valuable attribute. In this sense, this is the ‘rearview mirror’ property of blockchains - they are a record of the past, just like all other ledgers that came before it.

FIGURE 2: MERKLE TREE AND BLOCK-HEADER STRUCTURE⁹



⁸ For the purposes of this paper, issues associated with collusion risk are not discussed.

⁹ Courtesy of David Harding <https://dtrt.org/posts/blockchain-pruning-problems-and-solutions/> accessed 12 July 2019

However, despite the merits of these attributes, they do not obviate the issues of data capture and the validity itself of the data submitter(s). This goes in the first instance to the 'rubbish in / rubbish out' problem, and secondly to issues about the relationship between the data domain and the analogue world. These two points sit at the heart of most critical reactions to claims about product traceability using blockchain technologies, because - so the criticism goes - tracing data does not prove that the thing itself has any of the attributes claimed by the data, and that, in any case, data on a blockchain may simply be immutable nonsense.

There are at least two issues here:

1. the identity and veracity of the identity of the data messenger - call it an **Oracle**; and
2. the validity of the information - or what I would describe as the difference between **epistemological truth claims** and notions of **common knowledge**.

Let's start with **Oracles**. Oracles are a source of data about things or situations. Think of them as collectors and messengers. This data are foundational as records of the past; and are also necessary to drive state changes via self-executing smart contracts (which I discuss later). Oracles can *prima facie* gain status by virtue of some form of *exogenous authority*, such as a certification from a body of one sort or another, or by an act of law. In this sense, they are *institutional fiat in nature*. Oracles can also be devices, actively or passively collecting data about certain things or conditions. They can, of course be both.

If the data source is exogenous to the blockchain itself, risks of false identities (imposters) and bad raw data exist. The relationship between the analogue or material world and the digital ledger is what is at issue here. For many, this relationship seems to be the locus of a great deal of confusion and concern. As objects in the two domains - that is thought or language on the one hand and the analogue real object on the other - are irreducible to one another, the data record about things or conditions are not tantamount to being the things themselves. Because of this, some blockchain critics argue that blockchains can't deliver "true provenance" because "true provenance" can only be delivered via forensic or analytical science. This line of criticism is a bit like ships passing through the night, for the reason that this proposition merely shifts the epistemological onus from one set of thought objects (data on a blockchain) to another (the lab test report). The problem is actually the epistemological claim itself, and neither the hype nor the critic can resolve the reductive conundrum.

Rene Magritte's famous painting "This is not a pipe" illustrates this issue well. The painting is, of course, one of a pipe. Under the image of the pipe, Magritte writes: "this is not a pipe". That is because it is actually a picture of a pipe, not a pipe itself.

FIGURE 3: RENE MAGRITTE'S THIS IS NOT A PIPE



We actually do not need to solve this problem for blockchain-validated data to be meaningful in supply chain authentication and state change processes that support the core purposes of supply chains, and which can address sub-optimal outcomes occasioned by capricious action. A range of approaches can be mobilised to deal with risks associated with identities and data validity. From a decentralised ledger technology point of view alone, we would look to things such as decentralised validation protocols for networks of devices (for instance), decentralised protocols for identity validation, and the introduction of cryptographic tools like Zero Knowledge Proofs to validate Data Oracles without compromising their status (and, therefore, opening them up to imposter attack risks). We also would consider introducing integrations between digital data protocols and materiality interventions aimed at mitigating risks of artifact imposters and substitution.¹⁰

¹⁰ There is no need to treat blockchains in isolation from other approaches that can assist in securing confidence in supply chains. Packaging is an area of risk and also risk mitigation. BeefLedger adopts a pragmatic complementary approach, likening different interventions as threads that can contribute to the weaving of a robust fabric.

We can use cryptography to prove that some Oracle (X) made a message. We can also use cryptography to prove that some Oracle (X) made a message after another message. Blockchains can prove “when” the message was made. It can also prove that X did not make a message. Lastly, blockchains can prove that some set of messages is the entire set of messages that some set of participants made. In other words, blockchains, combined with cryptography, can increase the extent to which people can have confidence that they aren’t being cheated. And if they are being cheated, they aren’t being cheated in isolation. This is because Oracle’s messages - to become meaningful within a blockchain ledger and supply chain environment - must assume the status of common knowledge.

Let us now consider **Common Knowledge**. Something (Y) can be considered common knowledge if:

- Everyone believes that Y is true;
- Everyone believes that everyone believes that Y is true; and
- Everyone believes that everyone believes that everyone believes that Y is true And so on.

Common knowledge is the basis upon which much coordination between supply chain actors takes place. If a body of knowledge is accepted by a consensus mechanism to be valid, it then forms the basis upon which those actors who participate in the consensus mechanism can go about making decisions and undertaking actions. Blockchains and their consensus mechanisms are the means by which messages created and delivered are either rejected or accepted into the ‘cannon of common knowledge’. Note that there is no presupposition that common knowledge has any epistemological properties at all. Indeed, it is theoretically possible that some common knowledges may well not be accurate at all. *If we err, we err together*, is the consolation in these circumstances.

If we can be satisfied that a message (data) was made and sent by an Oracle, and complied with the rules of the consensus algorithm, then the relevant socio-economic network - in our case, a network of supply chain agents - can proceed with certain decisions and actions on the basis that the messenger and message are valid. In so doing, the message and its contents assume the status of common knowledge. Satisfaction that the identity of the messenger is, in this environment, a question for the multitudes rather than something that can simply be asserted by way of *fiat*.

Common knowledge is not the same as epistemological truth or truths. Epistemology deals with the relationship between two domains; the domain of knowledge on the one hand (sometimes

described as the thought or language domain), and the domain of the thing-in-itself (the analogue world). A whole range of debates about how it is that we can know something about the world has enlivened philosophers through the ages. We don't need to revisit these debates to observe that at any particular moment and place in social history, some knowledge and processes by which knowledges are produced, have assumed the status of common knowledge and common knowledge sources while others do not. For example, we accept the procedures of genetic testing as a means by which a knowledge about biological objects can be generated; or that chemical trace analysis makes certain truth claims about provenance in a very specific sense.

That said, a laboratory-as-Oracle cannot be taken as given. In the first instance, we have an identity issue. Fortunately we have tools to tackle this quite effectively, as I have suggested. The second issue is that the knowledge that the laboratory produces in fact comes in the form of an **artifact**, which emerges through a set of procedures utilising a range of equipment, operated by certain persons. *The artifact is the lab report*. All of these techniques, implements and agents are situated in specific socio-economic contexts, which can impinge on their conduct. How is it that we accept that the persons undertaking the procedure are capable? In what ways can we be confident that the right procedure was followed? Do we know that the equipment was in good condition and was appropriately calibrated? When was it last tested? What economic pressure are there impacting on the work of a lab? So many questions ...¹¹

If traditional approaches to institutional sources of truth confer an authority status on certain agents, then we have a need to validate the identity of the agent in the first place. But that still leaves open the question of the knowledge artifact. Whereas the agent may make a claim that its artifact is valid, a decentralised approach to common knowledge production would turn this authority on its head: only the multitudes can determine the validity of the artifact insofar as the applicable set of agents rely upon a common knowledge for action and coordination.¹²

It is for this reason that truth presuppositions from the traditional sciences are conflated with blockchain's claims to data validity. I have, ever so briefly, suggested that we are actually dealing with two quite different things. The way in which common knowledge is accepted in a blockchain ecosystem makes no claims about truth or otherwise in an epistemological sense. *Valid* does not necessarily mean *true*, with a Capital T.

¹¹ See Latour, B., and S. Woolgar, (1979) *Laboratory Life. The Construction of Scientific Facts* for a classic analysis of the production of scientific knowledge in a laboratory context. Also see <https://www.theguardian.com/science/2019/jul/05/eurofins-ransomware-attack-hacked-forensic-provider-pays-ransom>, and <https://www.theguardian.com/uk-news/2018/dec/06/fifty-convictions-quashed-after-investigation-at-uk-forensics-lab> for recent examples of lab "manipulation".

¹² We can in fact treat data production processes in not dissimilar ways to those in which we are dealing with the issues and challenges in material product supply chains. Data production chains involve similar issues.

We only need valid common knowledge for supply chain functionality.

Economics of Data Production

If blockchains can be valuable records of the past, from our perspective, this is only the beginning of their potential. The capacity to formalise mechanisms for common knowledge formation amongst a group of agents opens up the possibility of not only historic validation but also the framing of desired future possibilities. Recalling that blockchains are good at ensuring data states are valid, we can now unshackle blockchains from being merely ledgers of the past and liberate their potential as specifiers of desired future states. But not only can desired future (data) states be specified, we can utilise blockchains to operationalise mechanisms that link analogue behaviours “on the ground” to the realisation of desired futures wherein the extent of success can be confirmed by valid data states. We thus have the potential in the first instance to limit the possibility of capricious claims on performance; and secondly, we can design mechanisms by which behaviours that increase the probability of desired future states being realised can be encouraged.


Valorisation mechanisms can be created through the application of mechanism design ideas and embedded in the **state change mechanisms** that are endemic to blockchains. Aligning valid data states with value distribution protocols implies the creation of new markets that facilitate adaptive behaviours towards new desirable outcomes.

We, therefore, introduce another powerful tool that is given rise to by decentralised ledgers - namely, the possibility of creating coherent units of value via native cryptocurrencies - which could, perhaps, result in new behavioural “markets” within supply chain activities that reward valid data contributions and punish invalid (false) data contributions. The crypto economics dimension is ‘forward looking’, whereas for the most part, claims about the blockchain and what it can do for supply chain veracity are focused on a blockchain’s capacity to be a record of the past (‘backward looking’). What we can start to think through / imagine / design, from a ‘forward looking’ perspective, is how to maximise the probability that future desired states are realised and are relatively stable via the mobilisation of game theoretic ideas anchored in modes of behavioural incentivisation and disincentivisation.

Specification of desired future states. By aligning payments to realisation of desired future (data) states one effectively has, *prima facie*, a **mechanism** by which behaviours can be shaped (see Figure 4). Rather than presuppose the establishment of trust as a condition precedent of successful supply chains, we posit behavioural dynamics that simply assume actors will prefer to

avoid punishment or loss, and access rewards instead. Notice I do not say maximisers in any narrow sense. This is because we accept that actors in fact can and do frame rewards and punishments in complex, socially embedded ways, in which financial gain (loss) from a particular transaction is but one dimension - albeit in supply chain calculus a crucial one.

FIGURE 4: MECHANISM SCHEMATIC - BEEF SUPPLY CHAIN



To illustrate an alternative facet, we can think about social opprobrium and ostracisation that is a cost, which can arise due to conduct that is viewed as unacceptable by the relevant social setting. An illustration of this line of thinking can be found in Adam Smith's *The Theory of Moral Sentiments*. There, he argues for a socialised understanding of acceptable behaviour, in which a whole range of rewards and punishments help shape each person's disposition towards others. If ideas of virtue are a little fuzzy, we can render them calculable via instruments such as peer reviews and public ratings, in which reputations become legible and calculable.

Another example could see the value of improved ecological outcomes in the production of food valorised and paid via explicit accounting, and the linking of consumer value and payment to activities (validated by data states) further up the chain. We could do similar things on the basis of data states about animal welfare as well as to reward actions that contribute to the provision of "healthy food" versus "unhealthy food".

Smart contracts (as persistent coded procedures stored and executed on a blockchain) are the state change tools to achieve these kinds of valorised links between actions and rewards. We can achieve this by aligning payments with satisfaction of data preconditions, which are in themselves validated as common knowledge via blockchain protocols. Data states have the validity and dependability of the blockchain consensus mechanism behind them. Furthermore, smart contracts are also executed on a blockchain so that the self-executing / self-enforcing properties are also quarantined from capricious alteration, censorship and malfeasance. Both of these properties are inherently capricious-proof.

Design Application: BeefLedger examples

I now want to pull some of these ideas together and illustrate how we are applying them to addressing specific sets of challenges in the beef supply chain, and in particular to issues associated with the China export market.

Let me begin by sketching the key features of the China beef market, insofar as they impact on our discussions today.

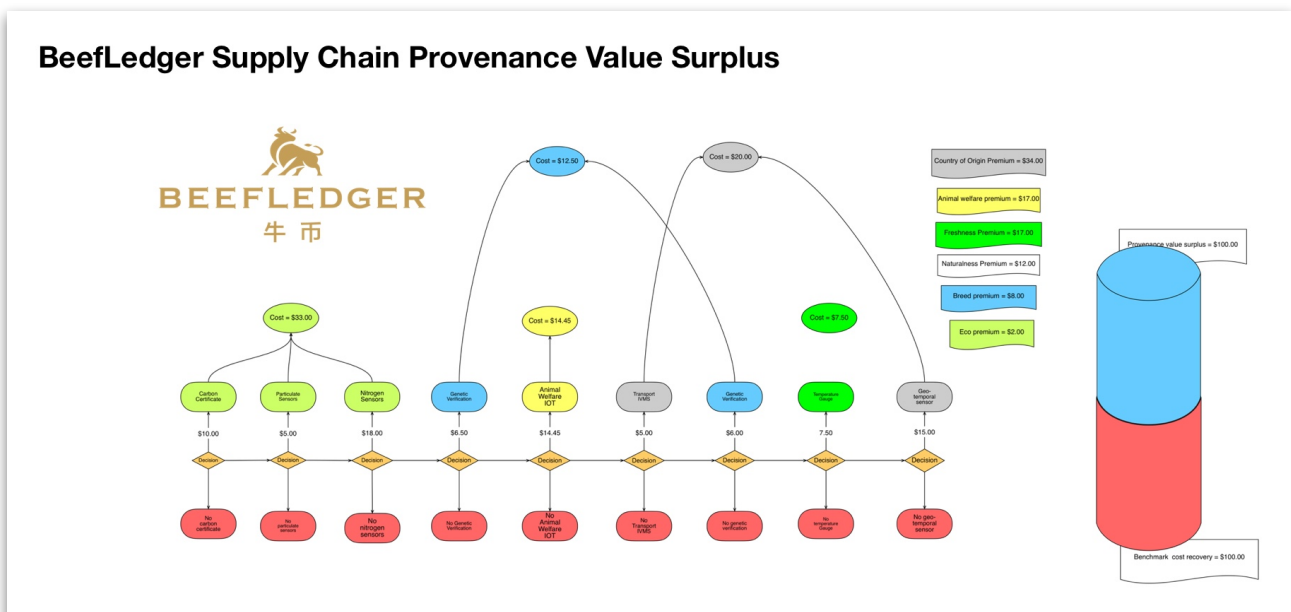
1. Firstly, demand for beef is growing rapidly. Demand for imported beef is also growing rapidly. At the beginning of the decade, China consumed 3.8kg of beef per person per year on average. This is now around 6kg, and is expected to reach the global average of 11kg by 2024. By way of comparison, Australia consumes around 25kg.
2. Secondly, China cannot meet its demand for beef from domestic supplies. Chinese researchers estimate that by 2024, China will experience a supply-demand gap equivalent to 80 million head of cattle. Again, to place this in context, Australia's total beef cattle herd is around 24 million head; having fallen from a peak of close to 29 million about 4 years ago. Australia's annual beef slaughter is in the order of 7-7.5 million head.
3. In conditions of supply shortages, two things seem to inevitably follow. Firstly, prices rise. This is especially the case for imported beef. As the price of imported products rise, in conditions of limited supply, counterfeiting and misrepresentation risk grows.

Chinese consumers already pay a premium for imported beef. Or, to be more accurate, for beef that *claims* to be imported. The difference in price for an otherwise identical cut of beef, between a domestic product and an imported product, cannot be simply explained by reference to the additional costs of production, transportation and customs duties. While these additional costs do account for some of the extra cost of the imported product, there is something above and beyond

these factors. Conceptually, if we adopted a hedonic price approach, we could estimate statistically the factor composition of price. This is a piece of work that we will be doing more systematically in the future.

For now, however, I simply observe that consumers *are paying* a premium for imported products, which can - by way of hypothesis - be explained by reference to the idea of **provenance value**. Meat and Livestock Australia research identifies that for Chinese buyers of imported beef, the single most significant factor is *country of origin*. This is followed by *animal welfare* considerations then by notions of *naturalness*. A chemical scientist may insist that provenance strictly speaking can only be determined via trace analysis; that may well be the case, but for consumers, a **common knowledge** basis of *country of origin* is arguably adequate and economically warranted. A provenance value accounting schematic is presented at Figure 5.

FIGURE 5: PROVENANCE VALUE ACCOUNTING SCHEMATIC



To illustrate this in practice, therefore, we can assume that amongst the attributes of provenance value, the ability to deliver consumers with valid *country of origin* data is a valuable part of the final transaction. Indeed, data concerning *country of origin* is directly correlated with the presence of premium pricing. (We can note that existing premium pricing in effect implies something about provenance value, even though data concerning things like country of origin is largely absent save for a labelling claim.)

The remainder of this paper illustrates a number of hypothetical cases by reference to Table 1.

Table 1: Provenance Value Hedonic Matrix

	Cost 1	Cost 2	Cost 3	Cost 4	Total Cost	Case 1	Case 2	Case 3	Case 4
Price paid by consumer						\$200.00	\$200.00	\$200.00	\$250.00
Benchmark cost recovery						\$100.00	\$100.00	\$100.00	\$100.00
Provenance Value Surplus						\$100.00	\$100.00	\$100.00	\$150.00
Benchmark ROR						\$10.00	\$10.00	\$10.00	\$10.00
Country of Origin Value	\$0.00	\$5.00	\$15.00	\$16.00	\$36.00	\$34.00	\$34.00	\$34.00	\$34.00
Animal Welfare Value	\$0.00	\$14.45	\$0.00	\$0.00	\$14.45	\$17.00	\$17.00	\$17.00	\$17.00
Freshness Value	\$0.00	\$7.50	\$0.00	\$0.00	\$7.50	\$17.00	\$17.00	\$10.00	\$10.00
Naturalness Value	\$0.00	\$0.00	\$0.00	\$5.00	\$5.00	\$12.00	\$12.00	\$11.00	\$11.00
Breed Value	\$0.00	\$6.50	\$4.00	\$4.00	\$14.50	\$8.00	\$8.00	\$16.00	\$16.00
Eco Value	\$0.00	\$10.00	\$5.00	\$18.00	\$33.00	\$2.00	\$2.00	\$2.00	\$52.00

Let us hypothesise a hedonic price composition for premium Australian beef of a given weight (W), as follows:

- A. Price paid by consumer = \$200.00
- B. Benchmark cost recovery of the supply chain as a whole = \$100.00
- C. Provenance Value Surplus = \$100.00, comprising:
 - 1. Benchmark rate of return (say 10%) = \$10.00
 - 2. Country of Origin Premium = \$34.00
 - 3. Animal Welfare Premium Value = \$17.00
 - 4. Freshness Value = \$17.00
 - 5. Naturalness Value = \$12.00
 - 6. Breed Value = \$8.00
 - 7. Eco Value = \$2.00

We can then posit two cases for otherwise identical products.

Case 1 has no common knowledge data validation of any of the **Provenance Value Surplus** (PVS) attributes. In Table 1 no costs are incurred to create the common knowledge data.

Case 2 has common knowledge data validation of all of the PVS attributes. Certain costs are incurred (Costs 2 in Table 1). This data is sourced from valid Oracles. A mechanism - implemented via a smart contract or a daisy-chain of smart contracts - can be designed so that there are explicit rewards for supply chain actors to invest in the necessary assets to perform Oracle functions, based on the consumer-determined valorisation formula.

In the face of Case 1 and Case 2, consumers are able to make discrete choices. If consumers had heretofore been willing to pay Price A for Case 1 beef, and are now presented with both Case 1 and Case 2, it is highly probable that consumers will *ceteris paribus* prefer Case 2. This is because the consumer makes gains (valid data) without incurring additional costs. Those that prefer Case 1 will insist on an impairment of value (price paid) of some amount between \$1 and \$90. At full impairment (when there is no basis of Provenance Value), the supply chain reproduces at Cost Recovery + Benchmark Rate of Return; in this case, \$110.00.

An actor or actors may implement an IOT system to capture data concerning the geospatial and temporal coordinates of the animals and products as they proceed through the supply chain. So long as the total cost of these systems does not exceed \$34.00 there is a case for actors to install the systems. Or, IOT devices could be installed to collect data that relates to animal welfare. The cost of this is, say, \$14.45. The smart contract distributes the Animal Welfare Premium Value quantum to those actors in the supply chain that implemented the IOT system in question. At the Animal Welfare Premium Value quantum specified, the actors implementing the IOT system earn a net surplus of \$2.55 (17.6% ROI).

Let us also assume that a producer incurs a cost of \$6.50 to capture some form of data on Breed. In terms of Case 2, that producer will receive \$8.00 of the relevant PVS. However, let us now assume that the data provided through this method is partial; in this situation, being absent of genetic verification post slaughter. The cost of post slaughter genetic verification is, say, \$4.00. Under Case 2 hedonics, there is no business case for incurring this additional cost (\$6.50 + \$4.00 > \$8.00).

Let us now consider **Case 3**, the hedonic price structure of which is as follows:

- D. Price paid by consumer = \$200.00
- E. Benchmark cost recovery of the supply chain as a whole = \$100.00
- F. Provenance Value Surplus = \$100.00, comprising:
 - 1. Benchmark rate of return (say 10%) = \$10.00
 - 2. Country of Origin Premium = \$34.00
 - 3. Animal Welfare Premium Value = \$17.00
 - 4. Freshness Value = \$10.00
 - 5. Naturalness Value = \$11.00
 - 6. Breed Value = \$8.00 \$16.00
 - 7. Eco Value = \$2.00

Note that the consumer stills pays the same price (D) as per Case 2. However, we see that the consumer values Breed Value higher than other factors such as Freshness Value and Naturalness Value. Indeed, compared to Case 2, Breed Value is now \$16.00 compared to \$8.00 provided that Breed Value data includes post slaughter DNA testing. In Case 3, the producer who was reluctant to incur the additional post-slaughter genetic testing (cost = \$4.00) is now in a position where this additional cost is justified as it provides access to an additional \$8.00 of provenance value. Total cost = \$6.50 + \$4.00 (\$10.50) < total Breed Value = \$16.00.

I want to conclude with one additional Case. **Case 4** hedonics is as follows:

- G. Price paid by consumer = ~~\$200.00~~ \$250.00
- H. Benchmark cost recovery of the supply chain as a whole = \$100.00
- I. Provenance Value Surplus = \$150.00, comprising:
 - 1. Benchmark rate of return (say 10%) = \$10.00
 - 2. Country of Origin Premium = \$34.00
 - 3. Animal Welfare Premium Value = \$17.00
 - 4. Freshness Value = \$10.00
 - 5. Naturalness Value = \$11.00
 - 6. Breed Value = \$16.00
 - 7. Eco Value = ~~\$2.00~~ \$52.00

Note that Case 4 sees a premium placed on ecological impact value of \$52.00. With all other values held constant as per Case 3 (a very demanding consumer indeed), the consumer's desires can only be satiated by a price (G) of \$250.00. This assumes that costs incurred to achieve the desired eco data state is lower than \$52.00.

Similar mechanism design ideas can also be implemented to re-align incentives in cattle transportation in ways that reward actions that contribute to optimal animal welfare and animal condition outcomes (during transit and on delivery). Our early research is focused on thinking through the transaction involving producer-transporter-processor through a principal-agent frame. Our working hypothesis is that there is a moral hazard implied in existing remuneration arrangements wherein the transporter is rewarded on timeliness, with animal condition not explicitly featuring in the economic rewards payment regime. If anything, animal condition on delivery only features indirectly via mechanisms of driver reputation. New reward mechanisms could be imagined that explicitly align gain-share opportunities resulting from better animal condition outcomes. These gains are realisable in the first instance due to less than baseline

weight loss outcomes, and also due to existing processor payment 'grids' that deduct for degraded product, such as bruising. By reducing weight loss, and reducing the incidence of degradation, a greater gross amount is payable by the processor for any given animal. The above benchmark gain can then be shared between producer and transporter.

What I have illustrated here are mechanisms that align data states with actions that can be taken by agents within the supply chain, and which can be evaluated in terms of a cost-benefit consideration. We can achieve this by discretely valorising data, and apportioning relative value via a smart contract-executed mechanism. In effect, we create an explicit price signal that can incentivise behaviours that can be validated via particular data states. In effect, we specify a desired data state, valorise this based on known or anticipated consumer propensities and empower agents to respond accordingly.

Conclusions

The capacity of blockchain technology to support supply chain improvements comes not so much from the technology's capacity to deliver a ledger of past events, though this is certainly useful especially in conditions where there are doubts about the product journey. Rather, it comes from the ability to deploy the technology in forward-looking ways to specify desired (future) data states and create incentive mechanisms by which actions required to meet those data states are economically justifiable. By removing or at the very least severely limiting the possibility of capricious action, the risks associated with incurring additional costs to pursue the required actions are also diminished. None of this presupposes truth or trust. Instead, the power of blockchains in supply chains is premised on more mundane foundations: common knowledge and capricious-free dependability.