FARM FARE

January 17, 2022

Greetings,

As you begin to review the following paper, as a practitioner in the food system, I wanted to provide some context for the topic and this suggested technological path forward.

By way of introduction, I am a champion pig shower. I supported a farmer training program in Charleston, WV. I co-operated a food hub in northeast Ohio. I am a food systems-first individual. I don't have a technology background nor did I anticipate finding my way to this side of the work. But, as I began to have a deeper understanding of the supply chain and understand the components foundational to building an inclusive, just and sustainable food system, I knew I wouldn't be able to connect those components without technology. Thus began my dive into aligning technology with user needs and creating functionality to support a new food system.

As with most things, it's not that simple, especially when it comes to the food system. At Farm Fare, we focus on capabilities that enable food hubs and related aggregators to seamlessly share their inventory to better serve institutional customers and a region-wide production planning strategy. Networking capability includes a suite of tools that enables hubs to share sales, invoice across hubs, manage commission fees and network pricing, etc. Point being, we don't have custom CSA capabilities, our system is more elegantly designed for wholesale customers and to gain functionally of the network tools, you have to use our ERP system.

I say this out of transparency, that our system, like almost every other system, isn't designed to manage all the dynamic complexities of aggregator operations. As an honest, collaborative and strategic software company - we won't ever have all those features. And we think that's in customers' best interest. A good product can't be everything to everyone in such a dynamic and custom industry. And secondly, we can build integrations to other softwares, but to what end? At Farm Fare, we are currently building an integration with Local Food Marketplace, but in the end, this is still just connecting two of the many softwares a hub may use.

The following commentary presents a new lens on navigating this issue. Aggregators need a diversity of softwares to manage their operation. Without a seamless way for data within these systems to "talk" to one another, regional food systems will be stuck with clunky technology that is time and resource intensive.

Please take a moment to review the opportunity a public food blockchain presents to regional food systems. At Farm Fare, we believe utilization of this type of public infrastructure will be transformational for both the ways regional food systems do their work, and the acceleration at which impact can be created. At the close of this paper is a QR code; should you feel compelled to join the conversation about moving towards a regional food blockchain with data standards, please complete the short form to share your questions, ideas and shared interest.

All my best.

Cheers,

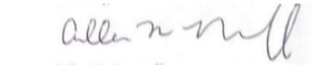

Cullen Naumoff, Co-Founder & CEO
Farm Fare



Winter 2022, V1

Blockchain For Regional Food Systems

The case for a modern day "highway" for regional food system data

Food hubs[1] are enabling evolved supply chains business models that support fairer, more resilient and collaborative local food systems.

Food hubs are underserved by technology. The challenge is that food hubs have complex needs, and the industry is not very standardized. This translates to needing multiple applications, which rarely work together. Hubs spend immense time copying data back and forth between systems.

At the same time, blockchain is creating an emerging model for building shared software and data ownership. Today's software solutions fail to serve food hubs in many ways that blockchains excel at.

1 https://www.usda.gov/media/blog/2010/12/14/getting-scale-regional-food-hubs#:~:text=Our%20working%20definition%20of%20a,value%20chain%2C%20food%20hubs%20are

What Is Blockchain & How Is It Different

Traditional Database Applications

Traditional web applications are built around centralized databases. Applications are then built on top of that database, which users interact with. Only one company owns, controls, and dictates changes to both the database and applications.

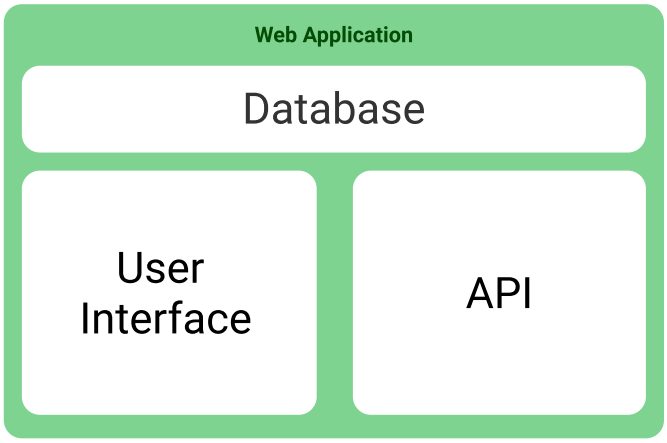

This model works quite well for a self-contained service. Where this model struggles is when collaboration is needed. For two companies to allow collaboration between their applications, each must write an API (application programming interface). Then each company must write integrations to communicate with the other's API.

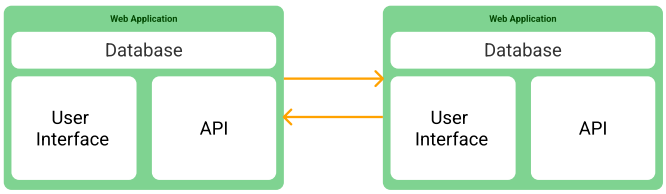

Inherent in writing this API is also determining truth. For example, imagine you have two applications that need to work together to manage a set of shared resources — in this example, a tractor rental application and a crop planning/planting application. The planting application would like to reserve the correct number of tractors on certain planting dates. What happens when a request fails? When two requests are made for the same tractor at the same time? Or when a reserved tractor breaks down and won't be available in time? How do you ensure consistency between these applications?

For traditional database applications, the solution is a custom and complicated solution specific to that API.

Now let's add a third collaborator, a shipping service in this example. They need to know the dates from the planning application and the tractor dimensions to determine how they will move this equipment and when. Now, instead of two applications and two integrations, we have three applications and nine integrations. Adding a fourth collaborating application brings us to 16 integrations.

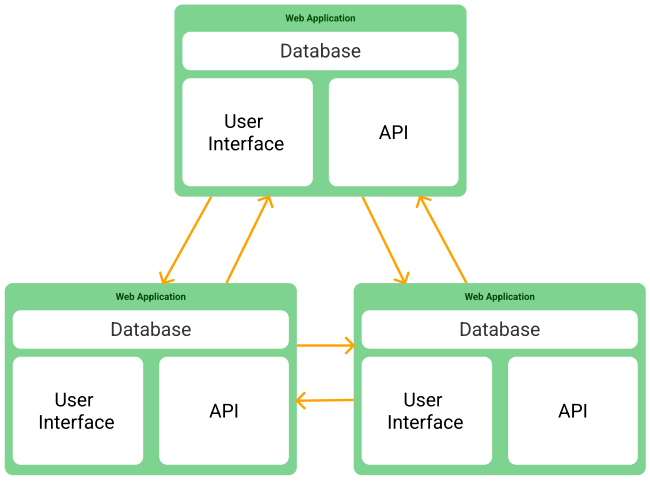

a network with 3 nodes

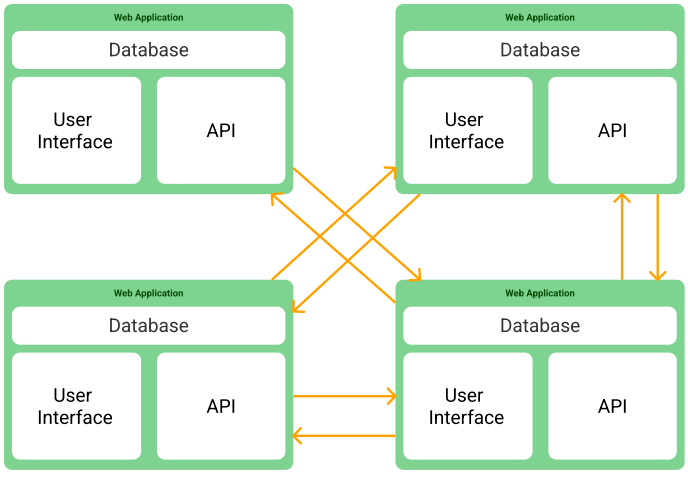

a network with 4 nodes

As you can see, this model leads to exponential growth in costly integrations that quickly overwhelm collaborators.

Blockchain Applications

A blockchain, by contrast, starts at the opposite end of this problem. Unlike a central database, blockchain assumes a team of collaborators. In a blockchain, each collaborator stores a separate copy of the data, but before any change is written to the blockchain the collaborators must reach agreement on what is true.

Using the example above, if the planning application wants to reserve a tractor, it submits the request to the group; then, once the group has decided that it is a valid request (i.e. a tractor is available and the user has the funds to rent it), each application records the data.

In a blockchain, the process by which the group determines what will be written to the blockchain is called a consensus algorithm. Different blockchains use different consensus algorithms, but they all result in each node/computer having the same verifiable record of the data.

In a blockchain application, before an application is written, before an API can be exposed, the blockchain supplies communication and synchronization.

Blockchain communication and synchronization are also standardized. Every single node in the network submits, reads, and stores data in exactly the same way. That means the cost of adding another participant is not every original collaborator implementing the new collaborator's API and the new collaborator integrating with all of theirs with an exponential increase in implementation cost. With blockchain, the cost of adding a new collaborator is the new collaborator integrating with the blockchain directly.

Blockchain as Public Infrastructure

The network cost of automating collaboration of traditional database centered applications follows the structure:

$$cn(n-1)$$

n = Number of collaborators
c = cost of an API implementation

$$c2(n-1)$$

n = Number of collaborators
c = cost of an API implementation

Whereas the network cost of automating collaboration with a blockchain is:

$$cn$$

n = Number of collaborators
c = cost of an API implementation

And the cost of adding a single collaborator is:

$$c$$

n = Number of collaborators
c = cost of an API implementation

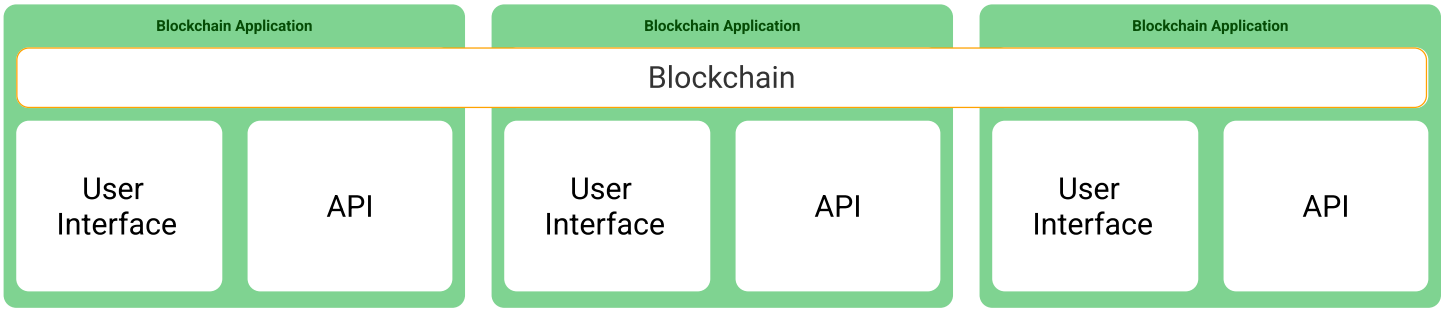

a network of blockchain nodes

To build intuition around the growth of those costs, let us assume a cost of $10,000 for an API integration (a low figure for an API integration, but illustrative of the cost growth).

Participants	Traditional Network Cost [Cn(N-1)]	Traditional Last Collaborator Cost [C2(N-1)]	Blockchain Network Cost [Cn]	Blockchain Last Collaborator Cost [C]
1	$0.00	$0.00	$10,000.00	$10,000.00
2	$20,000.00	$10,000.00	$20,000.00	$10,000.00
3	$60,000.00	$20,000.00	$30,000.00	$10,000.00
4	$120,000.00	$30,000.00	$40,000.00	$10,000.00
5	$200,000.00	$40,000.00	$50,000.00	$10,000.00
6	$300,000.00	$50,000.00	$60,000.00	$10,000.00
7	$420,000.00	$60,000.00	$70,000.00	$10,000.00
8	$560,000.0	$70,000.00	$80,000.00	$10,000.00
9	$720,000.00	$80,000.00	$90,000.00	$10,000.00
10	$900,000.00	$90,000.00	$100,000.00	$10,000.00

Graphically, that data looks like this:

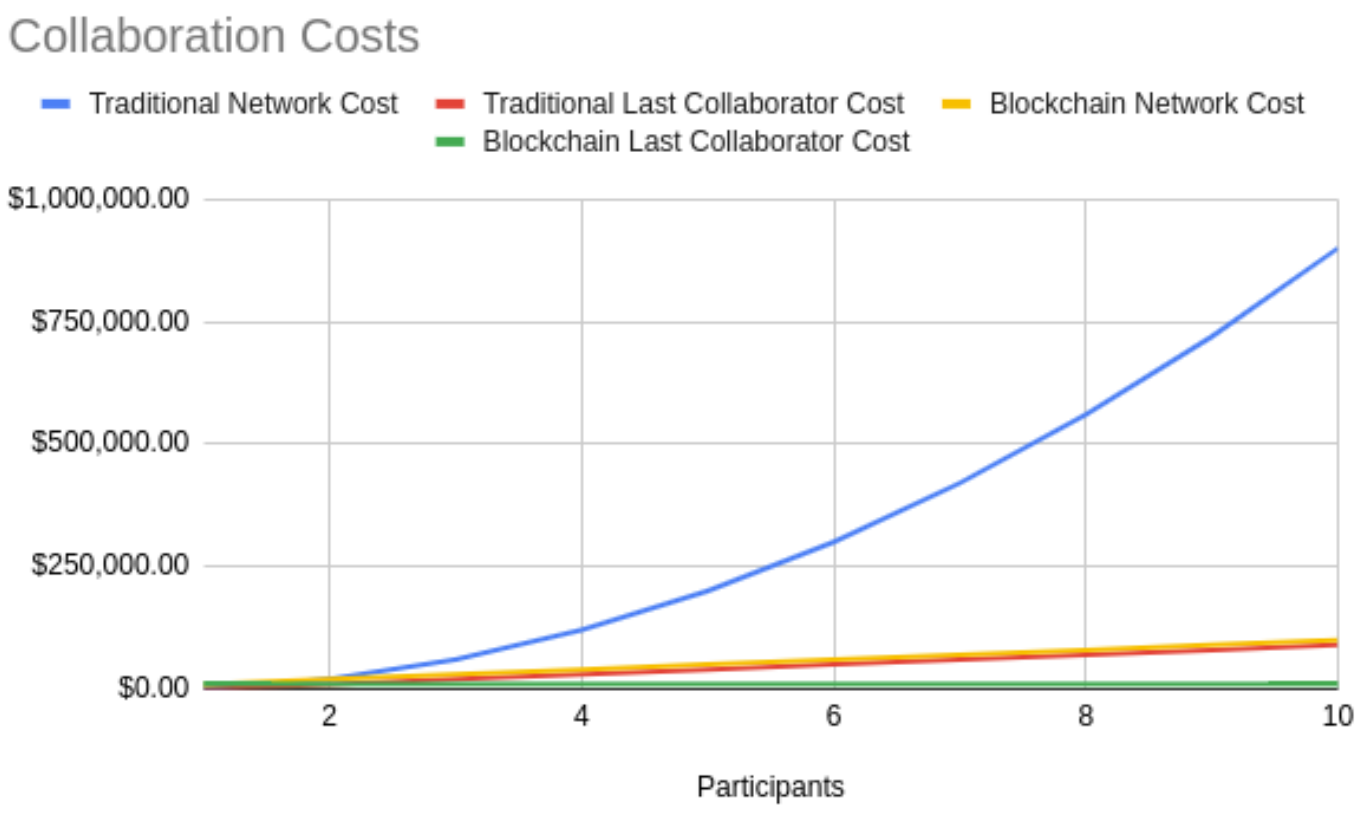

You may notice the total network cost with blockchain is very close to the cost of adding a single collaborator with traditional database applications.

Due to the cost dynamics of collaborating through a public blockchain, we should think about a food systems blockchain as public infrastructure, similar to the highway system.

It would be very expensive for a supplier to purchase all the land between them and a purchaser to build a single-purpose road. Building a single-purpose road to every customer would incur a similar c(n-1) cost structure we see with API integrations.

Building a public road system has enabled every purchaser to reach every buyer. Sharing roads for multiple purposes means the cost of the road gets distributed over the largest broadest user base and allows those interconnections to happen at a small fraction of the cost. In the same way, a food hub's e-commerce system sharing a blockchain with their inventory and logistics systems would allow those systems to communicate and synchronize at a greatly reduced cost to every participant.

What Are the Private Benefits of Public Data

Aside from the cost efficiency of this infrastructure, the data itself also provides a number of collective and individual benefits. By sharing the quantity and provenance of food moving through the food system, each participant gains better insight and can therefore make better decisions.

Traceability

Blockchain is already well established in traceability applications. A diverse and universal food blockchain would allow customers to track not just primary ingredients, but also the provenance of each ingredient that goes into processed or value-added products. Similarly, tracking at that level would allow for very targeted recalls, as the path of each individual item could be pinpointed.

Crop Planning

Today, most farms are left to guess about overall market trends. This is particularly true in regional food systems where robust regional market data does not currently exist. Frequently, farmers grow a crop not because there is a market for it, but because it grows well. This leads to waste, while customer demands go unmet. The farmers, of course, have no way to identify supply and demand trends; they can only see the performance of the crops they grew, not the ones they didn't. With the entire market captured in a public data store, farmers and their supply chain partners would be empowered to better match their production to the needs of customers. This would help farmers sell more of what they grow and not waste seed, water and amendments for crops that won't sell.

Financial Services

Today's lack of visibility into local market data poses challenges for institutions to support small regional specialty crop growers, such as with crop insurance. While specialty crop insurance policies do exist, the lack of market data often leaves these insurers unable to value a planned harvest — and therefore unable to insure it. A public directory of market data would allow insurers to quantify their risk and offer coverage to a broader audience of specialty crop farmers.

Federal Programs

The same factors that limit financial services in the food system also limit small regional farmers from participating in federal programs. As an example, 2020's food box programs initially intended to rely mostly on regional farmers. However, due to a lack of market data, the government was unable to establish fair market prices and couldn't create contracts for fair payment in many regions. This led them to fall back on the larger distributors and leave out many smaller local producers across the country.

What Are the Concerns About Private Data in a Public Context

While some data absolutely should be public, other data is more sensitive and should only be available to select parties. Blockchains achieve this through encryption. By recording encrypted data to the blockchain, the data gets recorded by every collaborator but can only be read by collaborators who have been given the decryption key. This model gives us a number of additional benefits.

Validity

Because of the process of reaching agreement, and the fact blockchains do not allow data to be changed once written, collaborators can verify that a piece of data is correct even without being able to read the contents of the data. This process builds trust between all collaborators.

You can think of this like depositing valuables at the bank. The bank gives you a lockbox and a key. You place your valuables within the box, the bank records that this box belongs to you but has no idea what is inside of it, and only you can open it. The difference here is that the key is math and the vault is located on every participating computer.

Resilience

Because every collaborator retains a complete copy of the blockchain, everyone's data is automatically backed up. Even during catastrophic failure, as long as one collaborator's copy survives, everyone's data persists.

Data Sovereignty

The ability to store private data publicly creates an opportunity for customer data sovereignty. In a traditional traditional database application, the owner of the database owns all the data. They have full control of the data and can read, write and modify it at their discretion. If a user has submitted encrypted data to the blockchain, only collaborators who have been given the decryption key can read that data. This enables a model where the users of the system, not the owner of the system, own the data. Only collaborators given the keys can read the data, the network as a whole decides what is written, and no one can change data that is already written.

Zero Cost Onboarding

Perhaps least intuitive, but most valuable is the impact of the blockchain in onboarding users. While each food hub will need to load its data onto the blockchain the first time, subsequent application onboardings can reference data already on the blockchain. Most enterprise software systems have large upfront costs in getting a company's data into the system. Instead of weeks spent exporting, reformatting and importing a company's data into the new system,onboarding can be as simple as giving a decryption key to the new application.

For example, if a food hub's inventory system is already writing to the blockchain, they could add an e-commerce interface that would be synchronized with that inventory system just by sharing that key.

What Are the Risks and Current Trends of Blockchain

While blockchain is gaining more and more traction in tracing our supply chains, it is unfortunately, being applied very differently from the vision described here. Rather than public, universal and equitable, the solutions currently being implemented are private, fragmented, and proprietary.

Private

The current wave of blockchains is generally built, owned and operated by a single company. This means the company that owns the blockchain controls who has access and can see the data. This eliminates financial services, government programs and countless other participants from benefiting from the data, thereby limiting their ability to serve local food systems.

Fragmented

These existing solutions are also oriented to either the final retailer in the supply chain or are limited to a specific product. This leaves all of them with only a subset of the data. So even if you gain access to one chain, you will see a much smaller, less complete set of data. Access to the region's beef sales[1] doesn't help an insurer value a field of carrots.

This also removes much of the efficiencies of a shared blockchain. Instead of a single integration giving access to all partners, farmers would likely need to interact with a chain per buyer, pushing cost and complexity towards the traditional database model.

Proprietary

Being owned by a single entity also means any activity with the blockchain must be approved by the operator of the blockchain. This means the participants do not participate as equals. The owner of the blockchain determines who gets to participate and when; leaving the owner holding all the cards and the rest of the participants subject to their whims.

What Are the Challenges to More Collaborative Blockchain

Resilience and Security vs. Transactions and Storage

There are a couple of potential paths to implementation, each with its own benefits and drawbacks. For this discussion, let's examine two large categories of approaches: Leveraging an existing general-purpose blockchain or creating a new task-specific blockchain.

Most general-purpose blockchains ensure security by having such a large number of participants that no one person could own a majority of them. This approach also ensures a very large number of replications of the data in that blockchain so data is never lost, and plenty of computers are available to process transactions. These are all great properties of a blockchain.

The downside of the large general-purpose blockchains comes from their usage. Being large and reliable means they process lots of transactions. Those transactions mean more cost in electricity used to process them and more required data storage required to store them. The participants processing those transactions are generally compensated in the native token of the blockchain. This compensation is taken from fees charged to the person requesting the transaction, and these fees can get very large. Depending on volume, the largest of such blockchains have seen transaction fees as high as $70; but even a more reasonable $4 fee is an awful lot to pay to record the sale of a pint of strawberries.

A task-specific blockchain, on the other hand, will have far fewer participants. That means there are fewer replications of the data, so there are fewer participants that would have to fail (as in losing power or a computer component shorting out) to lose data. This makes them less resilient. The benefits of a task-specific blockchain are in the cost structure. Since this blockchain only holds transactions related to this one task, there are far fewer transactions to process and store, greatly reducing these costs. Participants can also be incentivized through different means; for example, it is so beneficial for the hubs to know what inventory is available on the network that the blockchain could be run without transaction costs.

1 https://beefchain.com/

Acquiring Collaborators

Blockchains don't just enable collaboration; they require it. Without collaborators validating and replicating this data, and without a group to reach consensus, data can easily be manipulated. Without other participants, data would not be replicated and preserved through system failure. Building a group of collaborators is a critical step in the success of a blockchain.

Part of gathering a group of collaborators is incentivising their participation. While many chains charge transaction fees which are partially distributed to node operators, those node operators are there only to operate nodes, not to pursue some other business interest. If we respect the data sovereignty of farmers correctly, it should be worth the cost of participation to get access to farmer inventory data and the actions of the rest of the network. Similar to how service providers today buy servers and databases and build integrations to other services at significant expense, a blockchain can provide those functions and do it at a reduced cost.

Funding development

Perhaps the biggest hurdle is paying for initial development. Once such a blockchain exists and is in place, the operating costs should be both small and borne by the participants, but that still leaves upfront development costs.

A properly incentivized blockchain would have similar public utility to other public infrastructure we rely on, like highways and libraries. It also has similar cost structures — too high for any single participant, but extremely cost-effective at scale. It is also critical that no one owns this blockchain. One owner of such a blockchain would leave all participants subject to their whims. The participants must control the blockchain and transactions must be executed on a level playing field. Inherent in this method of operation is the need for governance. There are many emerging models for governing decentralized entities with multiple owners. In working towards this public blockchain future, identifying a form of shared governance will be critical to operating this infrastructure.

governance will be critical to operating this infrastructure.

How Can Food Systems Move Forward

While a public food blockchain will create many efficiencies for users, particularly small players that can't afford investments in highly integrated systems, it will take significant collaboration to move into an implementation phase. Collaboration includes technology companies writing an API to move data onto the public blockchain, ensuring there are a relevant number of nodes to replicate the data, not to mention ultimately funding the development.

As a result of the public nature of this approach, government should have buy-in on the solution. The public sector will be a central player to organizing stakeholders and providing a portion of the capital for development. This opportunity is primed for a public-private partnership. The public sector can de-risk the opportunity by providing some start-up capital and the convening force, and private sector players can push the initiative to the finish line with their newfound buy-in.

Blockchain technology enables the communication and data integrity that food hubs need to achieve the level of collaboration required for an equitable, resilient and cost-effective food system. Beyond those basic needs, blockchain also brings benefits to players at all levels of the food system. However, these benefits are only possible if we can overcome current trends in supply chain blockchains and build a blockchain that is public, universal and equitable.

Want to help? If you would like to be part of the collaborative to support a data standardization and blockchain effort for regional food systems you can sign at https://forms.gle/QqM9W1ohn45moiGn6 or use the QR code below:

2 Some examples of blockchain governance structures:
- https://sovrin.org/library/sovrin-governance-framework/
- https://ethereum.org/en/governance/
- https://cardano.org/governance/

Acknowledgements

The following entities and individuals were key contributors to the thinking, creation and authorship of this brief. We are grateful to their leadership on this important topic.

Anthony Coble, Farm Fare
Mitchell Bennett, BlueTell Technologies
Daniel Conway, The Conway Family Foundation
Cullen Naumoff, Farm Fare
Jules Repusseau