Bitcoin and Credit Card Networks

A Disruption Theory Analysis

Master's Thesis

Mads Krogh

(93098)

MSc Business Administration and Ebusiness

May 15, 2019

Supervisor: Raghava Rao Mukkamala

Number of characters: 163.989

Abstract

Bitcoin is multi-dimensional by nature, embedding complex economic incentives and decentralization in code, unlike anything before it. This makes it inherently difficult to grasp, in the context of similar existing centralized systems, and much remains to be understood about its design, dynamics, and use cases. Meanwhile, its low transaction throughput is increasingly used as proof of Bitcoin's inherent inferiority to credit card networks as well as other "more efficient" cryptocurrencies. This analysis builds on a superficial one-to-one comparison of Bitcoin and credit card networks and assumptions around market demand and technological progress. Such assumptions are at best unrigorous and at worst harmful, leading to potential misunderstandings around the technology, fractioning within the Bitcoin community and misguided investments in overhyped blockchain projects.

Following a deductive line of reasoning, this work engages in a rigorous technical comparison of Bitcoin and credit card networks to test the assumptions around Bitcoin's inferiority. Disruption theory is used to analyse the two fundamentally different competing systems, linking technical design with use cases and market needs. Our findings suggest that Bitcoin, at present, is superior to credit card networks in niche markets, where transaction throughput is in lower demand than features such as censorship resistance and store of value. Moreover, while credit card networks remain superior for mainstream western consumers, where throughput is in high demand, Bitcoin has the potential to grow transaction throughput and market share with innovative solutions such as the lightning network while preserving its unique features. This is almost by definition what describes a disruptive technology. We argue that this disruptive potential is unique and not transferable to other cryptocurrencies. Bitcoin and the surrounding community is special given the refusal to compromise on the core values which makes it unique.

Contents

Introduction

Despite the increasing adoption and understanding of Bitcoin, its utility, use cases and ultimately its *raison d'être* remains highly disputed. Developing from the dream of cypherpunks to public laughing stock, illegal darknet currency, speculation vehicle and, for some, digital gold, large concerns have most recently been raised around Bitcoin's ability to fulfill its "true vision" and become a peer-to-peer electronic cash to substitute state issued money and credit card networks. It is increasingly treated as a fact that Bitcoin in its current design won't scale to handle the thousands of transactions per second demanded from a global point of sale payment system. Hence, it is doomed to remain an worse alternative to credit card networks.

This narrative has spread inside the Bitcoin community itself to and throughout the wider cryptocurrency ecosystem. It has caused fractioning amongst main Bitcoin actors and the forking of Bitcoin Cash along with several core Bitcoin contributors and miners. Outside Bitcoin, new cryptocurrencies such as Ripple has developed with efficiency and transaction throughput in mind, in their quest to compete with credit card networks, trading off much of Bitcoin's original decentralization ethos. Overall, these projects and the wider public shares a view of Bitcoin as a poor implementation of the "revolutionizing" technology that is blockchain. This assumption has been the origin of much dispute within Bitcoin and new blockchain related projects hyping the technology while promising revolutionizing outcomes to investors and the public.

We wish to test the assumptions around Bitcoin's infeasibility as a substitute for credit card networks. The focus on current transaction throughput, one metric out of many, seems too simple for the comparison of such complex and fundamentally different systems. While such a metric obviously is significant, its context, how the performance is achieved, the demand for such performance and the tradeoffs it imposes on other features is of equal importance. This is true, given that electronic payment systems have a deep and broad foundation which includes the technological component as well as a part economic, part business component, as shown later in the work. The computer science-related parts of these systems and their designs are closely tied to their overall economic incentives, which strongly

impacts the behaviour of users. A fruitful comparison of these systems should acknowledge their vast technological differences and how these shape the behavior of actors.

Disruption theory offers a useful framework for such a comparison. Its distinction between sustaining technology and disruptive technology, respectively beneficial and threatening to incumbents, incorporates both technological design and market demand. If Bitcoin indeed is but a poor implementation of "blockchain technology", competing on the same product features and in the same dimensions as credit card networks, for instance transaction throughput, it can indeed be said to be sustaining and, in line with the popular critique, it will have a hard time displacing the incumbent system. Meanwhile, Bitcoin can be considered disruptive to credit card networks if, as a product of its technical design and innovation, it is superior in different dimensions, while possessing the ability to become "good enough" in the incumbents main competitive dimension over time. Hence, the application of disruption theory to issue at hand begs the following question:

Is Bitcoin a disruptive or sustaining technology to credit card networks?

In the quest to determine whether Bitcoin should be categorized as disruptive or sustaining to credit card networks, this paper will take the following shape. Chapter 1 accounts for disruption theory and motivates its choice in the context of related theories. Chapter 2 outlines the methodology of the subsequent analysis. Chapter 3 takes a deep-dive into the technical underpinnings of credit card networks, using Visa as an example, and Bitcoin. It accounts for their foundational technologies and how these are put to practice in transactions. Chapter 4 combines the technical insights of the previous chapter with market analysis to determine whether Bitcoin is disruptive or sustaining to credit card networks. The first part of the analysis compares the two systems on a number of technical dimensions while the second part frames the comparison in the context of user needs and the target market. Finally, chapter 5 discusses the findings of chapter 4 and asks to what degree these findings are transferable to other cryptocurrencies.

The findings of chapter 4 indicate that Bitcoin is disruptive to Visa and other credit card networks given its existing unique technical design paired with emerging innovations in the Bitcoin ecosystem. These innovations, in particular the lightning network, reside in higher layers of the Bitcoin protocol stack and

have the potential to mitigate the shortcomings of the Bitcoin base protocol in terms of transaction throughput, while preserving the unique features which differentiate Bitcoin from Visa and other fiat-based credit card networks. Moreover, we suggest with our discussion that this disruptive capability is rather unique to Bitcoin in the cryptocurrency ecosystem, given its realistic expectations around the strengths and weaknesses of a blockchain-based system. This is the generator of Bitcoin's layered ecosystem, with a highly conservative and sustainable innovation philosophy at the base protocol, and fruitful experimentation at the higher layers.

1. Theory

In this chapter, we describe the main theory of this work, Disruption theory. To motivate its choice and to better understand it and its application, we start with a brief review of related theories in the area of technological design, potential, and evolution.

1.1. Analyzing Technological Potential

This paper provides a thorough analysis and comparison of Bitcoin and credit card networks. As it later will become apparent, these are two fundamentally different systems competing for similar use cases. Additionally, one of these systems (Bitcoin) is decentralized and non-corporate without any formal, industry, organization and governance structure to analyze. Hence, to undertake the analysis we are looking for a theoretical framework which has the ability to describe the competition between highly differentiated technologies without comparing the organizations where they reside. Initially, there are three classical theoretical and potentially relevant bodies which have been developed to describe how technological evolution emerges.

1.1.1. Dominant Design Theory

Formulated by Abernathy and Utterback (1978), the dominant design theory describes how a single technology or design ends up winning within an industry. The theory suggests that there, within any new industry, initially is a period of ferment characterized by iteration and changing designs as competitors struggle to find the best technology for the market. Eventually, a dominant design does emerge which, being the best or at least good enough, is accepted as the standard by the whole industry and its customers. With the identification and acceptance of the dominant design, the speed of innovation declines as the industry puts its energy towards adapting organizations, production, value chains and marketing to this technology (Abernathy & Utterback, 1978). The dominant design theory mainly describes the evolution of technology as a new market and industry emerges. Thus, it offers limited utility towards predicting how a new design such as Bitcoin might impact an existing dominant design like credit card networks.

1.1.2. The Technology S-Curve

The theory of the technology s-curve describes the progress of technological evolution over time. The theory is best summarized in the s-curve graph with engineering effort on the x-axis and innovation on the y-axis. Initially, the rate of innovation is slow but as time passes by it accelerates until, at some point, the curve starts flattening again before stagnating almost completely. The key insight of the S-curve is that innovation ultimately is subject to decreasing returns to engineering effort given natural constraints such as size, complexity, and materials. Approaching these limits, more and more effort is demanded to innovate on the technology. What naturally follows is then that technologies in this stage of their innovation lifecycle are vulnerable to new, less mature, and more innovating technologies, which again become vulnerable to substitution by even newer technologies as they reach their innovation limit (Foster, 1986). S-curve theory lends itself well to the comparison of well-defined competing technologies. Meanwhile, in our specific effort to analyze two highly differentiated systems in terms of design and features, the s-curve does little to describe how and where they might compete.

1.1.3. Disruption Theory

Disruption theory has become the main force within a stream of research one might call technology and market trajectories (Christensen, 1998). Similar to s-curve theory, this approach emphasizes the rate of improvement in technological performance over time, which it calls "trajectories". Contrary to s-curve theory, it doesn't see all technological improvement as equal but instead considers them relative to market needs. If the performance of a technology grows to exceed the demand of customers further innovation will be of less value. On the contrary, improvement beyond demand will most likely make the product less appealing for consumers as the new improvements don't outweigh potential price increases, increased complexity, and lower reliability. It follows that a new technology which can provide good enough performance that meets market demand in the same dimension of performance, while offering better performance in associated dimensions will outcompete the incumbent technology. Disruption theory enables the comparison of competing technologies. Its inclusion of different dimensions of performance allows us to analyze and compare highly differentiated technologies to each other. The emphasis on market demand and impact allows for predictions of winners between competing technologies. Hence, this paper uses disruption theory as the main theoretical framework for the subsequent analysis.

Formalized by Clayton Christensen in his book The Innovator's Dilemma (2013a), the current popular understanding of disruption theory articulates and addresses one main question: Why do great companies fail? Specifically, it addresses *"the failure of companies to stay atop their industries when they confront certain types of market and technological change."* (Christensen, 2013b, p. ix). These are companies that do everything after the book; they strategize, innovate and execute. They constantly pay attention to the competition, listens attentively to their customers and invest aggressively in innovation, and yet they lose market dominance. Through the analysis of historical data from different industries, using the aforementioned technology and market trajectories approach Christensen suggests that the answer lies in what type of technology these companies are faced by; when challenged by new *sustaining* technology the market leaders excel while *disruptive* technologies displace the incumbent technologies and firms completely (Christensen, 2013b).

1.1.3.1. Sustaining technology

The majority of technologies are sustaining. They improve on the previous technology by doing the same things better. Staying within the context of trajectories, they move up the same trajectory as the incumbent technology is on. A sustaining technology can be either evolutionary (technically similar) or revolutionary (radically different) and so the technological differentiation of new technology cannot itself tell what market impact it will have. For instance, the first automobiles were arguably a sustaining technology to the incumbent transport technology of horses. These early cars remained expensive luxury items which only a few rich individuals could afford, and competed on the same performance dimensions as horses such as distance and comfort. Incumbent technologies and the market leading firms that produce them thrive on sustaining innovation as these technologies fit into the value chain and can be marketed to existing core customers (Christensen, 2013b).

1.1.3.2. Disruptive technology

Occasionally disruptive technologies come about. Contrary to their sustaining counterparties they don't offer improvements on the main performance dimensions valued in the market. As a matter of fact, they perform significantly worse than the incumbents in these dimensions. As a result, disruptive technologies start out as niche products marketed to groups who for some reason can't use the mainstream products, whether these are too expensive, too complex or lack some specific important feature. While ignored by

mainstream consumers and producers, a disruptive technology takes root in this new market segment until at some point establishing itself as the new low end of the larger market. With its own niche segment locked down, the new entrant using the disruptive technology have strong incentives to move upmarket and capture more profits by improving the technology according to what is valued by mainstream consumers. As the technology reach a good enough level it displaces the incumbents; the new technology is now good enough where it matters for most customers and it comes with the added advantages which made it popular to the niche market in the first place. Hence, the new technology attacks the incumbent from underneath (Christensen, 2013c). To continue our example from the previous section, while the initial automobile was sustaining in its nature, the Ford Model T truly disrupted the transportation market; it was good enough in terms of comfort and distance while highly superior in price and reliability.

An obvious question remains: In order to counteract new entrants, why doesn't the incumbent just move down market by selling "worse" editions of its existing technology adopt the disruptive technology itself? According to disruption theory, incumbent firms and technologies trap themselves in the high end of the market for different reasons. First, small niche markets such as those where disruptive technologies initially take root don't solve the growth problems of large companies and as markets that don't exist can't be analyzed market leaders will tend to dismiss these opportunities. Second, companies depend on their most profitable customers why they are naturally inclined to develop (sustaining) technologies for the needs of core customers while killing (disruptive) initiatives that these customers don't want. Hence, they will also be inclined to move upmarket, not downmarket, as the most profitable customers reside here. This causes the aforementioned over optimization on the performance dimensions which causes performance to exceed the needs of the average consumer. In turn, this creates space for a new entrant from below.

But it is not only core customers that denies incumbents from responding to disruption. The suppliers, sub-suppliers, producers of compliments, and retailers (that is both the industry and the value chain) makes up the value network in which the company operates. Within each value network is embedded certain metrics of value which are the main attributes of product performance across the network. For an incumbent to adopt a disruptive technology effectively means leaving its value network and joining, or even building, an entirely new network, based on new value metrics (Christensen, 2013b).

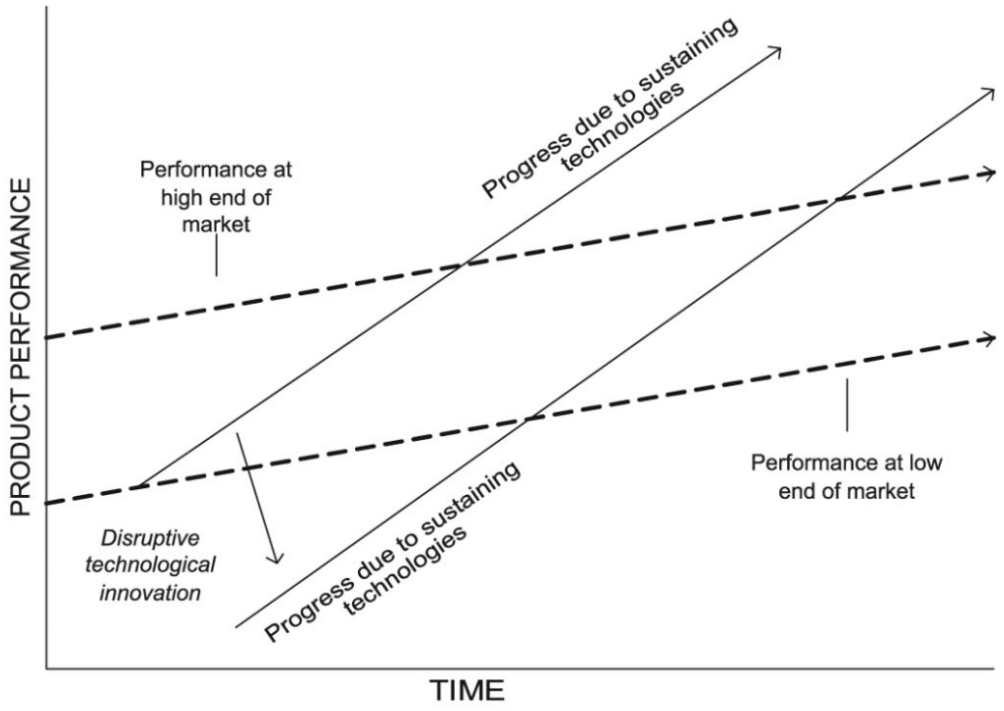

Figure 1.1: Conceptual graph of disruption theory analysis from Christensen's The Innovator's Dilemma. Over time the performance of the incumbent technology grows to exceed demand, leaving room for a new entrant. Meanwhile, the disruptive technology improves to a level good enough for an increasing share of the market. (Christensen, 2013a)

1.1.3.3. Locating Disruptive Technologies

Knowing how to describe and spot a disruptive technology after the fact, how do we make educated guesses about the disruptive potential of upcoming technologies such as Bitcoin? The obvious answer is to apply the most common traits of disruptive technologies to our case and with the data currently available evaluate to what degree the technology currently is in possession of these traits. We briefly touched on the features of disruptive technologies above and based on Christensen's foundational work more studies have followed to refine our view of the archetypical disruptive technology (Christensen et al., 2001, Christensen et al., 2015). Accordingly, about the potentially disruptive technology, we would ask questions such as; "is it cheaper with lower margins, or does it have some other differentiating feature", "does it perform poorly to mainstream consumers needs", "does it foster new, difficult to predict, markets and services" and "does it have the potential to improve performance and become good enough for mainstream consumers"?

This disruptive archetype is derived from the aforementioned technology and market trajectory analysis. Christensen applies empirical evidence to the model by studying specific industries over time. The main study behind his theory revolves around the hard disk drive industry which, with its rapid rate of change, is optimal for studying multiple technological and business shifts over a relatively short time span. It starts with a thorough account of the origin of disk drives, how they are engineered as well as their technical features and the reasoning behind those features. Central to the analysis, it also includes the performance measures of each version plus the core practical features. This account encompasses all significant versions of disk drive technology over a forty year period and lays the foundation for the subsequent grouping of the disk drive technologies and transitions between them into sustaining and disruptive technological changes. In practice, this account and classification materialize in a mapping of the performance trajectories of different disk drive technologies over time. A conceptual representation is given in Figure 1.1 where the black line at the top represents an incumbent technology, and the bottom black line a disruptive technology (Christensen, 1993).

Besides the supply-side, that is the technological performance, another important part of Christensen's analysis is the demand-side, i.e. customer demand for performance. It is this demand together with supply which determines if there is room for a disruptive technology at the lower end of the market. While performance is relatively apparent (companies tend to advertise is in public), demand can be more difficult to measure. To plot demand trajectories in his analysis, Christensen uses the performance (still disk capacity) of the median-priced computer system sold in each market category, the reasoning being that this represents what the average consumer really demands (Christensen, 1993). A similar method is used by Christensen in his analysis of the mechanical excavator industry as demand here is calculated from the average bucket size (performance) bought by contractors (consumers) (Christensen, 2013e). In Figure 1.1, this demand is represented by the area between the dotted lines which represents the demand interval between the low-end and high-end of the mainstream market.

Additional to mainstream consumers, a second part of the demand-side analysis is the low-end part of the market. For the disruptive technology to initially take hold it needs to have one or more unique features which make it valuable to some niche market segment, not currently satisfied by the mainstream

technology. Christensen reaches this conclusion by assessing the initial foothold of technologies that later turned out to be disruptive. Disruptive disk drive technologies where usually smaller and more rugged than the competition, thus enabling new use cases. A good example is the 5.25-inch drive introduced by Seagate in 1980. It was inferior to the existing 8-inch drive in speed and capacity but its smaller size and lower price made it economically viable to fit into desktop computers and it even enabled the first portable and laptop computers. Hence the technology established its foothold with consumers that couldn't afford the existing *minicomputers* or was physically constrained in their use of the machines (Christensen, 1993). In sum, one can't find a disruptive technology without knowing the market in which it operates - what are the competitors, who are the customers, what do they demand and are those demands met.

In our quest to determine if Bitcoin is disruptive to credit card networks, we will use a similar approach to that of Christensen; account for the technologies and map their trajectories given the standard performance dimension, determine mainstream market demand, and locate potential new performance dimensions introduced by Bitcoin. Worth mentioning is that Christensen, besides the descriptive side of disruption theory, also formulates a prescriptive side of the theory. This part aims to instruct managers and companies on how to deal with disruptive technology in order to survive. However, this work will not address this part of the theory but will exclusively use the descriptive part, described above. This is in part due to the physical scope of this paper as well as the fact that Bitcoin more than other technologies, similar to the internet, is decentralized and runs contrary to centralized companies in the first place. For the same reason, while using the example of Visa, this work addresses the credit card industry as one, not a single company. We thus leave the potential question of adopting Bitcoin from a business perspective to future research.

1.2. Studying Bitcoin

Given the relatively short history of Bitcoin, which at the time of writing is about ten
years, the body of academic literature that addresses Bitcoin remains relatively small. The existing literature is separable into two broad though overlapping categories, one focusing on exploring and describing Bitcoin the design and/or technology behind Bitcoin, and the other analyzing and addressing

issues, shortcomings and solutions with the technology either from an inside Bitcoin point of view or an outside societal perspective.

Given how little we still understand Bitcoin the first category of literature is mainly occupied by exploring the design of the technology. Scheuermann and Tschorsch provide one of the most thorough examples of this as they, similar to the original Bitcoin white paper, in their technical review of Bitcoin traverse through the foundational technologies and design decision behind Bitcoin one by one (Nakamoto, 2008). They expand on each of the core parts of the Bitcoin protocol which they define as decentralization, proof of work, the blockchain, transactions, scripts, and recapitulation. Furthermore, the paper frames Bitcoin from in a broader computer science context with respect to security, network and privacy (Tschorsch & Scheuermann, 2016). Additional literature in this category includes explorations from the perspectives of game theory (Kroll et al., 2013), application development (Garay et al., 2015) and monetary economics (Yermack, 2015).

The second body of Bitcoin literature focuses on similar perspectives as the aforementioned literature though with more prescriptive analysis of issues internal and external to the technology. Croman et al. analyzes Bitcoins scalability and compares different proposals to scale the technology to more users and use cases (Croman et al., 2016). Similarly, Decker and Wattenhofer study the importance of information propagation in the Bitcoin network and how to optimize propagation to avoid network forking (Decker & Wattenhofer, 2013). Others investigate some of the claimed features of the technology and the circumstances under which these could be invalidated. The privacy and untraceability features have for instance been tested both from the perspectives of users and outsiders such as governments (Meiklejohn et al., 2013; Reid & Harrigan, 2013, Fleder et al., 2015). Attention has also been given to Bitcoin's security model and how it performs under different circumstances in terms of decentralization of mining power, the number of transactions and block confirmations (Gervais, 2016).

1.3. This Work

So where does this work fit in? As mentioned, our analyze will remain descriptive thus leaning toward the former body of research. Similarly, this work will explore and frame Bitcoin, however, to our knowledge, in a new light; specifically, credit card networks and the competitive differences between the two systems,

as described by disruption theory. The first part of this exploration demands a deep technical dive into Bitcoin as well as credit card networks to create a well-defined ground for comparison between these two highly differentiated systems. In the process, we draw on existing descriptions of Bitcoin and likewise develop our own novel framings of Bitcoin in the light of credit card networks.

2. Methodology

In this chapter, we proceed to establish the philosophical and methodological foundation of the subsequent research analysis presented in this work. We do this to uncover any potential biases, shortcomings, and limitations of the analysis. Moreover, we build on top of the previous literature review, which justifies our choice of theory and presents the reasoning behind our philosophy, approach to reasoning and methods.

2.1. Positivist Philosophy

This work is built on a purely positivist foundation. Ontologically we assume the existence of a real world independent from humans which is observable and measurable using the appropriate methods and tools. As for the subsequent analysis the positivist or naturalist philosophy of science will form the foundation. In an ontological context, that is the study of existence and being, the positivist philosophy assumes the existence of a "real" world which exists independent of human interaction. It follows that it is possible to uncover the "truth" corresponding to the state of affairs in the real world. This is opposed to the constructivist philosophy which assumes that the real world merely is what we perceive it to be (Moses & Knutsen, 2012). Our positivistic world view is expressed through the factual exploration of Bitcoin and credit card networks systems as well as the goal to confirm or deny the former technology as being disruptive to the latter. The uncovering of a relationship between observations and the real world is likewise in line with the epistemological stance of positivism, that is how knowledge is perceived (Moses & Knutsen, 2012).

The positivistic epistemology states that knowledge about the real world is acquired through the identification of relationships between different variables based on empirical evidence, very similar to how the theory of disruption was developed in the first place. Together with the alignment of theory, positivism is also suitable for this work given that we are trying to analyze mainly computer-based systems and correlations between system features and real-world impacts. The belief in through observation and collection is expressed in the positivist methods, such as the comparative-and methods and these will be the tools utilized in this work (Moses & Knutsen, 2012).

2.2. Deductive Reasoning

The deductive approach has the distinct strength compared to inductive logic of building on clear and easily testable assumptions or premises. The canonical example of a deductive argument goes as follows: *All men are mortal. Socrates is a man. Therefore, Socrates is mortal.* By stating the premises behind the conclusion, deductive reasoning doesn't try to infer something from past observations but builds on proofs which each can be rejected if they turn out to be false. Hence it avoids *the problem of induction*; it can never be proved right as a single new observation can totally invalidate any inductive inference (Moses & Knutsen, 2012). A deductive chain of logic is the main approach used in computer science when analyzing system characteristics, where clear and stable relationships exist between variables. For the same reason, this work follows the deductive reasoning in its exploration and comparison of Bitcoin and credit card networks. Practically, we establish a clearly defined set of premises around our comparison of the systems by establishing equality between different elements of the two systems and their functionality, within categories such as state, security, and consensus. Establishing a clearly defined grounds for comparison between the systems is crucial given the large technical differences between them.

2.3. Comparative Case Study

This work builds on a deductive comparative case study. First, each of the two systems will be thoroughly unfolded before we proceed to the analysis. The analysis will compare the systems on a number of dimensions which are more or less central to both systems. With these premises established the application of disruption theory then follows. Hence the analysis might be summed up in the following way: *If two cases differ with regards to [independent variable] then they should differ in the following ways with regard to [dependent variables].* Practically, the independent variable might be how each system handles state while the dependent variable might be transaction and fund security. Conversely, this might then affect adoption from different parts of the market, a central part of disruption theory. Thus, we established a logical chain of reasoning, starting from technological design and ending at market adoption. While empirical data at certain points in the analysis is used to back up and strengthen claims, especially around the market analysis and innovation trajectories, it is the goal that the findings generally should rely on logical reasoning as opposed to induction from certain data points.

3. Fundamentals

In this chapter, we take a deep dive into the technical underpinnings of credit card networks and Bitcoin. We asses the individual parts that make up each system and look at how these interact with each other. Further, the technical features are tied to the practical workings of the systems, both with respect to the average use case and to edge cases. In our case, this means exploring a point of sale (POS) money transaction in each system, the mechanisms involved and potential weaknesses which arise under particular circumstances. Hence, this section lays the foundation for the following analysis.

3.1. Credit Card Networks

We start our technical inquiry by exploring credit card networks and the system which they represent. Though several networks exist we choose to focus on the Visa network in particular. This choice is justified by the large similarities between how large networks such as MasterCard and Visa function as well as the fact that Visa remains the undisputed leader in market share with more than 50% (Forbes, May 23, 2017). We start with some brief though important background information on the monetary system which Visa builds upon before diving into the Visa payment system and how it facilitates electronic payments.

3.1.1. Fiat Money and Banking

Fiat money including the fiat banking model is a huge part of why Visa works as it does. It dictates a lot of the design decisions behind this system and must, therefore, be understood before we can deal with Visa itself. Fiat money or currency includes currencies such as the US Dollar which is minted by a central bank and controlled by the state. By definition, Fiat money, also known as “paper money” is an intrinsically valueless medium of exchange. That is, unlike gold or silver, previously used as money, fiat has no underlying use case or value besides being a medium of exchange. Instead, fiat money derives their value from the support of the state in two main ways. First, to give value to its money, a state can guarantee to accept this money, and only this money, as tax payment from citizens, companies etc. that are liable to pay taxes to the state. Second, the state can likewise guarantee to accept the money as legal tender meaning that the money must be accepted, if offered, in payment of debt in the country. Third, to build trust in its

money and the associated banking system, discussed subsequently, the state and/or its central bank can guarantee to redeem all money deposits, to prevent bank runs (Lerner, 1947).

An important feature of fiat money is money supply given the ability to change this variable compared to for instance gold where supply is more fixed. In a fiat money system, there are two main ways of regulating the amount of money currently in the system. The majority of injection or inflation stems from the banking system and the concept of fractional reserve banking. At its foundation, a bank works by accepting money deposits (taking debit) and lending the same money to others (giving credit) (Bank of England, 2014). Usually, states regulate banks by dictating a fraction of the deposits which banks must hold in reserve, hence fractional reserve banking. If this fraction is 10% the banks can lend out the other 90%. These 90% most likely end up in another bank where 90% of the 90% can be given as credit, a mechanism which is called the multiplicator effect. Hence, when a bank gives credit it is essentially creating new money given the multiplicator effect. The second main type of money supply regulation is that of central banks and this is usually the one employed by states. Effectively central banks own the money press why they can inject or remove money supply as they want. They do so, either by adjusting the interest rate, which affects how much credit banks give, through purchasing or selling government bonds or by more extreme measures such as quantitative easing (McLeay et al., 2014).

The common denominator when dealing with fiat money is the state. Whether it is about giving value to the money, regulating their supply or controlling their use, the state has the power to and in many cases the need for controlling the money. All activity that takes place using fiat money, at least electronically, is to some degree connected to and regulated by the state. State institutions enforce the state's sole right to control and make money. Trust in the money is induced through law enforcement agencies and the legal system. A consumer that has been cheated in a transaction involving fiat money can through the courts force the other party to reimburse them. In the US the Department of Homeland Security can seize any persons banking accounts and the funds in them without notice if, for instance, the owner has been deemed a national security threat. The power of fiat money is therefore highly centralized in the state and from this derives the need for a large institutional body to enforce said power, a power which at its foundation relies on the states monopoly on violence (Davidson & Rees-Mogg, 1997)

The state-centric nature of fiat money dictates what it for the average person actually means to possess money electronically in a bank account. At the most basic level banks keep track of account balances using a ledger, tracking all debit and credit. From a technical perspective, this is implemented through a relational database containing a wide range of relational tables with different information such as customer information, account information, transaction information, account balances and more. Though the ledger has many copies for security and accessibility purposes the database containing the ledger is highly centralized and it needs to be in order to ensure it against fraud and data corruption/loss. There is most likely a single master version which determines the state of all replicas. Logs are recorded and stored to be the entire database or parts of it recoverable if a failure should occur.

Given the highly centralized nature of the database, elaborate security measures are demanded to secure the ledger from malicious and accidental interference. Techniques such as firewalls, data encryption, access authorization, backups, and replication are employed in this respect. Given how centralized the ledger is, changing it, for instance in order to decrease an account balance as the result of a credit card payment, is as simply updating a value in a table in the master database. Little enforcement or control, therefore, happens inside the database (Batiz-Lazo & Wood, 2002; Consoli, 2005; Samakovitis, 2012). Hence, like with fiat money in general, the structure around the technical system is integral to the success of the system. An example is when a loan is made from the bank to a customer. Instead of just transferring the money by changing the database (increase the account balance of the customer, decrease that of the bank), the transaction has to be accompanied by an elaborate contract which sums up the agreed upon terms of the loan, which later can be enforced by state institutions if disputed.

3.1.2. The Life of a Visa Transaction

The above picture of fiat money based banking represents one side of the credit card model; offering credit to your customers, managing their accounts and processing their payments. We could call this card issuing. The other part of credit cards is all about enabling merchants to accept payments through the same credit cards and get reimbursed. We could call this merchant acquiring and this is where credit card networks such as Visa comes into the picture. Any bank can issue credit cards to customers with accounts and any merchant is free to accept cards from any bank. The problem becomes when all cards have different rules, design, and interfaces which each merchant have to adapt to. Credit card networks as Visa

solves this problem by defining a standardized schema for the structure of a payment. Visa is first and foremost a schema for defining credit card transactions. Comply with the schema and you can send or receive with other parties that adopt the same schema. Second, Visa is a transaction router. That is, if your transaction complies with the schema and takes place on the with Visa accepted payment methods, the Visa network will route the transaction from point A to point B. With Visa consumers and merchants using different card issuers can transact with one single standardized payment method. Below we proceed with a step-by-step walkthrough of a transaction on the Visa network to see exactly how this is done.

3.1.2.1. Authorization

The authorization phase is the first of two major phases of a Visa transaction. At the face of a transaction, to a consumer, this phase makes up the transaction in its entirety. It covers the whole POS process for the payer, the rest is taken care off later by other parties. Step 1 of the authorization phase is initiated when the Visa cardholder "sends" a payment, by presenting their card to a merchant, for instance through an HTTP request online or swiping the card in a compatible card reader. Being a relatively simple technology, a large number of security measures are used to secure payments, including the primary account number (PAN), acting as a unique identifier, the PIN code and/or a signature (Security Research Labs, 2018). Having received this information, the next step for the merchant is to process it and, based on this, request an authorization from the merchant bank, in this context known as the acquirer. This is step 2 (Visa, 2016).

The acquirer is the bank in which the money received from Visa payments are stored, in other words, the bank of the merchant. The point of involving an acquirer bank is to transfer some amount of transaction risk away from the merchant by transferring the funds to the merchant immediately, even when they are still to be received from the consumer. Instead the merchant bank acquires the receivable (future payment) from the merchant for a fee. Step 3 sees the acquirer submit the authorization request, received from the merchant, to the Visa network. As step 4, Visa receives the payment request from the acquirer and routes it to the card issuer, that is the consumer bank, for a fee. This is the before mentioned point of standardization where matching schemas between credit card payment, merchant bank and consumer bank creates a standardized interface for all parties to transact (Visa, February 13, 2018).

In step 5 the issuing bank receives the payment request from Visa. The bank then checks in the digital ledger the account involved in the transaction, mainly to determine if there are enough available funds. Moreover, the bank looks in an exception file where special flags about the account can be raised with respect to account suspension and more. If all these checks are positive the bank freezes the payment amount in the cardholder's account, to be unfrozen when the transaction is actually exercised (see 9.1.2.2. Clearing and Settlement) or abandoned, and responds to Visa with an approval. In step 6 Visa forwards the approval to the acquirer which as step 7 forwards it to the merchant. Finally, as step 8, the merchant receives authorization response and completes the transaction accordingly (Visa, February 13, 2018).

3.1.2.2. Clearing and Settlement

The second main phase of a Visa transaction is the often unseen and overlooked process of clearing and settlement. While it from the authorization phase in many ways appears as if the transaction has been completed, this is not the case. Authorization is nothing more than the approval of available funds with a promise to pay these funds at some future time, justifying the aforementioned need for an acquirer to transfer risk. Clearing and settlement is what actually moves the money from one account to another (Visa, February 13, 2018).

We pick up the transaction walkthrough immediately after step 8 in the previous section when the merchant has received the payment authorization and the POS interaction is completed. Step 1 of the clearing and settlement phase starts when the merchant, at the end of the day, proceeds to send all approved transaction authorizations for that day to the acquirer in a batch. This is the action of capturing the authorizations. As step 2, the acquirer credits merchants account and submits the transactions to credit card network for settlement, that is the deliverance of the promised funds from the consumer to the merchant. Step 3 facilitates settlement by paying the acquirer and debiting the card issuer before, as step 4, the credit card issuer posts the transactions to the cardholder’s account and subtracts the funds from their balance (Visa, February 13, 2018).

Finally, the funds can now be said to have changed hands and the transaction has been completed. Meanwhile, given the technical, contractual and institutional features of the fiat based banking system, a

transfer of funds can never be said to be 100% final. Given the right circumstances, contracts and institutions can dispute and roll back transactions, something we will touch upon in the next section.

3.1.2.3. Edge Cases

In a fiat based electronic cash system, a transaction can never be said to be truly final, neither theoretically nor practically. Given the importance of state institutions and contracts in this system and the centralized and mutable nature of the transaction ledger, the finality of a transaction is dictated by the contracts that guards it. If it represents an edge case by falling outside of what might be considered a normal transaction, it is subject to different conditions than normal transactions.

The prime example of such edge cases is chargebacks. As a part of the Truth in Lending Act in the US, chargebacks mainly exists as a consumer protection device. It obligates merchants to reverse a specific previous transaction and return the charged funds to the consumer, given the right circumstances. These circumstances can be grouped into four categories: 1) fraud, when the consumer claims to not have authorized a transaction, or identity theft; 2) consumer quality disputes, when the consumer claims not have received the goods as agreed upon at the time of purchase; 3) processing errors, when the consumer claims to have been charged the wrong amount, charged more than once or never to have received a refund; 4) authorization, when the transaction was declined or there were insufficient funds. In the Visa network, the chargeback processes is initiated when a cardholder files a transaction dispute with the card issuer, based on one of the above circumstances. Through Visa, the card issuer forwards the dispute to the merchant which either accepts the dispute and repays the disputed amount or rejects it and submits supporting evidence to Visa which makes the final decision (Visa, October 10, 2018).

Chargebacks protect the consumer but leave merchants vulnerable as they never can rely on a transaction being final. In the case of fraud-based chargebacks, risk arises both from genuine fraudulent transactions and transactions which the consumer claims to be fraudulent, but in reality are not. The main merchant protection remedy of Visa is machine learning models which attempt to identify and reject truly fraudulent transactions. That is, when the transaction passes through Visa (step 4 in the authorization phase) the transaction and its metadata are analyzed and compared to previously fraudulent transactions to either forward it or reject it on the spot (Steensen, 2018). This solution essentially seeks to minimize

one of the issues with a highly manipulatable ledger, where transactions generally are governed by contracts and institutions.

3.2. Bitcoin

Since its invention in 2008 much has been written about Bitcoin, what it is, what it represents and how it works. However, given the novelness of Bitcoin, the technologies that constitute it and their unique combination and interactions much remains to be discovered about the inside dynamics of Bitcoin beyond the obvious contours of the surface. At a high level, this work takes it starting point from a very textbook definition of Bitcoin from the first page of Mastering Bitcoin by Andreas Antonopoulos (2017a):

> *Bitcoin is a collection of concepts and technologies that form the basis of a digital money ecosystem.*

We believe this definition captures the essence of Bitcoin while leaving enough room for exploring the system in accordance with the scope of this work. Adopting this definition, we can explore the subparts, the concepts and technologies at the basis, for later analysis and comparison. We start our exploration with what arguably are the three main pillars of Bitcoin; the network, the blockchain, and proof of work. This part can be likened to the exploration of fiat money and banking in the previous section. Similarly, we follow this by a step-by-step walkthrough of a Bitcoin transaction before wrapping up with a look a new part of the Bitcoin ecosystem, the lightning network.

3.2.1. Core technologies

While many elements, concepts, and technologies of Bitcoin can be deemed important and even indispensable, three parts stand out; the network, the blockchain, and proof of work. All essential, yet useless without the others, they encompass a large part of how and why Bitcoin looks as it does.

3.2.1.1. The Network

The Bitcoin network has a peer-to-peer (P2P) architecture. It consists of nodes (computers) which follow the Bitcoin protocol and speaks the same language by running compatible software, for the most part, bitcoin core, the current reference implementation of the protocol (Bitcoin Github repository).

However, other implementations are usable as long as their transactions, blocks, and communication are compatible with bitcoin core. Being P2P all nodes are equal, share the same burden and have the same rights and functionality. The network has a flat mesh topology meaning that there are no central servers; nodes connect directly with their peers. This topology is by design meant to foster decentralization of power; all nodes free to do as they want on the network as long as they follow the protocol. Besides the P2P network, other protocols exist which extends or builds on top of the network, for instance, the lightning network which is touched upon in one of the following sections (Antonopoulos, 2017d).

While all nodes in the P2P network are equal they can choose to take on different roles by using one or more of the functionalities in the Bitcoin protocol; routing, the blockchain, mining and bitcoin wallets. From these functionalities, we can further define two main categories of nodes: full nodes and SPV (simplified payment verification) nodes. All nodes include the routing functionality in order to participate in the network. They validate and propagate transactions and blocks and help other nodes stay connected by introducing themselves and others to new peers. Full nodes, also store a complete copy of the blockchain which they continuously keep up-to-date as new blocks are received. Hence, full nodes have the ability to autonomously and authoritatively verify or reject any transaction without external reference and help. On the contrary, for ease of hosting, SPV nodes only store block headers and its own transactions, not other transactions, why they can only verify new blocks on the surface but not the validity of the transactions within them. Nodes with mining functionality can both be full nodes or SPV nodes if they participate in a larger mining pool, and these nodes work to create and mine new blocks to earn bitcoin (Tschorsch, 2016).

When we start up our Bitcoin node for the first time the protocol first connects to the P2P network and establishes an encrypted TCP connection to a random node already on the network (found by querying from a list of DNS seeds) and completes a handshake. This node forwards our address to other nodes making us more connected and known in order to establish a diverse set of paths into the Bitcoin network. If we choose to be a full node, we then start to construct on top of the genesis block, statically embedded in the software, downloading the blockchain from our peers and verifying all transactions and headers while spreading the load across our peers. Our node will keep looking for and downloading blocks until our block height is equal to that of our peers (Antonopoulos, 2017d).

3.2.1.2. The Blockchain

The blockchain is probably the most well-known and hyped part of Bitcoin, yet the simplest and least novel of the three core technologies. Stripped from the other core technologies, the blockchain is simply a data structure, though a data structure which fits perfectly into Bitcoin. While the current dominant implementation of the Bitcoin software, Bitcoin Core, uses Google's LevelDB, the blockchain can be stored in something as simple as a flat text file. The standard linked list data structure is easily recognizable in the blockchain, as the latter represents an advanced implementation of the former, specifically a back-linked list. As a linked list, the blockchain is made up of nodes or blocks each containing some data and a pointer to the previous block which represents the chain. Bitcoin's blockchain contains two types of information which are embedded into each block; metadata and transaction data. While we take a deeper dive into the anatomy of a transaction in one of the following sections, a transaction essentially contains inputs, that is addresses where the bitcoin currency is sent from, outputs, the address where the bitcoin currency is being transferred to, and a digital signature verifying that the sender has actual possession of the bitcoin (Tschorsch, 2016).

Besides transaction data, the block has a header which contains metadata about the block and its transactions, specifically current mining difficulty, mining timestamp, nonce which acts as a mining proof (see section on proof of work) the root of a Merkle tree summarizing all transactions and the double SHA256 hash of the previous block header. The last two are worth exploring in a bit more detail. First, a Merkle tree is a data structure which allows one to summarize and quickly validate large amounts of data with very little effort. Visually, in Bitcoin, it works by arranging the hashes of transactions in a block side by side, as the leafs of a binary tree, and then hashing every pair of two with each other and treating the new hashes as the leaves. Doing this repeatedly the number of hashes is halved with every iteration until only one root hash remains, summarizing every single of the original transactions. This root is unique to this exact set of transactions and changing one of them will, therefore, invalidate the root. The Merkle root, therefore, allows everyone on the network to almost instantaneously validate that the transactions of the block haven't been changed, as opposed to checking every single one of them (Antonopoulos, 2017b).

Second, the hash of the previous header has the function of creating a continuous and chronological transaction history as well as allowing for easy validation of new transactions and protection from attempts to change transactions in old blocks. By hashing the previous header we ensure that nothing can be changed in that block, neither transactions as summarized in the Merkle root, nor any of the other metadata, without invalidating it to all subsequent blocks which builds on the same blockchain. Moreover, as the header of the previous block itself contains the hash of the block before it, and that a hash of its previous block, we ensure that our new block and all its followers effectively will reject all changes to any previous block no matter how obscure and minor this change might be (Antonopoulos, 2017b). Chaining blocks through headers (including Merkle tree of transactions) means any change to a single transaction in the past will change (invalidate) all subsequent blocks.

3.2.1.3. Proof of Work

The arguably most important pillar of Bitcoin, as well as the most difficult concept to grasp, is proof of work. As mentioned, the blockchain alone only holds the transaction data of Bitcoin and while this data structure is perfectly suited to its use case there is theoretically no reason why it couldn't be substituted by another data structure in an alternative implementation of Bitcoin. On the other hand, proof or work, also known as mining, is arguably what makes Bitcoin special, as a functioning system in itself and as a technological breakthrough in computer science (Antonopoulos, 2017c).

Without proof of work Bitcoin would be vulnerable to Byzantine faults; in a distributed computing system with imperfect information, how do the participants agree on something in order to avoid catastrophic system error when some nodes are unreliable or malicious? Achieving Byzantine fault-tolerance, generally considered as one of the hardest problems to solve in distributed computing systems, means that the system continues to function correctly even as some nodes are malfunctioning, not obeying the protocol or acts maliciously (Kleppmann, 2017a). To Bitcoin specifically, the problem lies in agreeing on the current state of the blockchain, that is what bitcoins belong to what addresses, when many nodes are incentivized to lie in order to gain more Bitcoin. For instance, Alice might make a transaction with Bob giving Bitcoin and receiving a good, then proceeding to signal the other nodes on the network a state where this didn't transaction didn't take place and, if successful, ending up with the good plus the same amount of Bitcoin as before. This is a variant of what is called a double spending

attack and failure to avoid these would mean catastrophic system error for Bitcoin given its function as a trusted currency. Additionally, even without malicious nodes such as Alice, the imperfect information of a distributed system such as Bitcoin makes it prone to forking, where different nodes have diverging sets of the blockchain, incompatible with each other (Tschorsch, 2016).

To Bitcoin and the distributed systems branch of computer science, proof of work offers, if not a theoretically complete solution, then a practically good-enough solution to Byzantine faults double spending, and forking. The goal is to secure the blockchain and the means is the creation of new bitcoins and harvesting of transaction fees. When a miner, a node wishing to gain more bitcoins through the creation of blocks, receives unconfirmed and valid transactions over the network it stores them, called the memory pool. When a new valid block is announced over the network, the node collects transactions (usually with most fees and with size less than the maximum block size) from the memory pool into a new candidate block, including a special coinbase transaction which rewards the node itself with new bitcoins and fees. It then creates the block header for the candidate block, at this point without a valid proof of work. The amount of the coinbase transaction is an important part of the Bitcoin protocol individually calculated and enforced by nodes. It is a decreasing function of the block height decreasing 50% every 210,000 blocks and failure to comply with it deems the block invalid (Antonopoulos, 2017c).

Having constructed a candidate block the node starts to mine it, repeatedly hashing the header of the previous block while adjusting an input variable (nonce) in order to get a hash output which is below some target value. This target value, or difficulty, is adjusted so that the average block mining time converges on 10 minutes. The SHA256 algorithm is deterministic yet completely unpredictable in its output given an input. Being unpredictable, the network as a whole can take the nonce and hash as proof of some amount of work put into mining the block. Being deterministic, once a miner finds a nonce that generates a hash under the target value, this can quickly be verified by other nodes by running the function with the same nonce and block header. It takes many hashes to produce the nonce and only one to verify it. Hence, when the first mining node finds a valid nonce and hash it adds them to the header of the candidate block and immediately broadcasts is to the network as it wants to be first and capture the coinbase transaction. When other nodes receive the block they check the validity of its data structure,

hash and nonce, timestamp (less than two hours in the future), block size and transactions (Antonopoulos, 2017c).

Having validated the block, other miners immediately abandon the current block they are mining to put together a new block, and the cycle continues. While they could keep mining the old block and potentially create a fork with two blockchains their best interest is know to build on top of the new block as it by default is the chain with most work that is valid. Every miner, therefore, have financial incentives to follow the same blockchain (with most work), or else their mined blocks will end up as part of an invalid chain, and their work in the form of mining hardware and electricity will have been wasted. For the same reason, double spending attacks become financially unsustainable as Alice would have to go one block down the chain, before her transaction with Bob, and then overtake the current valid chain through mining. Such an attack, known as a 51% attack, is only possible with a large share of total mining power (say above 30%) and would likely lead to decreased trust in Bitcoin and a decrease in the bitcoin price. While not proved, it thus seems unlikely that the financial gains of such an attack would outweigh the losses (Tschorsch, 2016).

3.2.2. The Life of a Bitcoin Transaction

Contrary to fiat currency no coins technically exist in Bitcoin. Instead, all we have are addresses and transactions. A coin is in the words of Nakamoto "a chain of digital signatures" (Nakamoto, 2008, p. 2). A signature is the verification of a transaction from one address to another. By tracking these transactions, first from Alice to Bob and then from Bob to Carol bitcoin wallet software makes it easy for Carol to establish how much bitcoin her addresses currently control. But this is really just a high-level abstraction of the underlying blockchain. It is Carol's control of the most recent address in the chain which gives her the ability to sign the next transaction in the chain. There are no bitcoins in her wallet and through the use of new addresses for each transaction, there is no apparent way for anyone else besides Carol to tell what bitcoins are controlled by her.

The closest we can get to defining a bitcoin as a static entity at a given point in time is the output of the most recent transaction in a chain of signatures, known as an unspent transaction output (UTXO). This is true given the structure of a Bitcoin transaction; it takes an input and produces an output which then

acts as the input of the next transaction in the chain. The transaction input consists of one or more UTXOs, essentially representing some amount of bitcoin which are available to spend. Specifically, the input contains the transaction ID(s) of the transaction where each UTXO comes from as well as the index, specifying which UTXO(s) from that transaction. If the transaction is successful we say that the UTXOs have been consumed, and they no longer classify as UTXOs. It also contains a signature verifying and unlocking the UTXO(s) in the shape of an unlocking script. While different types of unlocking scripts exist, the basic idea is to verify the ownership of the input. Using public-key cryptography, the owner of the input signs the hash of each of the previous transactions (where the UTXOs came from) with their private key. The output of this unlocking script is compared to the output of the locking script, a spending condition placed on the UTXOs when the previous transaction was made involving the corresponding public key. If the keys indeed are a pair and the scripts produce the same outputs, the transaction is validly signed to the network (Antonopoulos, 2017e).

Given a validly signed set of inputs, the output of a transaction creates one or more new UTXOs available for consumption as new transaction inputs. The output consists of two parts; an amount of bitcoin which the UTXO represents and the previously mentioned locking script involving the public key of the receiving address. Often a transaction will produce an output with more than one UTXO even when involving only a single receiving address. This is because UTXOs by design are indivisible why they only are consumable in their entirety by a transaction. Hence, if Alice wants to transfer Bob 2 bitcoins with one input UTXO of 3 bitcoins, she has to create two output UTXOs; one for Bob with 2 bitcoins and one back to herself with 1 bitcoin. Moreover, the output of a transaction will almost always be smaller than its input. The difference between the two is what is collected by the miner as fees. Strictly speaking, all bitcoins originate from a transaction output, not input. That is, the origin of all current UTXOs can be traced back to a coinbase transaction, put in a block by the miner of that block. Coinbase transactions represent money creation in Bitcoin as they produce UTXOs without any input (Antonopoulos, 2017e).

The corresponding unlocking script, described above, mirrors the Pay-to-PubKeyHash (P2PKH) script, which represents a vast majority of transactions. Meanwhile, a unique feature of Bitcoin transactions is the ability to program other types of scripts with different functionalities. One example of this is Pay-to-ScriptHash (P2SH). Contrary to P2PKH where the sender of the transaction specifies the

unlocking script of the output as sends it directly to the receiver address, P2SH transactions allow the receiver to specify the output script as the sender sends the transaction to a hash of script (called a redeem script). This allows the receiver to receive bitcoins to an address that is secured in various unusual ways without letting the sender know about this, enhancing security, privacy, and anonymity. For instance, this could be to a multisignature address, another transaction type, compatible with P2SH. While one might call P2PKH transactions single signature transactions, as they involve only one signature to verify, multisignature addresses take m signatures out of n specified to be verified. This enhances security around addresses where it is not ideal to have a single person controlling that address (Tschorsch, 2016). Finally, a category of noticeable and highly useful transaction scripts are timelocks, a scripting primitive used to restrict the consumption a UTXO until some time or block height. For instance, scripts can use the nLockTime opcode to define the earliest time a transaction may be added to a valid block while CheckSequenceVerify specifies the relative amount of blocks from a transaction was added to the blockchain, that its output can be spent. As seen in the next section, this is extremely important in dual payment channels where we want to avoid one of the parties stealing all of the money in the channel (Antonopoulos, 2017e).

3.2.3. The Lightning Network

Recalling the definition of Bitcoin at the start of our exploration of the system, Bitcoin is a digital money ecosystem. While the Bitcoin protocol represents the foundation of this ecosystem, there is room for other specialized protocols extending or building on top of the base protocol. The most impactful and noticeable element in the ecosystem is arguably the lightning network, a layer two solution which builds on top of the Bitcoin protocol. The lightning network is the most successful implementation of the broader concept of payment channels, an idea which has gained relevance in recent years as the Bitcoin network has gained adoption. The increasing adoption in terms of transactions has at times seen Bitcoin fees increasing significantly, potentially threatening the viability of bitcoin as a medium of exchange and a potential substitute to fiat currencies and credit card networks. This is because of the inherent constraints of the current design; one block every 10 minutes and a block size limit of around 1MB. While some have argued for an increase of the block size as the way to increase transaction throughput in the blockchain, payment channels and the lightning network has a different design philosophy; keep the blocksize unchanged by moving transactions off it from the base layer to the second layer (Poon & Dryja, 2016).

The lightning network works by locking funds in UTXOs on the blockchain and then establishing bidirectional transaction channels between different parties. The parties can transfer funds safely and trustlessly between each other in the channel an unlimited number of times without publishing the transactions to the blockchain. Furthermore, parties can route transactions between parties they have channels with. Hence, if Alice has a channel with Bob and Bob has a channel with Carol and Alice wants to send Carol money she can do so through Bob who is connected to them both. A payment channel is established through a funding transaction which is committed to the blockchain. Both parties of the channel send their funds to a 2 of 2 multisignature address meaning that none of them can take the funds without the consent of the other. Immediately before the parties make the funding transaction they also construct a refund or commitment transaction which only is valid some time in the future using the nLockTime opcode, which they only send to the other party and gets signed. This is to avoid that their funds will get locked up in this address if the other party disappears (Poon & Dryja, 2016).

With the initial commitment transactions signed and the funding transaction on the blockchain, the parties can start transacting via their payment channel. With every new transaction they exchange new commitment transactions dividing the funds of the funding transaction between them, each with a slightly shorter nLockTime than the previous commitment transaction, ensuring that the most up-to-date state of the channel always will get executed should one of the parties disappear. They also exchange something called a revocation key for the previous commitment transactions. Another important element of the commitment transactions which we haven't touched upon is that they include CheckSequenceVerify in the redeem script. The script specifies a delay of 1000 blocks after the transaction has been made before the party closing the channel (i.e. broadcasting the commitment transaction) can claim the funds. Meanwhile, another party can redeem the funds immediately if they have a revocation key. This redeem script along with the exchange of revocation keys for old commitment transactions means that if any party attempts to cheat by broadcasting an old channel state the opposite party will have 1000 blocks to use their revocation key to claim all the funds in the channel (Poon & Dryja, 2016).

The lightning network uses the time dimension of the Bitcoin blockchain to provide highly scalable and completely trustless payments. As long as the blockchain is decentralized and resistant to double spending and 51% attacks it is practically as safe as using the base protocol, and much cheaper. With a well-connected network, fees will be able to stay around zero; routing involves only a few intermediary nodes and new routing nodes can easily be added in the face of higher demand (Poon & Dryja, 2016). Moreover, lightning routing uses onion routing which is highly privacy conversing and as nothing besides the funding transaction and the final commitment transaction is recorded on the blockchain, transaction becomes practically impossible to track, even using correlative machine learning techniques which recently have been applied with some success to the blockchain (Harlev et al., 2018).

4. Analysis

Proceeding in our deductive chain of reasoning, this chapter moves from technically defining Bitcoin and credit card networks to comparing and categorizing them in the context of disruption theory. In line with disruption theory, we first analyze and compare their technical design and features given a number of dimensions from distributed computer systems. Second, we apply the core part of disruption theory as we use our established premisses around their technical differences to further establish how Bitcoin and credit card networks compete on key product features and conversely how their strengths and weaknesses map to various customer segments of the market. Third, we proceed to analyze how Bitcoin and Visa have in the past and will in the future innovate around these key product features. This brings us in the position to finally answer the question of whether Bitcoin is disruptive or sustaining to Visa and other credit card networks. As its last part, the chapter turns to a critique of the undertaken analysis and its findings, underlining the most obvious counterarguments and limitations.

4.1. Comparing Bitcoin and Credit Card Networks

As the first part of the analysis, we compare Bitcoin and Visa on a number of dimensions. As mentioned in the methodology section, given the large differences between the two systems, these dimensions have been picked based on relevancy, from both systems as well as computer science in general.

4.1.1. Consensus

The concept of consensus is central in computer science, specifically within the branch of distributed systems where numerous dispersed entities have to agree upon a shared world view (Kleppmann, 2017b). The concept is native to Bitcoin, given its highly distributed design in terms of state and protocol. All Bitcoin nodes are equal in so far as they can run the protocol implementation they want thereby verifying or rejecting the transactions and blocks they wish to. The design of the protocol means that they can do so without consulting or complying with any single centralized entity. This was demonstrated in 2017 with the events around the Bitcoin and Bitcoin Cash chain split followed by the failed SegWit2x network upgrade. A number of large transaction processors and miners both times attempted to change the protocol of the whole network without broad consensus by increasing the block size limit. Both attempts

failed; the first attempt led to a permanent split in the network while the second was called off before execution (Wirdum, November 8, 2018).

The Bitcoin protocol, the blocks, and transactions which are accepted, or rather changes to this, is therefore driven by broad consensus. Through its consensus model, the protocol is resistant to Sybil attacks (subversion through the creation of pseudonymous nodes); there is no voting, there is only the software which nodes run and the network(s) which emerges from their choices. Meanwhile, proof of work is used in the continuous process of reaching blockchain consensus and this is only as decentralized as the decentralization of mining power on the network. Proof of work avoids persistent chain splits where different parts of the network diverge and become incompatible. Given the coinbase transaction and fees, it incentivizes all miners to follow the same chain. Worth noting is that the lighting network only utilizes the Bitcoin consensus model in the funding and closing transaction of the channel. In between funding and closing, it is solely the two parties of the channel which have to reach consensus about the next commitment transaction. If consensus can't be reached they can go to the Bitcoin network and use proof of work consensus and close the channel.

Compared to Bitcoin, the consensus model of Visa and fiat-based banking is radically different. In the context of the Bitcoin network topology, we can treat Visa as the network and the card issuer and merchant banks as the nodes. With respect to the schema, that is the protocol which banks must comply with to have their card transactions be compatible with the Visa network, the power is highly centralized in Visa and demands no consensus. Banks can suggest changes but ultimately Visa can do as it wishes as long as it doesn't lead to banks abandoning their network for another. In theory, this may not seem that different from Bitcoin, where nodes likewise can abandon the protocol implementation if they don't agree. In practice, however, the open source nature of Bitcoin makes this far easier to do than in Visa where banks potentially would have to build a completely new network from scratch, including merchant adoption. The result is that far less of the power in the Visa network is held by the nodes, compared to Bitcoin where any nodes can fork the protocol within seconds.

As opposed to being verified by every single (full)node, as in Bitcoin, the verification of a Visa transaction is centralized in one entity, namely the issuing bank which singlehandedly controls the balance of the

cardholder. Ultimately, there is additional oversight from contracts and state institutions which guarantees that banks don't cheat customers, but the point stands that the system is centralized. It is worth noting that the ledger within a card issuing bank indeed is replicated several times which creates the potential for utilization of decentralized consensus algorithms, through a centralized leader-slave typology seems more plausible as the utilized consensus method. However, such decentralization takes place at a lower level of the network and is expressed through a single centralized actor.

Meanwhile, the clearing and settlement process is slightly less centralized compared to the verification, involving issuing and merchant banks plus Visa as the intermediary. The centralization of the Visa network negates the consensus issues which Bitcoin solves through proof of work. That is, by having a single entity controlling each part of the overall ledger there is no risk of different entities having different ledgers and diverging from each other; there is only one.

4.1.2. State

State is a central concept of most computer systems. Whether a fully fledged database used to persistently store information in a large application or a simple data structure for creating session state around requests from the stateless HTTP protocol, the creation and use of state is essential to the value of these systems. In Bitcoin, state is persisted on the blockchain, in what we before defined as an advanced linked list which is stored on disk. New blocks refer to previous blocks and state is kept between them by nodes. The blockchain is then Bitcoin's durable state; it is remembered across nodes and failures (Helland, 2018). Given adequately distributed mining power, one might even be tempted to call this state immutable (Helland, 2015); though new data in the shape of blocks is added to the blockchain on average every ten minutes, the existing blocks don't change once added. As a block is preceded by more blocks the amount of work (as in proof of work) accumulates on top of it, making it exponentially harder to change it. Hence the general rule in Bitcoin that a transaction generally is immutable after around six blocks.

If the blockchain, a shared or at least converging and linked state between all full nodes, represents the durable state of Bitcoin, then the memory pool arguably represents Bitcoins application state; in memory data of the running application, not persistent across failures. For miners to be able to quickly assemble new valid blocks and for other full nodes be able to quickly validate these, all full nodes validate and store transactions as they receive them in RAM for quick retrieval, why it is called the memory pool. The

memory pool is then where transactions are stored intermediately before they are included in a mined block and persisted on the blockchain, at which point they are removed from the memory pool. In the context of the lightning network, we can arguably equate a session with a payment channel as it represents intermediate states of a persistent Bitcoin transaction. Similar to session state, payment channel balances exist across updates and only in the two participating nodes, the session endpoints (Helland, 2018).

Around a Visa transaction, durable state is stored by issuing and merchant banks in their relational electronic ledgers. Unlike Bitcoin, one cannot refer to this data as immutable to any extent given the ability of banks and state institutions to overrule and withdraw already made transactions, for instance in the case of the aforementioned chargebacks. As accounted for in the walkthrough of a Visa transaction, a number of steps are involved leading to the durable state of transactions. These intermediary steps are equal to the session state of Visa. In the authorization phase, the issuing bank will check the cardholders account for sufficiently available funds and, if present, lock the amount in the account. The merchant will thereafter receive payment verification from the bank, through Visa. The transaction isn't durable yet, we only established session state between the two endpoints, the merchant and the cardholder, represented by their banks. Similar to a lightning channel, they may transact several times before the session is made durable, essentially updating the session state each time. It is with the end-of-day clearing and settlement that the session is terminated and the result is persisted in the ledgers of their banks.

4.1.3. Consistency

Both Bitcoin and Visa represent distributed systems, with state and operations spread across geographically distributed nodes. Hence, they are both forced to deal with the CAP theorem (Consistency, Availability, Partitioning) and choose between consistency and availability in the case of network partitions, where the state of different nodes diverge (Kleppmann, 2017b). In turn, this choice is made through the type of transactions that a system implements.

We argue that Bitcoin through its design implicitly chooses availability over consistency given its use of BASE transactions (Basically Available, Soft state, Eventual consistency) (Kleppmann, 2017c). First, being fully distributed, any given node sending broadcasting a transaction to its peers can never be certain that it will be received the first time and recorded on the blockchain. Similarly, an affirmative response

from a peer may not reach the node and may thus try to resend the transaction. Such issues can be remedied through more peers, but no guarantees can be given and the network will remain basically available. Second, given the same P2P structure and propagation of information Bitcoin has soft state from the perspective of a node; the memory pool and blockchain can change when there is no new input from new transactions and mined blocks. Third, in line with its soft state, Bitcoin is eventually consistent. Given the ability for miners to choose transactions to mine and fork the blockchain, purposely or accidentally, the probability of a transaction being durable rises with time. From the time of broadcast, in seconds the transaction will have spread across the network, within minutes it will be included in a block and within hours it will be practically durable as blocks are mined on top of it (Decker et al., 2016). Bitcoin's sacrifice of consistency makes nodes vulnerable to double-spend attacks immediately after the transaction, forcing them to wait for several confirmations for enough consistency. However, it also means that the network never pauses if data isn't consistent, making it fully available to transact on at any given time (Bitcoinuptime.com).

The lighting network, and specifically each payment channel, makes a different trade-off by prioritizing consistency through ACID transactions (Atomicity, Consistency, Isolation, Durability) (Kleppmann, 2017c). First, a lightning transaction has atomicity as it either fails or succeeds as a whole. This is particularly relevant and visible in transactions involving multiple intermediary routing nodes, all collecting some routing fee. The design of a lightning transaction means that no one collects their part of the transaction unless the entire transaction is successful in reaching the receiver. Second, in line with the above example, a lightning transaction is fully consistent at any point in time and can only bring the database (the shared latest commitment transaction) from one valid state to another, no in-between state in the middle. Third, each lightning transaction is fully isolated from other transactions and there is no intermediary state of other transactions they can access. A lightning node can issues as many payment invoices as it wishes but these will be validated in isolation at the time of payment. Fourth, lightning transactions are durable in so far as mining power is distributed and a node doesn't go offline for more than 1000 blocks or roughly seven days, in which case the counterparty would be able to broadcast an old commitment transaction revert more recent transactions (Poon & Dryja, 2016). Opposed to Bitcoin, lightning's prioritization of consistency guarantees full and practically instant consistency between the

nodes of a payment channel making transactions between available nodes extremely fast. However, if the other node is unavailable so will transacting with it be.

Just as the Bitcoin and the lightning network make different tradeoffs around the CAP theorem, so does the subsystems of the Visa network. Moreover, the banking system involves many other transaction channels such as wire transfers and ATMs which blurs the image further. Here we choose to closely draw the boundaries, limiting our system to Visa transactions to and from banks and merchants. Again it seems appropriate to separate Visa transactions into verification and clearing/settlement. The authorization process as mainly involving merchants and a single card issuer is ACID in nature prioritizing a consistent central ledger, similar to the functioning of the lightning network. Each authorization request is routed, via Visa, from the merchant to the master database node of the issuing bank with atomicity, either fully succeeding or failing. Hence, consistency is enforced as the funds in the account of the cardholder either are fully frozen or not. When concurrent requests are received isolation is achieved through locking (e.g. snapshot isolation and two-phase locking), standard to centralized relational databases, making updates serializable and independent (Kleppmann, 2017c). Finally, given a centralized database, transaction authorizations are durable and can survive through recovery from system logs and database replication. Consistency is prioritized around authorization for several reasons. First, to combat fraudulent transactions it is important for the issuing banks to ensure that all transactions are valid. Second, involving relatively few parties and a centralized database it leads to a relatively small loss of availability and transactions speed. The sacrifice of authorization speed for consistency in account balances is favorable, given the banks main objective of decreasing fraud.

Compared to authorization, the clearing, and settlement process of the Visa network is far more complex calling for a different set of tradeoffs. This step sees Visa receiving transactions in batches from different merchant banks. At this point, it is not possible to depend on every transaction having been recorded correctly in the authorization phase; despite the authorizations having been subject to ACID between each merchant and bank, mistakes happen at scale in distributed systems when data is stored different places, say in different ledgers in different banks. Hence, across all these subsystems, in the clearing and settlement phase, BASE transactions are needed (Hoff, May 1, 2013). At the Visa network level, the overall functioning of the network through the day is more important than consistency, which though

desirable, would be too costly to implement across the whole system. Visa collects revenue from transaction processing and are incentivized to increase the number of transactions on the network, and reduce the friction of these though high availability. At the end of the day, with the majority of transactions settled without issues, the protocol works out the exceptions using logs and other types of logic. The network eventually becomes consistent as this happens. From a legal perspective reliance on BASE transactions are only possible given the power of centralized state institutions to reverse fraudulent or wrong transactions when these can't be appropriately handled by the protocol.

Both the Visa and the Bitcoin network prioritizes availability over consistency at their foundation using BASE transactions. Moreover, both use ACID transactions to achieve consistency in other parts of the network. Meanwhile, given how they handle consensus and state, these features are based on very different conditions. Specifically, Bitcoin uses proof of work to achieve eventual consistency in a decentralized manner while Visa uses law and state institutions to do the same. Similarly, the Visa authorization process relies on a centralized ledger to reach strong consistency while the lightning network uses cryptography and the stability of the underlying blockchain.

4.1.4. Privacy

Privacy is a highly relevant concept around systems such as Visa and Bitcoin which handle as important and personal data as how and with whom we transact. We define privacy as the freedom of the individual to choose whom to reveal their data to and what to keep secret.

Given the emphasis on privacy from Bitcoin and its supporters, right from the very beginning, it is no surprise that the concept is central to how the system is designed. Indeed, one of Nakamoto's clearest messages from the original whitepaper is the elimination of trusted third parties (TTPs) as the source of privacy, given that they represent security holes, easy to identify and attack (Nakamoto, 2008). Moreover, given the transparency of the blockchain and the broadcasting of all transactions, privacy is unachievable through encrypted communication with central coordinating TTP to vouch for the identity of transaction participants, similar to how SSL certificates on the internet are handled (Szabo, 2005). Instead, privacy without TTPs is native to the Bitcoin protocol through the use of public key cryptography for Bitcoin addresses. By keeping the private key and the identity behind the public key

secrets, a transaction can't be tied directly to a person. This is further strengthened by the generation of a new address for each transaction, trying transactions together becomes even more challenging. Worth noting is that unlike several of Bitcoin's other features, its privacy component does not rely on decentralized mining power, simply the current mathematical irreversible nature of the SHA-256 hashing algorithm.

Bitcoin's privacy level has, meanwhile, become somehow disputed with the rise and application of machine learning methods to tie together and deanonymize transactions. This risk was originally acknowledged by Nakamoto in the whitepaper around multi-input transactions where several previous transactions are tied to a single address (Nakamoto, 2008). Such transactions are to some degree necessary for Bitcoin users to consolidate funds and avoid UTXOs becoming smaller than the fees demanded to consume them (i.e. Bitcoin dust). As more transactions are tied together in clusters, deanonymizing them all becomes increasingly likely, making Bitcoin addresses more pseudonymous than anonymous (Harlev et al., 2018). To counteract deanonymizing of transactions several measures are being researched and implemented, most noticeably CoinJoin, where several transactions from different senders are aggregated, and Confidential Transactions, where transaction amounts are encrypted but still verifiable, both without the use of trusted third parties (Maxwell, n.d.; Maxwell, 2013). This makes it extremely difficult, though theoretically not impossible, to deanonymize transactions. Moreover, as deanonymizing techniques become more sophisticated the number of transactions and rounds of the CoinJoin can be increased, making the deanonymizing less certain each time. Perhaps the largest privacy gain in Bitcoin comes from the lightning network, as intermediate channel states here are not broadcasted to the network. Using the onion protocol for routing between notes, transactions here become practically anonymous and impossible to track.

Visa and fiat-based banking rely on a privacy model, diametrically opposite to Bitcoin. In this system, transactions are encrypted send and decrypted. The encryption ends points are TTPs with the merchants and banks trusted to store the data of cardholders in a secure and privacy-preserving manner. The raw nature of transactions at these endpoints makes the identity of transactors, amounts and other metadata fully available and potentially vulnerable. For the same reason, merchants and banks employ elaborate security measures to preserve privacy (Szabo, 2005). However, unlike Bitcoin Visa payments demands

consumers to reveal their identity and downgrade their privacy towards merchants and banks around each payment which they, depending on the law, can use for different purposes without sharing it. Moreover, given the authority of states over fiat based transactions, these systems have so-called golden keys, which gives access to personal information, for instance in the case of legal investigations. Besides representing potential areas of exploitation, these further downgrade the basic privacy of the system.

4.1.4. Protocol Stack

Based on our previous comparison of the individual elements and concepts of Visa and Bitcoin we here seek to contrast how each respective system ties them together. In technical terms, we compare the architecture and layers of their respective protocol stacks.

The Bitcoin base protocol serves as the foundational layer of the Bitcoin protocol stack. It is from this foundation that the security guarantees and immutability of transactions originates. Its security derives from a number of design decisions, including a ~10 minute block time to minimize accidental forks, resource-intensive proof-of-work algorithm to support favourable mining incentives, a small block size to maintain network decentralization in terms of full nodes and a non Turing complete scripting language to minimize attack vectors. Moreover, Bitcoin relies on broadcast data transmission as every single transaction and block is sent to and downloaded by all full nodes, enforcing the other design decisions across the entire network. While these features represent some of the most ingenious parts of Bitcoin they also make the system highly unscalable in terms of transaction capacity.

Similar to the internet protocol, data transmission in Bitcoin works similarly to how legacy ethernet hubs functioned in the early days of the internet, where data which entered the port of a hub would need to be replicated and broadcasted through all other ports. Given the later specification and implementation of the TCP/IP protocol most internet traffic today happens through unicast data transmission; instead of broadcasting all data to all nodes, unicast limits data transmission to the sender (server) and receiver (client) using IP for addresses and TCP for routing across the network. Needless to say, unicast over broadcast has greatly increased the scalability of the internet and its applications (Kozierok, 2005a).

The lightning network is becoming to Bitcoin what TCP/IP is for the internet. It builds a stack on top of the foundational Bitcoin layer, analogous to the internet and transport layers of the TCP/IP model. The Bitcoin blockchain is the lowest layer and assembles the lowest layer of the TCP/IP model, the link layer, in so far as that it similar to ethernet represents stable physical connection and communication between nodes, similar to how the Bitcoin network stays connected and updated as a whole. One layer up, the Bitcoin nodes are the internet layer, acting as the endpoints of unicast transactions. Lightning channels between lightning nodes are analogous to established TCP connections, the network layer of the TCP/IP model, in which nodes can conduct unicast transactions. These connections can span several intermediary routing nodes, just like TCP connections. The nodes in the IP layer can choose to transact through broadcast, using the underlying blockchain layer, or with unicast in the lightning layer (Kozierok, 2005b). Thus, lightning enables unicast transactions where only broadcast transaction before where possible.

Unlike on chain Bitcoin transactions, Visa is an example of unicast rather than broadcast data transmission. A Visa transaction is an end-to-end operation only between the merchant the card owner, and a fixed number of routing nodes in the middle. Visa itself then fits within a similar layer to lightning channels, that is the network layer, as it acts as a payment coordinator and router between the nodes in the internet layer. Noticeably, this is not only true for the authorization phase of Visa transactions but also for clearing and settlement which involves solely Visa, the issuing bank and the merchant. This is contrary to Bitcoin where authorization of commitment transactions are unicast but settlement always happens by broadcasting the commitment transaction to the network. The natural question then is if we can find Visa's equivalent to the Bitcoin blockchain, which fits in the link layer. One might be tempted to answer the nation-state, given its ultimate power over fiat currency, but as the TCP/IP model only is applicable to computer networks, it isn't constructive to include it as a layer. Instead, the banks which are the keepers of durable state in the fiat banking system seems the most analogous to the link layer. Banks can move funds around locally within their ledgers for which they use specialized protocols for local, physical transmission.

The resemblance of Visa and lightning transactions as the network layer components of their respective protocol stacks can be further extended into the dimensions of consensus, state, and consistency. In both

systems, it is the two endpoints of the transaction that reaches consensus, not the entire network. While the payment channel is open (in Visa the end of the day, in lightning when the commitment transaction is broadcasted), the intermediate state is likewise maintained by the endpoints only. Both Visa and lightning employees ACID transactions, as opposed to BASE, thereby favoring consistency over availability as they have to protect the validity of their state until said state is made durable. Given this tradeoff, both, therefore, rely on the underlying layer (which for Visa is banks and the state and in lightning is the blockchain) to provide availability and handle edge cases where balances don't add up.

4.2. Bitcoin: Disruptive or Sustaining?

In this second part of the analysis, we apply the business side of disruption theory to Bitcoin and credit card networks. Our aim is to determine if the former technology presents a potentially disruptive or sustaining force to the latter. The comparative analysis of the technical designs of the two systems forms the foundation for this application of disruption theory; only by fully understanding their technical features and how these originate in the design relative to each other can we determine on what dimensions the systems actually compete, for which customers and how this competition is likely to develop in the future.

The last point, namely future development, is worth emphasizing. Specifically, it is of great importance to keep in mind that "disruption is a process, not an event". (Christensen, 2006, p. 46) In our application of disruption theory, we make sure to keep in mind that disruption is dynamic, happens over time and is identified through continuous innovation trajectories. The application of disruption theory begins by identifying the main dimension(s) of competition of the incumbent system as a product of what is demanded by the current core customers. Next, other dimensions of competition, orthogonal to the main dimensions are identified. These dimensions can be combined with market analyses, to locate the current target markets of the systems, and to what degree they meet the demand from different market segments. Finally, based on the previous technical inquiry, we map the key performance dimensions and extrapolate their future development. This last step allows the reasoning about to what degree Bitcoin is disruptive to Visa and credit card networks, and how the target markets of the two might develop in the future.

4.2.1. Competing Dimensions

Here we examine the dimensions of competition between Visa and Bitcoin, linking them to their origin in the technical design of each respective system. Dimensions of competition is a central concept in disruption theory. It is through knowledge about how systems compete that we can assess if the new system offers anything unique compared to the incumbent and thus if there is a foundation to establish and grow market share.

4.2.1.1. Transaction Throughput

Transaction throughput, that is the number of transactions a system can handle within a given timeframe, is certainly one of the most important metrics for a global payment network. A network aimed at enabling POS transactions is of little use if it can't meet the demand for transactions in the first place, no matter what other features it might offer. Moreover, there is mounting evidence that even small changes in online payment processing have a disproportionately large impact on online bounce rates in developed countries, that is to what degree customers leave an online retail site before completing their purchase. For instance, Google research in 2018 reported that wait time of one to 3 seconds increases the bounce rate by 32% while a wait time increased the rate by 132% (An, 2018). In our analysis of innovation trajectories, we will thus treat transactions per second as the main dimension of competition on which the current payments industry, including Visa, is measured. Indeed, Visa seems to acknowledge the importance of transaction throughput and swiftness. In the companies 2018 annual report increasing scalability or transactions per second is mentioned as the first and most important step in the enhancement of the network (Visa, 2018, p.7). Visa seems highly focused on further improvements in transaction throughput through investments in server centers and mainframes. This is evident given the increase in capacity from 56.000 to the current 65.000 transactions per second.

Given Bitcoins security-decentralization-scalability tradeoff, the number of on chain transactions is highly restricted. Given a new block on average every 10 minutes and a theoretical block size of 4MB containing exclusive Segwit signatures it handles a maximum of ~7 transactions per second, several orders of magnitudes below what is required from a global payment system (Croman, 2016). In reality, Bitcoin is currently in the spring of 2019 handling around 2000 transactions per block, or ~3, transactions per

second (Blockchain, March 11, 2019). Moreover, depending on the security requirements of the party at the receiving end of the transaction, both parties will realistically have to wait between 10 and 60 minutes for the transaction to be sufficiently secured towards double spend attacks and confirmed.

Meanwhile, the lightning network, as Bitcoins network layer scaling solution, has the potential for an arbitrarily large number of transactions per second. This is true given the P2P architecture which means that there is no single node, like Visa, which all payments have to pass through towards their destination. However, the transaction capacity of lightning is constrained by other mechanisms. Theoretically, the number of nodes is a constraint and as each node has to handle more transactions there is bound to be a point where a routing node will meet its capacity. Second, lightning transactions relies on sufficient channel balances also called buffer capital. If Alice wants to send a payment of 1 bitcoin to Bob through Carol, whom they both have a channel with, Carol needs to have at least 1 bitcoin locked up in both channels. Hence, talking about transactions per second on the lightning network is only relevant for payments sizes where smaller than or equal to the size of the channels between the payer and the payee (Poon & Dryja, 2016). At the time of writing, the average capacity of lightning channels is 0.023 bitcoins (~$83) while the top capacity channel is 0.17 bitcoins (~$650) (1ML, March 11, 2019a; 1ML, March 11, 2019b).

4.2.1.2. Censorship Resistance

In the context of financial transactions, censorship resistance is the ability of a system to facilitate transactions without the ability for any party to censor or block the payments based on the sender, receiver or other feature of the transaction. In other words, is anyone able to transact freely despite what someone else might think about them or the transaction? While financial censorship resistance might not seem relevant for the large majority of credit card holders, it is crucial to a certain niche group of individuals. In 2011 Visa and Mastercard participated in the US-led payment blockade of whistleblower organization Wikileaks as they stopped processing donations to their account (Holden, October 24, 2011). Recently, a number of individuals have seen themselves banned from the crowdfunding platform Patreon, for some their primary source of income, based on their controversial opinions while the social network Gab saw themselves censored from Paypal and Stripe, given the content and actions of its users (Bowles, December 24, 2018; Robertson, November 5, 2018). On a larger scale, millions of Chinese have

recently seen themselves blocked from buying flight and train tickets due to their unfavorable state determined social credit scores, just the initial step in a much broader plan to censor payments of individuals who don't behave favourably (Cockburn, November 22, 2018).

No matter what one might think about the morality of such censoring, the value of a censorship-resistant system to the censored parties is indisputable. Censorship resistance is naturally linked to that of privacy. That is, in order to censor someone or something you have to be able to locate it. The ability for a system to hide the transaction metadata, therefore, increases its censorship resistance. Privacy is valuable outside the context of censorship too. The degree to which the individual transparently and easily can choose the level of information which they wish to disclose through a transaction is an important feature of a payment system, especially given the amount and sensible nature of the information contained in our transactions.

Depending on the point of view, Visa can be seen as rather censorship resistant or censorship prone. The former view is based on the removal of cardholders, or customers, from Visa. Given the highly centralized nature of Visa including card issuing and merchant banks, consumers don't have access to the internals of the system and have practically no ability to track payments and censor them. At the same time, cardholders themselves are highly vulnerable to censoring from a number of sources including states, Visa, banks, and merchants. As previously discussed, state institutions and agencies are able to confiscate fiat funds and bank accounts, as well as deny the creation of another account including the issuance of new Visa cards. Thereby, they are able to censor payments indirectly, either given their own reasons or on behalf of others, for instance through the court system. Visa can likewise censor transactions during the authorization phase as they pass through VisaNet. Indeed, Visa's main remedy towards transaction fraud is analysis and censorship of transactions using machine learning models which uses transaction data such as sender, receiver, location and transaction history (Visa, 2018, p. 10). Banks themselves use similar risk assessment models to determine if an individual can be accepted as a customer (Addo, 2018). Finally, merchants collect large amounts of information on customers at the point of sale, undermining their privacy and increasing the ability to censor. Merchants are likewise vulnerable. As they need a certified account with an acquiring bank to accept Visa payments, states and banks can block their ability to receive payments and seize their funds.

Compared to Visa, Bitcoin is a highly censorship resistant system. This is first and foremost true, given the pseudonymous privacy of the system, which makes it difficult to identify the sender and receiver of a transaction. With Confidential transactions in the works, even censoring based on metadata such as amount or script could become impossible. As mentioned, user privacy can be further enhanced through the lightning network. Importantly, the above is based on mathematical proofs, probability, and cryptography, not network decentralization. Should the assumption of privacy through pseudonym break, Bitcoin's censorship resistance relies on a decentralized network both in terms of full nodes and mining power. Given the current number of full nodes in Bitcoin between ~10.000 and ~64.000 (including non-listening nodes) the network is highly decentralized and censoring nodes, without a large majority of malicious nodes, by not accepting their transactions is highly unrealistic as all nodes by default have 6 random connection to the network (Bitnodes.earn.com, March 13, 2019; luke.dashjr.org, March 13, 2019). Miners can censor transactions by not including them in their blocks. However, with the current hashrate distribution no single mining pool has sufficient power to keep a transaction out of the blockchain consistently, which would require a large majority of hash power (bitcoinity.org, (March 13, 2018). I conclusion, Bitcoin is highly censorship resistant in its current form, with the potential to improve further with respect to privacy with the introduction of new transactions methods and the lightning network.

4.2.1.3. Store of Value

A good store of value is something which over time maintains, or increases, its purchasing power, that is its value in terms of other goods. This is also known as zero or negative inflation. While accurately predicting inflation as the purchasing power of currencies has proven extremely difficult we can theoretically look at the phenomenon as a supply/demand question. In the context of money, supply is the current amount of the money in circulation times its velocity (an indirect function of inflation amongst other things) while demand mainly is represented by real economic growth, though many factors besides this one influence demand. We say that the higher the rate of growth in money supply is above the growth in money demand the higher is the inflation rate, and vice versa (Friedman, 1970). Inflation and its societal consequences, whether it can be controlled and its optimal level, remains a highly debated and polarizing topic. Sidestepping this discussion, we here stick to the single individuals and what

inflation and a good store of value mean to them, isolated from the broad society. It is indeed these individuals and their needs which according to disruption theory determines the success of a product or technology such as Bitcoin. From this perspective, it seems hard to argue against the appealing nature of a good store of value. It sustains the value of savings and work done in the past and allows the individual to buy more goods and services in the future. A good store of value doesn't undermine a person's saving's purchasing power with time.

The current inflation rates of certain nation-state currencies highlight the need for money which stores value. The annual rate of inflation in Venezuela at the beginning of 2019 hit ~120.000%, making goods priced in Venezuelan bolívar hundreds of times more expensive over a year (Hanke, January 30, 2019). Practically, this makes 1.000 bolivar today worth less than 1 bolivares in a single year. This correlates with a huge increase in M1 monetary supply (coins and notes in circulation and other assets that are easily convertible into cash) in recent years (CEIC, March 21, 2019a). While this is the only current example of such extreme hyperinflation, the inhabitants of several other nations states are experiencing inflation, highly undermining to their savings. Countries such as Zimbabwe, Sudan, and North Korea are all experiencing annual inflation around or above 50%, while large economies and ~80 million populations such as Iran and Turkey have inflation rates of 40% and 20% respectively (IMF, March 21, 2019). Since 2015 Iran have seen an increase in M1 monetary supply of ~50% while Turkey saw an increase of ~60% in the same period (CEIC, March 21, 2019b; CEIC, March 21, 2019c).

While not directly representing the above mentioned inflationary fiat currencies, Visa does indeed build on top of these why it seems relevant to start here. As previously accounted for in our examination of fiat money and banking, the state has the power when it comes to adjusting monetary circulation, through the reserve requirements of banks and interest rates. The centralization of this power creates the potential for abuse through excessive increases in money circulation, the most directly adjustable variable in the inflation supply-demand model. There are several reasons why a nation-state might want to increase the circulation of money including financing of its own overconsumption, indirect wealth taxes as well as attempted acceleration in economic growth through micromanagement of the inflation rate. However, the reasons to increase monetary supply are less important than the centralized ability and willingness to actually do so. The above correlations do not prove that rapidly increasing money circulation is the main

reason behind high inflation rates. As mentioned we intentionally avoid this discussion. They do however back up the simple premise that, all else equal, state-induced increases in monetary circulation increases inflation and undermines the utility of fiat currencies as a store of value. A majority of developed nation-states currently have low and stable inflation rates and currencies that provide good stores of value. Meanwhile, in the currencies themselves, there is nothing that inherently prevents this from changing at any given time. Given this uncertainty, we cannot label them as good stores of value in the long term.

As with many other aspects of Bitcoin, its monetary growth and thereby inflation is dependent on the decentralization of the network and any future protocol changes. However, given the already accounted for high node count, distributed mining power and the ability for everyone to run the software they wish (by which money supply is determined), the current plan for monetary growth seems dependable. Bitcoin has a fixed monetary cap of 21 million bitcoins and an issuance schedule which follows a logarithmic curve. Specifically, the issuance rate per mined block, current 12,5 bitcoins, is halved every 210.000 blocks, or approximately four years, with the last coin being issued some time in 2140. Hence the current annual inflation rate of 3,7% (52.560 blocks * 12,5 bitcoins / 17.600.000 bitcoins) will decrease over time at a predictable rate towards 0. Bitcoins low inflation rate and predictable issuance schedule make it a good store of value in theory. Contrary to fiat based currencies, no single entity holds the power to increase supply and inflation. This doesn't mean that bitcoin, given some time period, will increase in price. As far as goods continue to be priced and generally purchased in fiat currencies, and bitcoin remains more speculative in its use than utilitarian, the price of bitcoin will remain volatile and depend on many other factors as well. However, with growth in adoption, the price should become increasingly resistant to such factors and grow to become a good store of value.

4.2.1.4. Fraud

Fraud and consumer protection is a significant part of a payment system. The risk which individuals associated with storing value and transacting through a system influence their propensity to use the system. A higher risk of fraud increases the mental transaction costs of using the system, thereby raising the barrier to customer acceptance and use (Szabo, 1999). Payment fraud is a significant and growing problem in developed countries such as the US. In 2017, 78% of organizations experienced attempted

and/or actual payments fraud while over 250.000 instances were reported by consumers in 2018 (Federal Trade Commission, February, 2019; J.P.Morgan, 2018).

Of these 250.000 instances, credit card fraud was the second largest category with 50.000 reports with a value of $131 million, only succumbed by wire transfers (Federal Trade Commission, February, 2019). Given what we already know about Visa, this vulnerability towards fraudulent transactions is not surprising. One of the central designs of the system in terms of consensus, state and privacy are indeed to favor quick and frictionless transactions while using state institutions to handle edge cases such as fraud through chargebacks. This creates a solid level of consumer protection around fraudulent transactions leading to loss of funds. The funds and accounts of consumers are easily recoverable given the systems centralized nature. Meanwhile, while being just as vulnerable to theft of funds, identities are far less recoverable given their transferable nature. In 2018, credit card fraud accounted for a large majority of identity fraud claims (Federal Trade Commission, February, 2019). Such claims can only reverse financial transactions but not stop the metadata about the transactions and transactions from propagating over the internet. Originating in the privacy design of the system, which forces consumers to reveal their identity at each purchase, this is a large drawdown in the consumer fraud protection of Visa.

Bitcoin is somewhat contrary to Visa in the context of fraud. Whereas Visa has strong consumer protection around the recovering of lost funds, Bitcoin has no remedies towards lost or stolen keys. If you lose your private keys or someone else gets there hands on them there is simply no one to contact, you are on your own. This is both a feature and a bug of the system. It makes fraudulent transactions practically impossible in the first place (in terms of the cryptography there is no such thing) and it makes funds unrecoverable if lost or stolen. This is of course because, given Bitcoins pseudonymous design, the only evidence existing for coin ownership are the private keys. The pseudonymous design, meanwhile, makes Bitcoin far more resistant to identity fraud than Visa. Even if funds are stolen and spend there is to link to the identity of the previous owner. In other words, there is no identity to steal.

4.2.2. Target Market

At this point, we have thoroughly accounted for and compared the technical aspects of Bitcoin and Visa as well as investigated how technical differences translate into the competing dimensions of electronic

payment systems. In this section, we proceed to distill our current knowledge on the performance and features of the systems into specific target markets, defined by the needs of consumers. The defining feature of disruption theory is how it evaluates technology, not as existing in a vacuum, but as relative to market demand. In their focus on the most profitable customers, incumbent disruption prone technologies tend to overshoot the needs of the average consumer in the main competing dimension, leaving parts of the market behind in the process as product complexity and prices increase. New disrupting technologies initially tend to perform much worse in the main competing dimension, instead serving the market segments neglected by the incumbent given the unique strengths and features inherent to their alternative technological design. Hence, this part of the analysis is of utmost importance in our pursuit of determining the disruptive potential of Bitcoin towards credit card networks.

4.2.2.1. Main Dimension

Given its intuitive meaning along with the importance which Visa ascribes to the measure, transaction throughput is arguably the main competing dimension of global electronic payment systems. In line with what is prescribed by disruption theory, Visa, the dominant incumbent, excels in this dimension to the point where it exceeds actual demand by far. Its current capacity of 65.000 transactions per second is more than 4 times peak demand in 2017 and more than 16 times average demand in 2018 (~124.3 billion transactions a year) (Visa, 2018). The reason behind the drive towards transaction throughput is linked with the company's core customers and their demands. Visa makes money by taking a percentage cut of every transaction, and the company is therefore incentivized to maximize the payment volume on their network. Contrarily, the actors that enable this payment volume are large sellers of consumer goods, which given the aforementioned importance of quick and frictionless payments to decrease consumer bounce rates, demand high and seamless transaction throughput. Visa acknowledges this reliance on core customers:

> *Because a significant portion of our operating revenues is concentrated among our largest clients, the loss of business from any one of these larger clients could harm our business, results of operations, and financial condition.* (Visa, 2018, p. 26)

This mechanism incentivizes Visa to design their network around the needs of large clients which drive payment volume, as opposed to expanding the reach and usability of the into smaller niche markets. Meanwhile, the base Bitcoin protocol with its ~7 transactions per second has little ability to compete with Visa for its core customers but as we know this is not uncommon for disruptive technologies. On the contrary, Visa's focus on transaction capacity and large clients creates an environment for potential neglect and alienation of market segments and consumers that don't have this need.

4.2.2.2. Alternative Dimensions

While the incumbent(s) mainly compete on one dimension for their core customers, there are other dimensions which have equal importance to other segments of the market. Censorship resistance is a good example of an alternate dimension. There is a good chance that the average consumer and business in developed countries give little thought to the censorship resistance of their payment system, given that so few actually have been in a relevant situation. Indeed, there is little proof that censorship resistance is a dimension which Visa attempts to compete on and the network remains inherently censorship prone. On the contrary, through the aforementioned transaction analysis, the network is enabling such censorship making censoring easier and more scalable. Bitcoin, however, is highly competitive in the alternative dimension of censorship resistance. All else equal, this makes it valuable to market segments which do demand this feature, including political groups, darknet markets, and citizens of authoritarian states. Indeed, censored actors like Gab and Wikileaks have already embraced Bitcoin while there is a growing interest for the technology in China which already accounts for a large part of global Bitcoin volume (McCormack, February 10, 2019; Spaven, June 25, 2013; Bambrough, December 31, 2018; Coin Dance. March 23, 2019).

A similar picture emerges around the alternative dimension of store of value. Visa's core customers, consumers, and businesses, mainly operate in developed markets where inflation generally is low and fiat currencies preservers purchasing power relatively well. These countries tend to have developed and well-functioning state institutions, the foundation of Visa and fiat currencies, and as a result a relatively more stable and responsible monetary policy. Visa does not provide a good store of value for individuals in countries where the opposite is true, and the local currency is highly inflationary. In these markets, Bitcoin has utility given its inflation resistance. Since it doesn't rely on the presence of institutions is

works agnostically no matter where in the world you are, and what type of state that governs the area. Bitcoin has in fact seen disproportionally large adoption in highly inflationary economies. The key example here is Venezuela which, despite being poor, technologically underdeveloped, and having crumbling institutions, has the highest Bitcoin usage per capita in the world (Ahlborg, February 8, 2019).

In the alternative dimension of fraud, Visa and Bitcoin offer different tradeoffs. Visa sacrifices privacy and security to make funds and accounts recoverable. Like in the case of store of value, this assumes strong state institutions which govern and labels fraudulent transactions. Visa, therefore, offers excellent consumer protection in such cases, to consumers and businesses in countries with advanced institutions, while neglecting those missing said institutions. Bitcoin, meanwhile, sacrifices fund recoverability for the sake of security towards fraudulent transactions, which no matter institutions secures funds far better than in the fiat Visa system. This makes Bitcoin a far more secure option in developed countries, but rather inferior in the eyes of consumers in developed nations. Additionally, as a side effect of said non-recoverability of funds, Bitcoin is far more secure towards identity fraud than Visa. This makes Bitcoin valuable to individuals in developing countries which values privacy highly, a group which arguably have been growing recently given the increasing number of customer data breaches from large business and other institutions.

Concluding on our assessment of the various target markets of Bitcoin and Visa, the latter continues to be the more valuable to the average Visa consumer in a developed western country in North America or the EU. This is true given transaction and protection against loss of funds. Meanwhile, given the focus of Visa on their core customers (consumers and large businesses), there is substantial room for Bitcoin to capture the parts of the global electronic payment market where transaction throughput and fraud are of less relative importance. These niche segments, whether criminalized and censored, subject to high inflation or fear of identity theft are and will find value in Bitcoin's design and features. Returning to disruption theory, such unaddressed niche markets is exactly what a disruptive technology needs to take root. However, the technology is not disruptive until it actually manages to move upmarket through improvements in the main dimension of competition. This is what we address next; can Bitcoin grow its transaction throughput, while still preserving its advantages in censorship resistance, store of value and identity theft?

4.2.3. Innovation Trajectory

Having established the current target markets of Visa and Bitcoin, given their strengths and weaknesses, we must keep in mind the quote from earlier; "disruption is a process, not an event." (Christensen, 2006, p. 46) In this section, we thus move on to examine the innovation trajectories of Visa and Bitcoin on the main dimension of competition, transaction throughput.

Following the methodology used by Christensen in his study of the hard drive industry, we have collected data on the transaction capacity of Visa over time, from official press releases and annual reports. A linear regression has been fitted to these data points to clarify the trend over time. The same has been done with demand for transactions, both for average demand, calculated based on annual transaction volume, and peak demand, the highest number of transactions per second during the year. The two trend lines representing demand indicates a realistic range for how many transactions actually are demanded.

The data, presented in Figure 4.1 indicates solid and rather stable growth in the transaction capacity of the Visa network since 2002. Specifically, there has on average been an increase in transactions per second of ~4000 a year. Compared to transaction throughput, the actual transaction demand has grown at a far slower rate. Here we observe an annual increase of ~142 in average transactions per second and ~850 transactions per second in peak demand. As a result, after being nearly equal in 2008, the transaction capacity was more than 4 times the peak demand in 2018 and more than 16 times the average demand. The data then support the conclusions from the previous section; Visa has focused on transaction capacity which it has grown to a level far above what is needed by most customers.

Evident from Figure 4.1 is Bitcoin's disappearingly small, modestly growing transaction throughput. The two data points represent the ~3 transactions per second limit, inherent to the initial 1MB block size limit, and the recent increase to ~7 transactions per second as a product of Segwit which practically doubled to efficient block size to ~2MB. Besides this event, the metric has been unchanged over the years. Increasing transaction throughput of the base protocol is more often than not a matter of changing the block size which involves trading off valuable decentralization and security. There are, however, some innovations which in the future have the potential to increase throughput, not by increasing the block

size but by decreasing the size of transactions. The most likely of these in the near future is the introduction of Schnoor which decreases the size of signatures, with time bringing a block capacity increase of up to 40% (Wirdum, April 4, 2016).

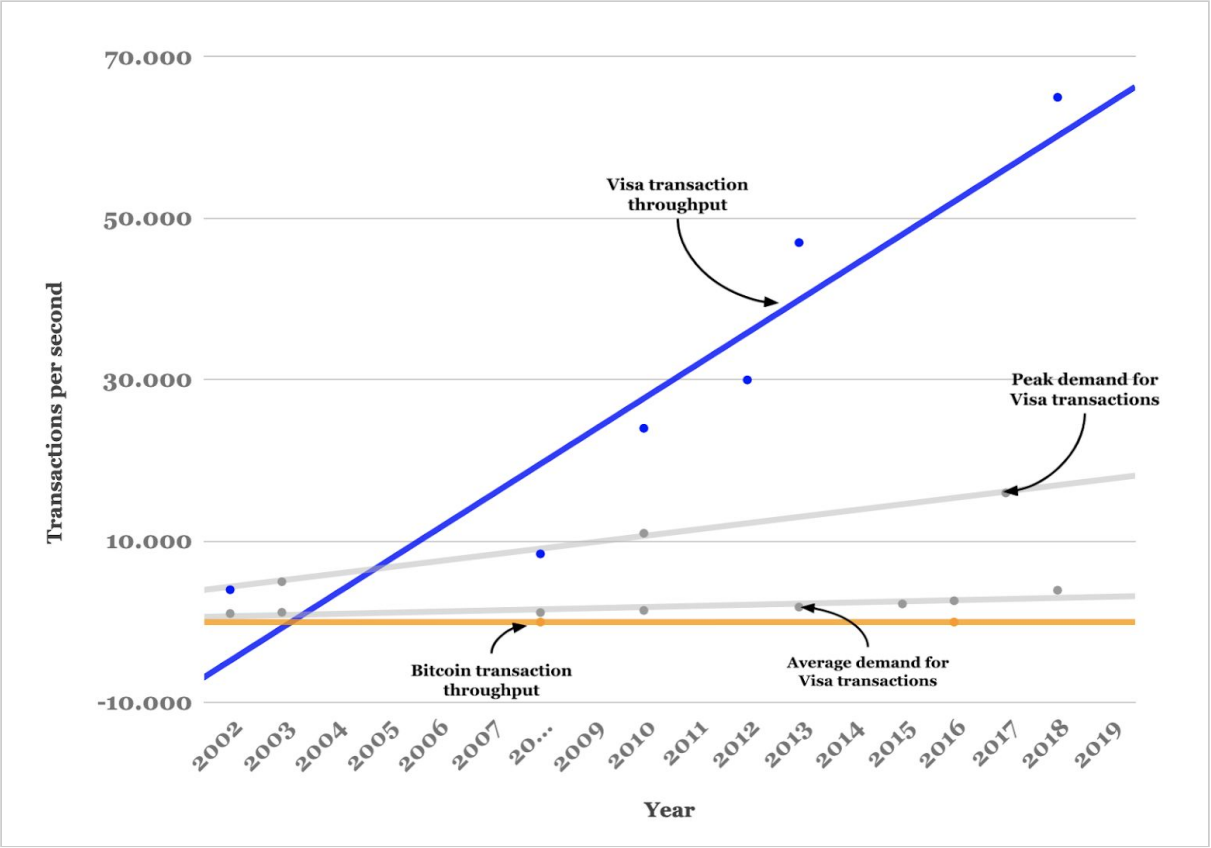

Figure 4.1: Intersecting trajectories of supplied transaction throughput of Visa and Bitcoin vs. market demand. (Visa, 2018; Trillo, October 10, 2013; Visa, April 9, 2019; Visa, 2016; Visa, 2013; Visa, 2010; Visa, 2008; Visa November 16, 2009; Trillo, January 12, 2011; Finextra, April 5, 2004; Forbes, September 16, 2002)

While not insignificant, a 40% increase in capacity does little towards scaling Bitcoin to thousands of transactions per second. With its initial niche markets, does Bitcoin have the potential to move upmarket by improving its transactions per second sufficiently, while maintaining its other strengths? Here we again turn to the lightning network, which increasingly is seen as the solution to scaling issues of the base protocol. Plotting lightning network throughput in the above figure is very difficult given that it is unicast P2P and there is no single bottleneck, like Visa net, which can be analyzed. The measure does not

really apply to lightning and theoretically, there is no transaction limit. Meanwhile, there are shortcomings in the current lightning protocol which practically makes some payments and adoption more difficult, which indirectly limits throughput.

Lightning throughput is a function of the number of nodes and channels and the capacity of channels. Atomic Multi-Path Payments (AMP) enables easier routing of larger payments, effectively increasing capacity. The protocol improvement allows payments to be cut up into smaller pieces which can travel through different routes to reach to the receiver. This is similar to how BitTorrent and TCP works. Besides allowing for larger payments with less capacity, it will also make the lightning network quicker; a single node can be slow, but in general, the network as a whole won't (Osuntokun, February 6, 2018). Submarine Swaps or loop is a part of a solution to the previously discussed feature of lightning's design; channels require inbound receiving capacity in order to receive funds. Loop Out allows individuals to move some funds from a channel to the Bitcoin blockchain in order to increase inbound capacity instead of closing the channel and opening a new one, highly useful for merchants which primarily receives payments. Loop In is valuable to consumers as it allows direct transfer from the blockchain into a lightning channel, increasing outbound spending capacity, again without closing and opening channels (Bosworth & Vu, March 20, 2019). Loop makes capacity increases easier and arguably creates a better user experience, important to increase the number of nodes. Noteworthy, these innovations are not merely smoke and mirrors but actual existing techniques and technologies build, tested and ready to be implemented in the years to come.

4.2.4. Is Bitcoin Disruptive?

Since January 2018, the lightning network has grown at an increasing rate from nothing to ~7.500 active nodes with ~4.000 active channels and an overall capacity of ~1000 bitcoins, or ~$4.200.000 (1ML, March 11, 2019b). With the selected innovations and many more in the works, it is increasingly looking like a realistic scaling solution for Bitcoin's transaction throughput (Wirdum, May 2, 2018). Significantly, this scaling maintains Bitcoin's censorship resistance, store of value and independence from institutions to secure funds. As we have seen lightning is trustless, and Bitcoin full nodes preserve their ability to autonomously validate transactions and blocks. The user experience around lightning will continue to improve as the ecosystem matures and make it a more viable payment system to more and more people. With the transaction throughput constraint removed from Bitcoin, it will continue to have a strong

foothold in its current niche markets, while slowly but surely gaining adoption with mainstream consumers in developed countries, given its additional features, absent from Visa and its peers.

Bitcoin is nothing like a sustaining technology. Recalling our previous definition, Bitcoin's use case is no to do the same things as Visa and other credit card networks. Its ability to dwarf the transaction throughput of Visa, in the long run, exists despite its main features, not because of them. Accordingly, Bitcoin diverges drastically from the business model of credit card networks, which effectively derives from a monopoly on transaction routing and transfer fees, highly incompatible with Bitcoin's decentralized and open market for transaction fees and lightning routing. Hence, unlike sustaining technologies Bitcoin fits very poorly into credit card networks' value chains, and there is little chance of them adopting it in the foreseeable future, if this is even possible.

If not sustaining, must Bitcoin be disruptive? Indeed, Bitcoin seems to tick all the boxes when it comes to being a disruptive technology. The incumbent technology is increasingly focusing on its core customers at the high end of the market, and how to sustainably innovate to satisfy them while neglecting the low-end of the market. Bitcoin, as the new entrant, has highly differentiated features in a number of areas which makes it valuable to certain niche groups of said low-end, while initially lacking performance and appeal in the dimension which the majority of the market is focused on. Meanwhile, Bitcoin's innovation trajectory, embodied mainly in the lightning network, mitigates its disadvantage in this main dimension over time, while maintaining its differentiating features. At some point, Bitcoin's performance becomes good enough for the majority of the market, and the market/incumbent has been disrupted. Of course, this is easier said than done. While we have shown that Bitcoin so far has ticked the boxes of a disrupting technology, the second phase of actually moving upmarket still lies ahead, undoubtedly full of challenges and setbacks. But if the historical data which disruption theory builds on is generalizable, Bitcoin has a very real chance of actually overcoming these barriers and with time become the dominant global payment system.

4.3. Counterarguments and Limitations

Where here turn a concise critique of our findings. While our reliance on a deductive line of reasoning brings significant upsides to the testability of our findings it is also associated with a certain type of risk as

our findings. Our main findings build on a number of assumptions previously defined and argued for in this work. If but one of these proves wrong it could falsify the entire following chain of arguments leading to the main findings. Moreover, though this generally is a well known and accepted weakness, and strength, of deduction this risk is arguably somewhat more significant in this work, compared to pure natural science where deduction so often is used. While we attempt to approach our research area from a direction of pure computer science we cannot refute that disruption theory involves an inseparable element of social science/business theory which we address in our analysis of the target market. Additionally, the systems defined and analyzed are inherently political, Bitcoin given the power games around consensus and decentralization and Visa as a product of its status as a business and the system's reliance on state-controlled fiat currency. As opposed to pure natural sciences, this element of social science from the theory and technologies brings an extra level of complexity into our investigation. It involves human decision making, not just ones and zeroes, inherently difficult to predict. While we largely avoid addressing this complexity, for the sake of simplicity and brevity, we cannot deny that it exists, significantly increases the risk that one of our assumptions, and as a product of this our main findings, could be wrong. We have tried to mitigate such added risk through the use of various data points, for instance with respect to market demand and supply but ultimately we cannot remove it entirely.

A prime example of said complexity and the risk that follows from it is the future security budget of Bitcoin, especially in the context of a growing lightning network. Bitcoins security budget refers to the amount of resources it would take to attack Bitcoin through a 51% attack, thus invalidating the assumption of persistent state and secure funds. It is a function of the current issuance of new bitcoins and transaction fees. The higher the issuance and transaction fees, the more it would cost for a miner to forego this income and mine on a separate chain. Though most miners are incentivized economically to not forego such profit, there might be special cases where, for instance, a state might feel threatened from Bitcoin and attempt to undertake an attack regardless of economic costs. All else equal, the disincentive to attack decreases with the security budget. As the bitcoin issuance rate falls over time the, fees have to increase proportionally to keep the level of security stable (Sztorc, February 14, 2019). Without any off chain scaling solutions fees would indeed be expected to increase as block space is limited. However, the introduction of the lightning network makes block space far less valuable. With growing adoption of lightning in the future, there is a plausible scenario where on chain transactions don't increase and fees

fall, diminishing the security budget. To be clear, this is not certain. An equally plausible scenario would be that the added activity in the lightning network increases on chain transactions as well, through commitment transactions. But if it is, it could seriously challenge the assumption made in this paper, namely that lightning scales transactions while keeping Bitcoin censorship resistant and safe from loss of funds. This, in turn, would partly falsify the finding that Bitcoin indeed is disruptive.

5. Are All Cryptocurrencies Disruptive?

Our findings suggest that Bitcoin indeed is disruptive to Visa and other credit card companies. Bitcoin's design contains unique features, in particular, censorship resistance, store of value, fraud protection, and privacy. What makes it truly disruptive is its ability to scale transaction throughput over time, while preserving said features which differentiate it from the incumbents. This makes Bitcoin attractive to a number of niche markets today, and the mainstream market in the future.

Do our findings imply that other cryptocurrencies are just as disruptive as Bitcoin, or maybe even more? While unique in many ways, Bitcoin is not the only cryptocurrency, nor alone in aiming to be a global electronic payments system and take on the credit card industry. Since Bitcoin's inception, countless crypto projects have built on Bitcoin's design and by changing parts such as the consensus algorithm and block constraints have tried to create a better Bitcoin. In this section, we discuss the question posed above. Given the scope of this work, we constrain our discussion to two cryptocurrencies besides Bitcoin, both aiming to disrupt global POS payments yet different in design from each other and from Bitcoin. However, the lessons from these two examples provide more general insight as to what makes a cryptocurrency disruptive.

5.1. Bitcoin Cash

Bitcoin Cash was forked from Bitcoin following the highly contentious block size limit debate in 2017. Given rising transaction fees at the time, a part of the Bitcoin community proposed an increase in the block size limit and unable to reach broad consensus Bitcoin Cash was forked from Bitcoin in August 2017 with an 8MB limit, later increased to 32MB. Being a recent fork of Bitcoin with its conservative development schedule, Bitcoin Cash highly resembles Bitcoin with the exception of the block size limit as well as segregated witness transactions and other minor fixes (n.a., April 5, 2019). So how is Bitcoin Cash interesting for our inquiry into the disruptiveness of cryptocurrencies? Because it allows us to investigate the impact of scaling philosophy and decentralization, all else being equal. As we will see, only by changing the block size limit, Bitcoin Cash makes fundamentally different choices from Bitcoin in these areas.

By increasing the block size limit to handle more transactions, Bitcoin Cash uses on chain scaling, as opposed to Bitcoin's off chain scaling with lightning. In this regard, on chain scaling is highly effective and responsive. Increasing the block size doesn't require the introduction of new technology and transaction throughput as a function of the block size limit is highly predictable. In the context of Visa, Bitcoin Cash is therefore theoretically able to compete on transactions per second, given a large enough block size increase. Its scaling solution is more straightforward than the lightning network. However, it does not preserve the unique design features, which is shares with Bitcoin, which makes it disruptive. A larger block size means more data which full nodes have to store in order to validate blocks and transactions. This data already is a considerable hurdle to running a Bitcoin full node and requires considerable storage and internet bandwidth to download and keep up-to-date (Wirdum, February 15, 2019).

Having chosen on chain scaling through the block size limit, Bitcoin Cash will be forced to significantly increase the current limit in the future, if adoption increases. Even with a 32MB limit amounting to ~100 transactions per second (based on 3 transactions per second with 1MB for Bitcoin), the network is far below the average level of Visa. Using the 2018 data of an average of ~4000 transactions per second, Bitcoin Cash would need a blocksize ~1,3GB (4000 transactions / 100 transactions * 32MB) causing the blockchain to grow by ~68TB every year (52.560 blocks * 1,3GB). Large requirements to run full nodes decreases the number of individuals that have the resources and willingness to do so. Hence, a larger block size limit decreases the decentralization of the network, no matter what innovation in storage and bandwidth costs one might assume (Houy, 2014; Wirdum, February 15, 2019).

Decentralization is the cornerstone of the features which make Bitcoin unique and in turn disruptive. Less decentralization means less effort to censor addresses, less effort to change the monetary policy and undermine the store of value and less effort to steal funds through chain through 51% attacks. And without these unique features, a technology has no niche markets to appeal to and grow from at first, and no way to outcompete the incumbent overtime. Thus, from our findings on Bitcoin, we cannot conclude that Bitcoin Cash is similarly disruptive to credit card networks. On the contrary, it fits far better into the

description of a sustaining technology as it tries to improve on the current main dimension of competition, transaction throughput.

5.2. Ripple

While a part of the original wave of cryptocurrencies which follows after Bitcoin's inception, Ripple has developed more or less independently from Bitcoin since its beginning in 2012. Indeed, Ripple is very different in key areas from Bitcoin. First, while the source code is open to view for all, Ripple Labs, the company behind Ripple, maintains sole control over the code base. Second, the Ripple consensus model is far more centralized than that of Bitcoin. Specifically, only selected nodes on the network, determined by Ripple Labs, have special validator status, meaning that they have the power to accept, reject and append transactions to the ledger. A large part of these validator nodes is run by Ripple Labs itself, while other trusted parties are planned to be banks and other organizations that aggregate payments. Validators aggregate broadcasted payments into proposals (similar to blocks) which they transmit to other validators. Consensus on the ledger is reached between validators, without any influence from other types of nodes, when more than 50% of validators accept the proposal (Armknecht et al., 2015).

Ripple, like Bitcoin Cash, scales on the base layer. Meanwhile, its inherent centralization and lack of block time (proposals can be added to the ledger as quick as they can be agreed upon) mean that its transaction throughput is as high as that of Visa. Each ledger proposal contains as many transactions as needed and reaches consensus within 4 seconds. Though labeled as a cryptocurrency given its use of public key cryptography, a public ledger and a P2P network, comparing Ripple to Bitcoin is like comparing apples to oranges. The network arguably resembles Visa to a higher degree, which likewise use a large distributed yet trusted net of servers and databases, with some consensus algorithm, to rout payments and coordinate balances. Though Ripple Labs claims the RIpple network to be decentralized, its ability to change the code base any time and control transactions speak against this (Armknecht et al., 2015).

The centralized power of Ripple Labs undermines the areas that separate Bitcoin from a sustaining technology. Any censorship resistance which comes from pseudonymous transactions is compromised when a single entity freely can determine what addresses are allowed to transact. Ripples inflation is likewise in the hands of Ripple Labs through its control over the code, making its store of value highly

questionable. And the security of funds is not much safer than with electronic fiat currency. We, therefore, cannot conclude Ripple to be a disruptive technology to credit card networks. It seeks to outcompete credit card networks in transaction throughput but fails to preserve the other unique advantages which Bitcoin possesses (Armknecht et al., 2015).

5.3. Why Bitcoin is Different

The fact that neither Bitcoin Cash nor Ripple manages to scale transaction throughput while preserving the unique features of censorship resistance, store of value and security does not by itself mean that they aren't disruptive to Visa and its peers. These systems may have unique features of their own which give them the ability to capture a niche market and from there move upmarket. Specifically, costs, which according to our investigation of disruption theory often is the differentiator of a disrupting technology, is one possible way of capturing market share from Visa. However, the problem with costs and cryptocurrencies is that there is nothing inherently effective or cost saving about a public ledger or blockchain. As we have shown, unicast is significantly faster than broadcast which is why the lightning network scales so well compared to the base Bitcoin protocol. Indeed, if the power of a P2P network is centralized like Bitcoin Cash or even more so Ripple, it can process payments faster than a decentralized competitor but by the same logic, it will never be more effective than a fully centralized non-P2P network like Visa, where no consensus has to be reached at all and transactions are stored in a relational database. Given enough centralization, a network such as Ripple might be able to reach or even surpass Visa's transaction capacity and efficiency but at that point, it will compete largely in the same areas as Visa and only have a sustaining impact.

What makes Bitcoin special from Bitcoin Cash, Ripple and other projects with similar approaches is first and foremost how its design recognizes and embraces the inherent use case of a blockchain and P2P network: easy validation of data in a trustless way. More than anything else, this, in turn, relies on a decentralized network. Bitcoin is valuable first and foremost because of this, and compromising here in order to improve other use cases is like throwing the baby out with the bathwater. Every important design decision of Bitcoin follows from this realization. Bitcoins layered architecture demonstrates this awareness perfectly. Whereas Bitcoin Cash and Ripple scale transaction throughput on chain, Bitcoin has denied

compromising on decentralization which instead has forced the emergence of a scaling solution situated on an entirely different layer of the protocol stack.

A layered structure is far more suitable than monoliths for protocols in general as they allow for different implementations at each layer without having the change the surrounding layers. This is clearly the case with the internet where the transport layer of the TCP/IP stack has both the TCP and UDP protocols, usable from the IP layer and valuable for different use cases. In Bitcoin, separation of layers keeps the base protocol stable, safe and with a strong focus on decentralization while scaling solutions can be implemented and tested on the above layers. In fact, assuming a decentralized base layer it would be difficult to innovate at all without more consensus agnostic layers as reaching consensus on significant base layer changes can be extremely difficult. Hence, one of the lessons from Bitcoin is that cryptocurrencies, like other protocols, should have a layered architecture layered not monolithic.

Besides protocol layering, Bitcoin's consensus model including proof of work is a clear prioritization of decentralization over cost and efficiency. The automatic difficulty adjustment keeps the block time at ~10 minutes and everyone is free to participate in the mining process. Moreover, all full nodes are validating nodes, and are fully able to do so given the 1MB block size limit and considerable block time. This means that transactions can't be confirmed in the same way as in a centralized system such as Ripple. In Bitcoin, chain forks are a real though improbable risk up to six blocks after a transaction. Transaction throughput is predictable, smaller and slower so that a wide set of decentralized actors in the future are able to easily download and validate new blocks and transactions.

In sum, Bitcoin is different, not because it is better at everything but because it is self-aware enough to know what makes it special and in turn ensure that any changes preserve its uniqueness. Easy and trustless validation through decentralization is a simple but powerful concept which enables the features so unique to Bitcoin. They can be leveraged through more efficiency but without them, Bitcoin is far less interesting. To achieve disruption of the credit card industry Bitcoin will one day have to do both at the same time, increase efficiency while staying decentralized. While the lightning network looks promising in this regard there is a real chance that it may not be the answer to Bitcoins scaling dilemma. In this case, other solutions might emerge out of a need for transaction throughput on different layers of the Bitcoin

protocol stack instead. In the context of disruption how these solutions might look is unpredictable and frankly uninteresting. Disruption is a dynamic process, not an event, and the path for Bitcoin toward becoming the dominant global payment system which protects users from inflation, censorship, and theft may involve several scaling solutions yet to be discovered.

Conclusion

Despite the novelty and complexity of Bitcoin and its foundational technologies, unsupported assumptions around its value and use-cases are often taken for granted. With this work, we set out to test one of the largest and most impactful of these assumptions, namely that of Bitcoin's inability to scale, its inferiority to credit card networks and consequently and superiority of other blockchain-based cryptocurrencies. In this context, disruption theory was found to be a fruitful analytical framework as it allows for comparisons of radically different technologies over time while taking market demand into consideration. Hence we formulated our research question as the following: *Is Bitcoin a disruptive or sustaining technology to credit card networks?*

In Chapter 3, we thoroughly accounted for the technical foundations of a credit card network (Visa) and Bitcoin. Fiat money and banking were presented to demonstrate how the foundation of credit card networks functions, including how states control fiat money, give them value, issue them and how banks, the endpoints of credit card transactions, store money via centralized ledgers. We also followed a Visa transaction from beginning to end, starting from authorization and ending with clearing and settling. Finally, the edge cases around credit card transactions were examined such as transaction chargebacks and confiscation of funds, ultimately made possible by the centralization of banking ledgers and the state power of fiat money. We then moved on to Bitcoin, initially examining three of its underlying core technologies: the P2P network of nodes with different functionality and roles, the blockchain data structure which stores transactions in blocks and the proof of work consensus algorithm, arguably the most important pillar of Bitcoin, which secures transactions and funds and prevents forks. We then examined a single Bitcoin transaction, different transaction scripts, and how these all are some combination of transaction inputs (UTXOs) and outputs, before looking at the lightning network, the most prominent example of a number of new emerging layer to scaling solutions.

In the first part of Chapter 4, we began the analysis by establishing a number of premises around the technical similarities and differences of credit card networks and Bitcoin, with respect to specific concepts from computer science. First, in the context of consensus, we found that Bitcoin distributes all of the power to the nodes of the network whereas Visa centralizes the power in the middle of the network with

Visa itself. Second, while Bitcoin keeps durable and practically immutable state within its distributed blockchain ledger, Visa transactions are made durable once a day, during clearing and settlement, and can never be considered immutable. Third, both Bitcoin and Visa use eventual consistency and BASE transactions to achieve overall distributed shared state while subsections, the lightning network, and Visa transaction authentication, use ACID transactions. Fourth, with pseudonymous addresses and Confidential Transactions and the lightning network Bitcoin has the potential to achieve close to complete privacy, far removed from Visa where all transactions are tied to an identity which is known by both banks and merchants. Fifth, Bitcoin uses a layered architecture with the lightning network to the Bitcoin base layer what the TCP/IP layer is to the internet, while Visa operates at a single layer, arguably at the same level as the lightning network.

In the second part of Chapter 4, we used the premises from part one to identify the competitive strengths and weaknesses with respect to features, performance and market demand, an important part of Disruption theory. First, in terms of transactions per second, Visa's centralized design allows for highly scalable transaction throughput with demand, whereas on chain Bitcoin transactions are highly restricted and difficult to increase significantly without incurring large negative decentralization tradeoffs. Meanwhile, the lightning network and other layer two solutions have the potential to mitigate this constraint. Second, Bitcoin's decentralization and privacy make it highly censorship resistant and reliable for customers, compared to Visa, where censorship frequently and easily is practiced by Visa and state authorities. Third, Bitcoin represents a safe store of value for users, with predictable decreasing inflation, given its consensus model as well as decentralized power and state while Visa builds on top of highly centralized and inflation-prone fiat currencies. Fourth, Visa offers good fraud mitigation, though limited preventive protection, for users in countries with access to public consumer protection institutions and laws while Bitcoin has built-in security against theft beneficial to users without such state consumer protection and/or towards identity theft.

We then sued these insights around market use cases and consumer demand to identify specific target markets the needs of which Bitcoin and Visa respectively currently fulfill. We considered these markets in terms of the established competitive dimensions and, in line with disruption theory, we identified the dimension currently dominant in the credit card industry, transaction throughput, before focussing on

the remaining, alternative competitive areas. The findings here suggest that Visa, as expected from an incumbent, has a strong focus on improving on on the main competitive dimension and hereby serving their current core customers, large businesses and consumers in North America and the EU. This, therefore, leaves room for Bitcoin to compete where this dimension is of less important relative to the alternate dimensions. This is indeed what Bitcoin does, in countries and market segments with a large need for a store of value given high inflation, consumer protection given lack of institutions and censorship resistance due to political prosecution.

Having established the existence of an initial market for Bitcoin as an alternative to credit card networks, we assessed the potential of Bitcoin to over time disrupt the incumbents. We observed the innovation trajectories of Visa and Bitcoin in terms of transaction throughput, finding that the former over time has increased supply far above actual demand. We also found that while Bitcoin currently is vastly below the average demand for transactions, the lightning network is a living example of an innovation which scales Bitcoin transactions to meet demand while preserving the unique features which makes Bitcoin special from Visa in the first place. Hence, we found that Bitcoin checks all the boxes of a disruptive technology, including an incumbent industry, focused on serving core customers, a niche market for Bitcoin to establish itself and grow from and a technological development which allows Bitcoin to scale and increasingly capture market share from credit card networks.

The limitations of our findings where thereafter addressed. These generally relate to the deductive approach of our reasoning, which works be repeatedly establishing clear premises and building on top of these. Hence, should one of the previous premises prove to be false, it invalidates the findings derived from it. In this context, we discussed Bitcoin's security budget and how the growth of the lightning network could invalidate the assumption made around Bitcoin's decentralization and immutability.

Finally, in Chapter 5 we discussed the relatability of the findings around Bitcoin to other cryptocurrencies, with special attention paid to Bitcoin Cash and Ripple. Both are cryptocurrencies which have taken different approaches from Bitcoin, one by scaling on chain and another through a centralized consensus algorithm. We argued that by choosing such tradeoffs with the scope of easy scaling these systems have made themselves irrelevant. They have sacrificed what made Bitcoin special in the first

place, easy and trustless validation through decentralization, to play the game of transactions per second and are thereby at best sustaining innovations to credit card networks. What makes Bitcoin special is the recognition in its design and innovation philosophy that its founding principles of decentralization and trustlessness must be preserved whatever it takes. This, in turn, is what has forced the Bitcoin community the scale with protocol on other layers, such as with the lightning network, allowing for conservative preservation of its strengths in the base protocol coupled with creative experimentation to mitigate the inherent downsides of Bitcoin and all blockchain based technologies that originates from it.

Literature

1ML. (March 11, 2019a). Channels - Top Capacity. Retrieved from https://1ml.com/channel?order=capacity

1ML. (March 11, 2019b). Real-Time Lightning Network Statistics. Retrieved from https://1ml.com/statistics

Abernathy, W. J., & Utterback, J. M. (1978). Patterns of industrial innovation. Technology review, 80(7), 40-47.

Addo, P., Guegan, D., & Hassani, B. (2018). Credit risk analysis using machine and deep learning models. Risks, 6(2), 38.

Ahlborg, M. (February 8, 2019). Nuanced Analysis of LocalBitcoins Data Suggests Bitcoin is Working as Satoshi Intended. Retrieved from https://medium.com/@mattahlborg/nuanced-analysis-of-localbitcoins-data-suggests-bitcoin-is-working-as-satoshi-intended-d8b04d3ac7b2

An, D. (Febraury, 2018). Find out how you stack up to new industry benchmarks for mobile page speed. Google. Retrieved from https://www.thinkwithgoogle.com/marketing-resources/data-measurement/mobile-page-speed-new-industry-benchmarks/

Antonopoulos, A. M. (2017a). Introduction. In Mastering Bitcoin: Programming the open blockchain (pp. 1-15). O'Reilly Media, Inc.

Antonopoulos, A. M. (2017b). The Blockchain. In Mastering Bitcoin: Programming the open blockchain (pp. 195-213). O'Reilly Media, Inc.

Antonopoulos, A. M. (2017c). Mining and Consensus. In Mastering Bitcoin: Programming the open blockchain (pp. 213-269). O'Reilly Media, Inc.

Antonopoulos, A. M. (2017d). The Bitcoin Network. In Mastering Bitcoin: Programming the open blockchain (pp. 172-195). O'Reilly Media, Inc.

Antonopoulos, A. M. (2017e). Transactions. In Mastering Bitcoin: Programming the open blockchain (pp. 117-145). O'Reilly Media, Inc.

Antonopoulos, A. M. (2017f). Blockchain Applications. In Mastering Bitcoin: Programming the open blockchain (pp. 117-145). O'Reilly Media, Inc.

Armknecht, F., Karame, G. O., Mandal, A., Youssef, F., & Zenner, E. (2015). Ripple: Overview and outlook. In International Conference on Trust and Trustworthy Computing (pp. 163-180). Springer, Cham.

Bank of England. (2014). Rulebook Glossary. Retrieved 13 July 2018 from http://www.prarulebook.co.uk/rulebook/Glossary/Rulebook/0/B

Bambrough, B. (December 31, 2018). Bitcoin Adoption Could Be Boosted By Surging China Demand. Forbes. Retrieved from https://www.forbes.com/sites/billybambrough/2018/12/31/bitcoin-adoption-could-be-boosted-by-surging-china-demand/

Batiz-Lazo, B., & Wood, D. (2002). Information technology innovations and commercial banking: a review and appraisal from a historical perspective (No. 0211002). EconWPA.

Bitcoin Github repository. Retrieved from https://github.com/bitcoin/bitcoin

Bitcoinuptime.com. Retrieved from http://bitcoinuptime.com/

Bitnodes.earn.com. (March 13, 2019). Nodes. Retrieved from https://bitnodes.earn.com/dashboard/?days=730

bitcoinity.org. (March 13, 2018). Bitcoin network hashrate. Retrieved from https://data.bitcoinity.org/bitcoin/hashrate/2y?c=m&g=15&r=week&t=a

Blockchain. (March 11, 2019). Average Number Of Transactions Per Block. Retrieved from https://www.blockchain.com/charts/n-transactions-per-block?timespan=all&daysAverageString=7

Bosworth, A., & Vu, B. (March 20, 2019). Announcing Lightning Loop Alpha: An Easier Way to Receive on Lightning. Lightning Labs. Retrieved from https://blog.lightning.engineering/posts/2019/03/20/loop.html

Bowles, N. (December 24, 2018). Patreon Bars Anti-Feminist for Racist Speech, Inciting Revolt. New York Times. Retrieved from https://www.nytimes.com/2018/12/24/technology/patreon-hate-speech-bans.html

CEIC. (March 21, 2019a). Venezuela Money Supply M1. Retrieved from https://www.ceicdata.com/en/indicator/venezuela/money-supply-m1

CEIC. (March 21, 2019b). Iran Money Supply M1. Retrieved from https://www.ceicdata.com/en/indicator/iran/money-supply-m1

CEIC. (March 21, 2019c). Turkey Money Supply M1. Retrieved from https://www.ceicdata.com/en/indicator/turkey/money-supply-m1

Coin Dance. (March 23, 2019). LocalBitcoins Volume (China). Retrieved from https://coin.dance/volume/localbitcoins/CNY/BTC

Consoli, D. (2005). The dynamics of technological change in UK retail banking services: An evolutionary perspective. Research Policy, 34(4), 461-480.

Christensen, C. M. (1993). The rigid disk drive industry: A history of commercial and technological turbulence. Business history review, 67(4), 531-588.

Christensen, C. M. (1998). 3.1 The Evolution of Innovation. The Technology Management Handbook.

Christensen, C., Craig, T., & Hart, S. (2001). The great disruption. Foreign Affairs, 80-95.

Christensen, C. M. (2006). The ongoing process of building a theory of disruption. Journal of Product innovation management, 23(1), 39-55.

Christensen, C. (2013a). The innovator's dilemma: when new technologies cause great firms to fail. Harvard Business Review Press.

Christensen, C. (2013b). Introduction. In The innovator's dilemma: when new technologies cause great firms to fail (pp. ix-xxvii). Harvard Business Review Press.

Christensen, C. (2013c). Value Networks and the Impetus to Innovate. In The innovator's dilemma: when new technologies cause great firms to fail (pp. 29-59). Harvard Business Review Press.

Christensen, C. (2013d). How Can Great Firms Fail? Insights from the Hard Disk Drive Industry. In The innovator's dilemma: when new technologies cause great firms to fail (pp. 3-28). Harvard Business Review Press.

Christensen, C. (2013e). Disruptive Technological Change in the Mechanical Excavator Industry. In The innovator's dilemma: when new technologies cause great firms to fail (pp. 61-76). Harvard Business Review Press.

Christensen, C. M., Raynor, M. E., & McDonald, R. (2015). What is disruptive innovation. Harvard Business Review, 93(12), 44-53.

Cockburn, H. (November 22, 2018). China blacklists millions of people from booking flights as 'social credit' system introduced. Independent. Retrieved from https://www.independent.co.uk/news/world/asia/china-social-credit-system-flight-booking-blacklisted-beijing-points-a8646316.html

Croman, K., Decker, C., Eyal, I., Gencer, A. E., Juels, A., Kosba, A., ... & Song, D. (2016). On scaling decentralized blockchains. In International Conference on Financial Cryptography and Data Security (pp. 106-125). Springer, Berlin, Heidelberg.

Davidson, J. D., & Rees-Mogg, W. (1997). The sovereign individual: how to survive and thrive during the collapse of the welfare state. New York: Simon & Schuster.

Decker, C., Seidel, J., & Wattenhofer, R. (2016). Bitcoin meets strong consistency. In Proceedings of the 17th International Conference on Distributed Computing and Networking (p. 13). ACM.

Decker, C., & Wattenhofer, R. (2013). Information propagation in the bitcoin network. In Peer-to-Peer Computing (P2P), 2013 IEEE Thirteenth International Conference on (pp. 1-10). IEEE.

Federal Trade Commission. (February, 2019). Consumer Sentinel Network Data Book 2018. Retrieved from https://www.ftc.gov/system/files/documents/reports/consumer-sentinel-network-data-book-2018/consumer_sentinel_network_data_book_2018_0.pdf

Finextra. (April 5, 2004). Visa reports strong growth in 2003. Retrieved from https://www.finextra.com/pressarticle/628/visa-reports-strong-growth-in-2003

Fleder, M., Kester, M. S., & Pillai, S. (2015). Bitcoin transaction graph analysis. arXiv preprint arXiv:1502.01657.

Forbes. (September 16, 2002). Visa's Vision. Retrieved from https://www.forbes.com/forbes/2002/0916/078.html

Forbes. (May 23, 2017). Strong Growth In Volumes Boosts Visa's U.S. Credit Card Market Share To Over 52%. Retrieved from https://www.forbes.com/sites/greatspeculations/2017/05/23/strong-growth-in-volumes-boosts-visas-u-s-credit-card-market-share-to-over-52/

Foster, R. N. (1986). Working the S-curve: assessing technological threats. Research Management, 29(4), 17-20.

Garay, J., Kiayias, A., & Leonardos, N. (2015). The bitcoin backbone protocol: Analysis and applications. In Annual International Conference on the Theory and Applications of Cryptographic Techniques (pp. 281-310). Springer, Berlin, Heidelberg.

Gervais, A., Karame, G. O., Wüst, K., Glykantzis, V., Ritzdorf, H., & Capkun, S. (2016). On the security and performance of proof of work blockchains. In Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security (pp. 3-16). ACM.

Hanke, S. (January 30, 2019). Hyperinflation -- A Kaleidoscope Of Uses And Abuses. Forbes. Retrieved from https://www.forbes.com/sites/stevehanke/2019/01/30/hyperinflation-a-kaleidoscope-of-uses-and-abuses/

Harlev, M. A., Sun Yin, H., Langenheldt, K. C., Mukkamala, R., & Vatrapu, R. (2018). Breaking bad: De-anonymising entity types on the bitcoin blockchain using supervised machine learning. In Proceedings of the 51st Hawaii International Conference on System Sciences.

Helland, P. (2018). Mind your state for your state of mind. ACM Queue, 61(10), 47-54.

Helland, P. (2015). Immutability Changes Everything. ACM Queue, 13(9), 40.

Hoff, T. (May 1, 2013). Myth: Eric Brewer On Why Banks Are BASE Not ACID - Availability Is Revenue. Retrieved from highscalability.com/blog/2013/5/1/myth-eric-brewer-on-why-banks-are-base-not-acid-availability.html

Holden, M. (October 24, 2011). WikiLeaks says "blockade" threatens its existence. Reuters. Retrieved from https://www.reuters.com/article/us-britain-wikileaks/wikileaks-says-blockade-threatens-its-existence-idUSTRE79N46K20111024

Houy, N. (2014). The economics of Bitcoin transaction fees. GATE WP, 1407.

IMF. (March 21, 2019). Inflation rate, average consumer prices. Retrieved from https://www.imf.org/external/datamapper/PCPIPCH@WEO/OEMDC/ADVEC/WEOWORLD

J.P.Morgan. (2018). Payments Fraud and Control Survey. Retrieved from https://commercial.jpmorganchase.com/jpmpdf/1320745402134.pdf

Kleppmann, M. (2017a). The Trouble with Distributed Systems. In Designing data-intensive applications: The big ideas behind reliable, scalable, and maintainable systems (pp. 273-310). " O'Reilly Media, Inc.".

Kleppmann, M. (2017b). Consistency and Consensus. In Designing data-intensive applications: The big ideas behind reliable, scalable, and maintainable systems (pp. 321-389). " O'Reilly Media, Inc.".

Kleppmann, M. (2017c). Transactions. In Designing data-intensive applications: The big ideas behind reliable, scalable, and maintainable systems (pp. 221-273). " O'Reilly Media, Inc.".

Kozierok, C. M. (2005a). Chapter 1: Networking Introduction, Characteristics, and Types. In The TCP/IP guide: a comprehensive, illustrated Internet protocols reference. No Starch Press.

Kozierok, C. M. (2005b). Chapter 8: TCP/IP Protocol Suite and Architecture. In The TCP/IP guide: a comprehensive, illustrated Internet protocols reference. No Starch Press.

Kroll, J. A., Davey, I. C., & Felten, E. W. (2013). The economics of Bitcoin mining, or Bitcoin in the presence of adversaries. In Proceedings of WEIS (Vol. 2013, p. 11).

Lau, J., Lombrozo, E., & Wuille, P. (December 21, 2015). BIP 141: Segregated Witness. Retrieved from https://github.com/bitcoin/bips/blob/master/bip-0141.mediawiki

Lerner, A. P. (1947). Money as a Creature of the State. The American Economic Review, 37(2), 312-317.

Luke.dashjr.org. (March 13, 2018). Software. Retrieved from https://luke.dashjr.org/programs/bitcoin/files/charts/software.html

Maxwell, G. (n.d.). Confidential Transactions - Investigation. Elements by BlockStream. Retrieved from https://elementsproject.org/features/confidential-transactions/investigation

Maxwell. G. (2013). CoinJoin: Bitcoin Privacy for the Real World. Post on. Bitcoin Forum. Retrieved from https://bitcointalk.org/index.php?topic=27924

McLeay, M., Radia, A., & Thomas, R. (2014). Money creation in the modern economy.

Meiklejohn, S., Pomarole, M., Jordan, G., Levchenko, K., McCoy, D., Voelker, G. M., & Savage, S. (2013, October). A fistful of bitcoins: characterizing payments among men with no names. In Proceedings of the 2013 conference on Internet measurement conference (pp. 127-140). ACM.

Friedman, M. (1970). A theoretical framework for monetary analysis. journal of Political Economy, 78(2), 193-238.

McCormack, P. (February 10, 2019). Gab's Andrew Torba on Why Bitcoin Is Free Speech Money. Retrieved from https://hackernoon.com/gabs-andrew-torba-on-why-bitcoin-is-free-speech-money-dcbe15be5e43

Moses, J., & Knutsen, T. (2012). Ways of knowing: Competing methodologies in social and political research. Palgrave Macmillan.

n.a. (April 5, 2019). Bitcoin Cash. Retrieved from https://en.wikipedia.org/wiki/Bitcoin_Cash

Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash system. Retrieved from: https://bitcoin.org/bitcoin.pdf

Osuntokun, O. (February 6, 2018). AMP: Atomic Multi-Path Payments over Lightning. Retrieved from https://lists.linuxfoundation.org/pipermail/lightning-dev/2018-February/000993.html

Poon, J., & Dryja, T. (2016). The bitcoin lightning network: Scalable off-chain instant payments.

Reid, F. and Harrigan, M. An analysis of anonymity in the bitcoin system. Security and privacy in social networks, (2013), 197–223.

Robertson, A. (November 5, 2018). Gab is back online after being banned by GoDaddy, PayPal, and more. The Verge. Retrieved from https://www.theverge.com/2018/11/5/18049132/gab-social-network-online-synagogue-shooting-deplatforming-return-godaddy-paypal-stripe-ban

Samakovitis, G. (2012). UK banking experts as decision-makers: a historical view on banking technologies. Journal of Technology Research, 3, 1.

Security Research Labs. (February 13, 2018). Payment terminals allow for remote PIN capture and card cloning. Retrieved from https://srlabs.de/bites/eft-vulns/

Sood, A., & Tellis, G. J. (2005). Technological evolution and radical innovation. Journal of Marketing, 69(3), 152-168.

Spaven, E. (June 25, 2013). Bitcoiners donate to WikiLeaks to support Edward Snowden. Coindesk. Retrieved from https://www.coindesk.com/bitcoiners-rally-behind-snowden

Steensen, J. (May 24, 2018). Machine Learning in the Payments Industry. Retrieved from https://usa.visa.com/dam/VCOM/download/merchants/chargeback-management-guidelines-for-visa-merchants.pdf

Sztorc, P. (February 14, 2019). Security Budget in the Long Run. Retrieved from www.truthcoin.info/blog/security-budget/

Szabo, N. (1999). Micropayments and mental transaction costs. In 2nd Berlin Internet Economics Workshop.

Szabo, N. (2001). Trusted third parties are security holes. White Paper. Retrieved from https://nakamotoinstitute.org/trusted-third-parties/

Trillo, M. (January 12, 2011). Visa Transactions Hit Peak on Dec. 23. Retrieved from https://www.visa.com/blogarchives/us/2011/01/12/visa-transactions-hit-peak-on-dec-23/index.html

Trillo, M. (October 10, 2013). Stress Test Prepares VisaNet for the Most Wonderful Time of the Year. Retrieved from https://www.visa.com/blogarchives/us/2013/10/10/stress-test-prepares-visanet-for-the-most-wonderful-time-of-the-year/index.html

Tschorsch, F., & Scheuermann, B. (2016). Bitcoin and beyond: A technical survey on decentralized digital currencies. IEEE Communications Surveys & Tutorials, 18(3), 2084-2123.

Visa. (2008). Annual Report 2008. Retrieved from https://www.visa.ro/media/images/visa_europe_annual_report_2008-32-9905.pdf

Visa. (November 16, 2009). Visa Opens New Data Center in the U.S. Retrieved from https://usa.visa.com/about-visa/newsroom/press-releases.releaseId.8236.html

Visa. (2010). Annual Report 2010. Retrieved from http://www.annualreports.com/HostedData/AnnualReportArchive/v/NYSE_V_2010.pdf

Visa. (2013). Annual Report 2013. Retrieved from https://s1.q4cdn.com/050606653/files/doc_downloads/annual%20meeting/Visa%20Annual%20Report%202013%20final%20website.pdf

Visa. (January 21, 2016). How a Visa Transaction Works. Retrieved from http://web.archive.org/web/20160121231718/http://apps.usa.visa.com/merchants/become-a-merchant/how-a-visa-transaction-works.jsp

Visa. (2016). Annual Report 2016. Retrieved from https://s1.q4cdn.com/050606653/files/doc_financials/annual/Visa-2016-Annual-Report.pdf

Visa. (February 13, 2018). CyberSource Payments. Retrieved from https://developer.visa.com/capabilities/cybersource/docs#creating_a_cardpresent_or_emv_authorization_request

Visa. (October 10, 2018). Dispute Management Guidelines for Visa Merchants. Retrieved from https://usa.visa.com/dam/VCOM/download/merchants/chargeback-management-guidelines-for-visa-merchants.pdf

Visa. (2018). Annual Report 2018. Retrieved from https://s1.q4cdn.com/050606653/files/doc_financials/annual/2018/Visa-2018-Annual-Report-FINAL.pdf

Visa. (April 9, 2019). Visa acceptance for retailers. Retrieved from https://usa.visa.com/run-your-business/small-business-tools/retail.html

Wirdum, A. V. (April 4, 2016). The Power of Schnorr: The Signature Algorithm to Increase Bitcoin's Scale and Privacy. Bitcoin Magazine. Retrieved from https://bitcoinmagazine.com/articles/the-power-of-schnorr-the-signature-algorithm-to-increase-bitcoin-s-scale-and-privacy-1460642496/

Wirdum, A. V. (May 2, 2018). The Future of Bitcoin: What Lightning Could Look Like. Bitcoin Magazine. Retrieved from https://bitcoinmagazine.com/articles/future-bitcoin-what-lightning-could-look/

Wirdum, A. V. (November 8, 2017). NO2X: Hard Fork “Suspended” Due to Lack of Consensus. Bitcoin Magazine. Retrieved from https://bitcoinmagazine.com/articles/no2x-hard-fork-suspended-due-lack-consensus/

Wirdum, A. V. (February 15, 2019). Is It Time to Take an Initiative to Decrease Bitcoin's Block Size Seriously?. Bitcoin Magazine. Retrieved from https://bitcoinmagazine.com/articles/is-it-time-to-take-an-initiative-to-decrease-bitcoins-block-size-seriously/

Yermack, D. (2015). Is Bitcoin a real currency? An economic appraisal. In Handbook of digital currency (pp. 31-43).