Bitcoin and the Academy

Korok Ray

Korok Ray
is a professor at Texas A&M and a research director at the Mays Innovation Research Center.
korok@mays.tamu.edu

KEY FINDINGS

- Bitcoin is a radical technology that fuses multiple disciples—such as economics, computer science, and mathematics—in a novel way.
- While all the underlying ideas within Bitcoin were part of academic research, universities today are not structured for such interdisciplinary innovation.
- To create transformational innovation in the future, universities need to reorganize around interdisciplinary innovation.

ABSTRACT

Bitcoin emerged from an anonymous creator who brought together several different academic disciplines with history in academic research, such as economics, computer science, and mathematics. The genius of Bitcoin was combining these disciplines into a single economic system. The academy has resisted its adoption partly because Bitcoin is unsuited to this kind of interdisciplinary innovation. Economists and computer scientists do not understand one another's disciplines deeply enough to create something like Bitcoin. For future radical innovations, the academy must reorganize around interdisciplinary innovation by breaking down disciplinary barriers to allow mixing in novel ways.

Since its inception in October 2008, Bitcoin has reached a market capitalization of over $1 trillion. Its growth has drawn both retail and institutional investment, as the financial community now begins to see it as a legitimate store of value and an alternative to traditional assets like gold. And innovations in second layer settlements like the Lightning Network make it increasingly possible for Bitcoin to serve as a medium of exchange.

Still, Bitcoin has a precarious and somewhat checkered history with the academy. Curricula in universities are largely devoid of any treatment of Bitcoin, instead leaving the teachings to student clubs and nonprofits. Over time this may change, as Bitcoin and the entire cryptocurrency market continue to grow, attracting attention from top talent in both engineering and business. But Bitcoin's absence from university is not a problem with Bitcoin itself, but rather with the academy, with its insufficient embrace of innovation, its emphasis on backward-looking data analysis, and its excessive preoccupation with individual disciplines rather than collective knowledge. Bitcoin can serve as an inspiration for what academic research can and should be. In fact, it presents a roadmap to change higher education for the better.

SIMILARITIES

One may wonder why to even assume a relationship between Bitcoin and universities. Technologists are in constant contact with the real needs of customers today, while faculty develop basic science that (may) have application far into the future. After all, innovations like Facebook, Microsoft, Apple, and even Ethereum were launched by young men who didn't even graduate from college. Yet, it's no accident Silicon Valley and Route 128 both emerged out of proximity to our nation's greatest coastal universities. So there's certainly a correlation between universities and the tech sector. But even so, Bitcoin is different. Bitcoin has an even tighter relationship with its intellectual and academic roots. To understand, we must peer into its history.

At the turn of the century, a ragtag band of cryptographers, computer scientists, economists, and libertarians (the Cypherpunks) exchanged messages over an internet mailing list. This was an obscure electronic gathering of a diverse cadre of scientists, technologists, and hobbyists to develop and share ideas of advancements in cryptography and computer science. Here's where some of the early giants of applied cryptography spent time, like Hal Finney, the architect of PGP.

It was on this mailing list that the pseudonymous creator of Bitcoin, Satoshi Nakamoto, announced his solution for an electronic payment system. After that announcement, he began to field questions from the forum on both the concept and its execution. Shortly thereafter, Satoshi provided the full implementation of Bitcoin. This allowed participants of the forum, like Hal Finney, to download the software, run it, and test it on their own.

The white paper itself bears similarity to academic research. It follows the structure of an academic paper, has citations, and looks similar to what any paper in computer science may look like today. Both the white paper and the conversations around it reference prior attempts at implementing the proof of work algorithm, one of the core features of Bitcoin. For example, it cites HashCash from 2002, also part of the corpus of knowledge that preceded Bitcoin. Hal Finney himself developed early prototypes of proof-of-work, such as reusable proof of work (rPow), trying to solve the problem of eliminating spam in emails.

Thus, Bitcoin didn't fall out of the sky, but rather emerged out of a long lineage of ideas developed over decades, not days or weeks. We tend to think of technology as operating at warp speed, changing rapidly, and driven by ambitious, young college dropouts. But Bitcoin wasn't based on "move fast and break things." It was and is the opposite: a slow, careful deliberation based on decades of real science practiced not by kids but more like their parents. The cryptography forum was similar in nature to an academic research seminar, where professional scientists politely but insistently attempt to tear down ideas to arrive at the truth. Though the concept of a white paper is now all the rage among alternative cryptocurrency coins and tokens, it's the hallmark method of communicating ideas within the professional research community.

Even though the cryptocurrency economy today occupies center stage in the financial press and a growing share of national attention, Bitcoin, when it emerged, was as far from this as possible. It was obscure, technical, and very fringe. In its long gestation from ideas around for decades but unknown except to a small circle of cryptographers, economists, and political philosophers, Bitcoin shares more in common with other radical innovations, like the internet, the transistor, and the airplane. And just like those innovations, the story of Bitcoin is the triumph of individual reason over collective misperception. Just as the Wright Brothers proved the world wrong by showing man could fly even though physicists claimed, just the year before, that it was mathematically impossible, so too did Bitcoin confound the naysayers by building digital scarcity for the first time ever.

Why focus on Bitcoin rather than some of the other cryptocurrency tokens, like the #2 Ethereum? If you look under the hood, the majority of the innovation of cryptocurrency came from Bitcoin itself. For example, Ethereum relies on the same elliptic curve as Bitcoin, utilizing the same public key cryptography. Bitcoin emerged over a long gestation period and secret development by a pseudonymous applied cryptographer and was released and debated in an obscure mailing list. For this reason, Bitcoin shares many similarities to the arcane academic circles that occupy modern universities. It was not a professional cryptographer who made Ethereum, but rather a teenager who even admits he rushed its development. Thus, it's only Bitcoin with deep connection to the academy, while the more incremental innovations crowding the cryptocurrency space now are more similar to the small advances taken in the modern technology sector.

DIFFERENCES

Bitcoin differs from the academy in important ways. Most significantly, it's fundamentally interdisciplinary in a way universities today are not. Bitcoin fuses three separate disciplines: mathematics, computer science, and economics. It's this fusion that gives Bitcoin its power and shatters traditional academic silos.

Public key cryptography has been the major innovation in applied cryptography and mathematics in the past 50 years. The core concept is simple: users can secure a message with a private key known only to themselves that generates a public key known to all. Therefore, the user can easily distribute the public key without any security consequence, as only the private key can unlock the encryption. Public key cryptography achieves this through hash functions, one-way transformations of data that are impossible to reverse. In Bitcoin, this occurs through elliptic curves over finite fields of prime order.

But public key cryptography isn't enough. Because Bitcoin seeks to serve as an electronic payment system, it must solve the double spending problem. If Alice pays Bob in Bitcoin, we must prevent Alice from also paying Carol with that same Bitcoin. But in the digital world, copying data is free and therefore, preventing double spending seemingly hopeless. For this, Satoshi utilized the blockchain, a construct from computer science. Cryptographer David Chaum laid the groundwork for the concept of the blockchain as early as 1983, in research that emerged out of his computer science dissertation at Berkeley.

The blockchain is a linked list that points backwards to the original (genesis) block. Each block contains hundreds of transactions, each transaction containing the ingredients for transferring Bitcoin from one account to another. The blockchain solves the double spending problem because it is distributed, i.e., publicly available to all nodes on the Bitcoin network.

These nodes constantly validate the blockchain with new transactions added only when all other nodes on the network agree (consensus). In our prior example, when Alice pays Bob, this transaction enters the blockchain, which all nodes observe. If Alice tries to use those same Bitcoins to pay Carol, the network will reject that transaction since everyone knows that Alice has already used those Bitcoins to pay Bob. It's the distributed, public nature of the blockchain that prevents double spending, a problem unique to electronic payments.

Indeed, Satoshi designed the blockchain specifically as a solution to double spending. It's inherently inefficient, as it requires the entire network to constantly validate and reproduce the same data. This is also why most applications of blockchain technology outside of Bitcoin make little sense, as it forces an inefficient solution custom built for electronic payments onto other applications that would be efficiently

solved with central databases. The notion of a blockchain as a reverse linked list by itself is not revolutionary in computer science, but its distributed nature specifically designed to prevent double spending is.

Even so, cryptography and blockchain aren't enough. There needs a reason for the network to secure the blockchain. This is where the economics of Bitcoin shine. Satoshi proposed a group of computers that would prove that the history of transactions did in fact occur. This proof requires costly work to be done. Satoshi solved this by setting up a tournament in which individual computers (called miners) would compete to solve a hard math problem. The winner would receive newly minted Bitcoin, which the network would release. The math problem must be sufficiently challenging that the only way to solve it is to deploy more computational resources. Bitcoin mining requires real computation and therefore real energy, similar to gold mining a generation ago. But unlike gold mining, the supply schedule of new Bitcoin is known by everyone.

The microeconomics of mining is the design of a contest that rewards new Bitcoin to miners that solve a puzzle. This is a form of a microeconomic mechanism, i.e., a game economists design where individual agents compete for a reward. The macroeconomics of Bitcoin pertain to the supply schedule, which adjust predictably over time, with the block reward reducing in half every four years. This forces the constraint of 21 million Bitcoins to ever be mined. This inherently limits the inflationary growth of the currency and imposes a constraint to which no fiat currency today must adhere. The difficulty of the underlying puzzle adjusts every two weeks regardless of the computing power of the network, providing a robust implementation despite exponential advances in computing power in the decades since Bitcoin launched.

This interdisciplinary feature of Bitcoin is existential, not incremental. Without any of its three components (public key cryptography, a backward linked blockchain, and a mining contest using proof of work), Bitcoin would not function. By itself, each of the three components consisted of a coherent body of knowledge and ideas. It was their combination that was the genius of Satoshi. So too will future radical innovations need to link together multiple disciplines in existential ways, without which their combination would not survive.

WHY NOT THE ACADEMY?

Why could Bitcoin not have emerged out of the academy? First, Bitcoin is inherently interdisciplinary, yet scholars at universities are rewarded for excellence in single domains of knowledge. While Bitcoin fuses ideas from computer science, mathematics, and economics, it is unlikely any single university faculty would have the breadth of knowledge necessary for interdisciplinary consilience.

Second, the academy suffers from incrementalism. Academic journals explicitly ask their authors for the incremental contribution their work provides to the literature. This is how knowledge advances, inch by inch. But Bitcoin, like other radical innovations in history (such as the airplane and the transistor) made giant leaps forward that would likely have not survived the peer review process of the academy.

Third, Bitcoin rests on libertarian political foundations that are out of favor among the mainstream academy, especially professional economists. Baked into the software are algorithmic representations of sound money, where the Bitcoin protocol releases new Bitcoin on a predictable schedule. This is very different from the world we live in today, where the Federal Open Market Committee has full discretionary authority on the money supply. The Cypherpunks who vetted Bitcoin v0.1 shared a skepticism of collective authority, believing technology and cryptography can provide privacy to individuals out of the watchful eyes of the government, or any large organization.

Most economists do not share this skepticism toward central authority. At least the social science community never took Bitcoin seriously. Besides, the Federal Reserve has an outsize role in both funding and promoting mainstream academic economic research. It recruits from top PhD programs, hires bank presidents and governors who were former professors of economics, and encourages its staff to publish in the same academic journals as the academy. It is no wonder the university of faculty, influenced by the culture of the Fed, would not embrace technology that radically replaces it.

I asked all living Nobel laureates of economics to speak at the Texas A&M Bitcoin Conference, and all but one declined. Some acknowledged they do not know enough about Bitcoin to warrant a lecture; at least they were honest about the constraints of the disciplinary model in which they have so successfully thrived. Others, like Paul Krugman, view cryptocurrencies as the new subprime mortgage (he also once predicted that the internet would have the same impact on the economy as the fax machine). Academic economists dedicated almost no attention to Bitcoin's rise, and even now remain ignorant of how the Bitcoin blockchain works, despite its being the only real innovation in finance in this past decade.

Bitcoin is first and foremost an intellectual contribution. It doesn't require a deep knowledge of industry, special insight into the current practices of firms, knowledge of idiosyncratic details of the labor and capital markets. It didn't build from existing practice, but rather off existing theory. Bitcoin emerged unapologetically out of the land of ideas, and should, in some sense, have come from the academy. An academic economist possibly could have designed the mining tournament, a computer scientist developed the blockchain, and a mathematician developed public-key cryptography. It takes an unlikely fellow (or team) to combine these three innovations. Universities develop faculty with deep expertise in the individual disciplines but do nothing to tie the disciplines together in the way Bitcoin does. For this reason, Bitcoin couldn't have emerged out of the university, even though it rests on disciplines well established within the university. The problem isn't the knowledge itself, but its organization. And therein lies the opportunity.

HOW DID WE GET HERE?

The academy, in its current form, is not suited for innovations like Bitcoin. After students enter graduate school, they learn the techniques of their own discipline, which they publish in specialized journals that earn them tenure and future academic recognition with a small set of peers within that discipline. These isolated corridors of knowledge have ossified over centuries since the early universities. How did this happen?

There are two primary trends in the academy since World War II. By far, the most important is the digital revolution. As computing power became accessible to anyone, the objective of science shifted from building theory to measurement. Suddenly, a wide array of social and natural science data was available to researchers on any laptop in the world. The growth of the internet spread data sharing and data availability, and advances in micro processing power made large analysis of data cheap and easy.

The academic community shifted en masse to data analysis and moved from trend to trend on 10-to-15-year cycles. The first cycle was on summary statistics and variance analysis; the second was on linear regression; the third on machine learning. When problems arose in the specific domain of each discipline, scholars rarely returned to their underlying theory for revision. Instead, they simply fed more data into the machine, hoping measurement error and omitted variables were to blame.

The growth of big data and statistics, in concert with machine learning, has led us to now, where artificial intelligence (AI) is a black box. No researcher can fully explain what exactly AI is doing. At the same time, questions have become smaller. Before, development economics as a field would ask, "Why is Africa so poor?" Now, research in the field asks whether placing a sign on the left or the right side of a bathroom door is more likely to lead to usage.

This preoccupation with causality is intellectually worthwhile but comes at a high price, as often researchers must narrow their domain to behaviors that are easily observable and measurable. The large, complex, and mathematical theories developed after World War II were largely untestable, and so empirical researchers abandoned those theoretical foundations. Where once academics held the intellectual high ground by asking the biggest questions of the day, now empirical research dominates academic journals. Experimental physicists and empirical economists alike cite mostly other data-driven work.

As computers filtered throughout our society, students were exposed to computation earlier in their lives. By the time they arrived in college and graduate school, they already had basic facility with data manipulation and analysis. Why bother with mathematics when some simple experiments and linear regressions can provide tables of results that can be quickly published? Over time, students gravitated toward data work as the academic profession slowly migrated away from math.

It became far easier for journals to accept papers with some small experimental or empirical fact about the world. Given that editors and referees make decisions on academic research on a paper-by-paper basis, there's no overarching evaluation of whether the body of empirical and experimental work truly advances human knowledge. As such, data analysis has run amuck, with teams of researchers making ever more incremental advances, mining the same core data sets, and asking smaller and more meaningless questions. Does rain or sunshine affect the mood of traders and therefore their stock picks? Can the size of a CFO's signature on an annual statement measure his narcissism and predict if he will commit fraud? (I'm not making this up—see the April 2017 issue of the *Journal of Accounting Research*).

One might think that advances in computation would have led research to verify some of the theories developed after World War II, but this has not been the case. In technical terms, many of those complex models are endogenous, with multiple variables determined in equilibrium simultaneously. As such, it's a challenge for empirical researchers to identify specifically what's happening, such as whether increasing the minimum wage will increase unemployment, as Economics 101 suggests. That has led to a turn to causality. But causal inference requires precise conditions, and often those conditions hold only in a few specific examples, like US states that adopted anti-abortion laws at different times. The Freakonomics revolution in economics may not dominate the Nobel prizes, but it certainly has influenced the majority of published social science research.

The chief problem with this data-driven approach is that it is ultimately backward-looking. Data, by definition, is a representation of the world at a point in time. The entire fields of business and economics research are now almost wholly empirical, where scholars race to either gather new datasets or use novel and empirical techniques on existing datasets. Either way, the view is always from the rearview mirror, looking back into the past to understand what did or didn't happen. Did low interest rates cause the great financial crises? Do abortions reduce crime? Does the minimum wage reduce employment? These questions are fundamentally preoccupied with the past, rather than designing new solutions for the future.

The second trend has been the shrinking of the theory community, both inside and outside the academy. The theorists have vastly shrunk in number, and also have refused to collaborate with their much larger empirical and experimental colleagues.

This tribalism led theorists to write ever more complex, intricate, and self-referential mathematical models with little basis in reality and possibly no hope for empirical validation. Much of game theory remains untestable, and string theory is perhaps the most extreme example of a self-referential world that can never be fully verified or tested.

Finally, academic theory trails technology by a long time. Often, mathematicians, physicists, and economists provide ex-post rationalizations of technologies that have already been successful in industry. These theories don't predict anything new, but rather simply affirm conventional wisdom. As the complexity of theory grows, its readership falls, even among theorists. Just like everything else in life, the tribalism of theory leads the community to act as a club, barring members who don't adopt its arcane language and methods.

Thus, we have arrived at something of a civil war. The theory tribe is shrinking year by year and losing relevance to reality, while the empirical/experimental data community grows, asking smaller questions with no conceptual guidance. Both academics and technologists are left in the dark about which problems to solve and how to approach them.

This also has led to a pervasive randomness in our collective consciousness, causing us to blow in whatever direction the winds of the moment take us. Economics has well-established theories of markets and how they function, yet technology companies are massive marketplaces unmoored in much of that economic theory. Computer science rests on a sturdy foundation of algorithms and data structures, yet the theory community is obsessed with debates on computational complexity, while trillion-dollar tech companies perform simple A/B tests to make their most significant decisions.

We've reached a tipping point in the scale of human knowledge, where scholars refine their theories to ever precise levels, speaking to smaller and smaller communities of scholars. This specialization of knowledge has led to hyper specialization, where journals and academic disciplines continue to divide and subdivide into ever smaller categories. The profusion of journals is evidence of this hyper specialization.

FROM SCIENCE TO ENGINEERING

Much future innovation will occur at the boundaries of the disciplines, given that much knowledge has already been discovered within existing disciplines. But there must be a greater transformation. Universities today still largely adopt the scientific method, establishing knowledge for its own sake and seeking to know the natural, physical, and social world. But we should not stop there. Given their fundamental knowledge, scientists are in the best position to engineer better solutions for our future. Moving to an engineering mindset will force academics to design and implement solutions to our most pressing problems. In the long term, it also will close the gap between the academy and industry.

The pressure students face to search for jobs and start companies, which takes a toll on their academic coursework, emerges because of the gap between the needs of the market and the academic curriculum. Were this gap to close, they could spend time in college building better solutions for the future, and this cognitive dissonance would dissipate.

This transformation has already begun in some disciplines, like economics. One of the most successful applied areas of economics is market design, which unambiguously adopted an engineering mindset and delivered three Nobel prizes in the past decade alone. These scholars came from engineering and adapted game theory to build better markets that can work in the real world—such as better ways to match

kidney donors to recipients, students to schools, or medical residents to hospitals. They also designed many of the largest auctions in use today, such as the spectrum auction of the government and the ad auction within Google. There's no reason the rest of the economics profession, or even the rest of higher education and academic community, cannot similarly position themselves toward adopting more of this engineering mindset.

Over time, closing this gap between the academy and industry will relieve much of the public outcry against escalating tuition and student debt. Once students and professors orient their research to develop better solutions for society, so too will their students and, in the long term, the companies that employ them. Students will no longer resent their faculty for spending time on research rather than teaching if that research directly creates technologies that ultimately benefit the students, future employers, and society at large. Over time, this naturally will close the skills gap that America currently faces. Universities no longer will need to focus on STEM skills explicitly, but rather focus on providing technological solutions that will ultimately draw heavily from the STEM areas anyway.

A CALL TO ACTION

How can we reform higher education to produce the next Bitcoin? Of course, the next Bitcoin won't be Bitcoin, but rather, a first principled innovation that conceives of an old problem in an entirely new way. I have three specific recommendations for university culture, priorities, and organizational structure.

First, the academy must more explicitly *embrace engineering*, even on the margin, more than science. The Renaissance and the Age of Reason have led American higher education to celebrate science and knowledge for its own sake. The motto for Harvard is "Veritas," or "truth," while that of the University of Chicago is "Crescat scientia, vita excolatur": "Let knowledge grow from more to more, and so human life be enriched." And these universities, based on the scientific and liberal arts traditions, have done much to establish the corpus of knowledge necessary for human progress.

But this past half century has been the age of the engineering universities, with Stanford and MIT competing to build solutions for the world, not just to understand it. This ethos of engineering should extend beyond engineering departments, but even, and especially, to social science. For example, require all freshmen to take a basic engineering class, to learn the mental framework of building solutions to problems. Economists have articulated the benefits of sound money for generations, but only through an engineered system like Bitcoin can those debates become reality.

This shift in engineering is somewhat happening within the social sciences. For example, the recent Nobel prizes given to Paul Milgrom and Bob Wilson in economics celebrated their work in designing new markets and auctions to solve real problems in resource allocation problems that governments and society face. This community of microeconomic theorists is still a small minority within the economic profession, yet its work blends theory and practice like no other field and deserves higher representation among practicing scholars. Universities should abandon the forced equity in treating all disciplines as equal, allocating an even share of faculty lines and research dollars to every discipline, no matter its impact on society. Instead, prioritize disciplines willing and able to build solutions for the future. This culture must come from the top and permeate down toward the recruiting decisions of faculty and students.

Second, *reward interdisciplinary work*. The traditional, centuries-old model of deep disciplinary work is showing its age, while most of the exciting innovations of our time lie at the boundaries of the disciplines. Universities pay lip service to inter-disciplinary work as a new buzzword across college campuses, but unless incentives for faculty

change, nothing else will. Promotion and tenure committees must reward publications outside of a scholar's home discipline, and especially collaborations with other departments and colleges. While large government agencies, like the National Science Foundation, have increased allocation of funding toward cross-disciplinary teams, when it comes time for promotion and tenure decisions, faculty committees are woefully old-fashioned and still reward scholars within rather than across disciplines.

Over time, I expect this to change as the older generation retires. But the most pressing problems of society cannot wait, and universities should pivot faster now. Unless promotion and tenure committees explicitly announce recognition for interdisciplinary work, nothing else matters.

Third, the academy must *aim high*. Too often, academic journals are comfortable seeking incremental contributions to the fund of knowledge. Our obsession with citations and small improvements inevitably leads to small steps forward. Academic communities have a reflexive desire to be self-referential and tribal. Therefore, scholars like small conferences of like-minded peers. Some of the biggest steps forward in the history of science came from giant leaps of understanding that only could have occurred outside the mainstream. Bitcoin is one, but not the only, example.

Consider the discovery of the double helix, the invention of the airplane, the creation of the internet, and more recently, the discovery of the mRNA sequence for the COVID-19 vaccine. True progress comes from unapologetically tossing out the existing intellectual orthodoxy and embracing an entirely fresh look. Our standards of excellence for our faculty and students must insist they aim to solve the biggest problems facing humanity. Too often this discourse is silenced from campus, and over time, it erodes the spirit of our young people. To overcome this, allocate research funding based on impact, and make these requirements strict.

The vast increase in wealth from the technology sector has put various pressures on campus. For one, it induces young students to drop out and start new companies, following in the footsteps of the young founders who dominate the technological and financial press. But this happens only because there's a rift between the rewards of the market and the activities of the University. Remember that Bitcoin emerged from a small community of intellectuals seeking to engineer a solution to an ancient problem using new technology. This easily could have occurred within the academy, and in some sense, it should have.

The corporate firm, either startup or established, is the natural locus for incremental innovation. The constant noise of customer needs, investor demands, and industry folk knowledge make it a natural place for small changes in society's production possibilities frontier. Radical innovation is uniquely suited to the academy, with its longer, more deliberate time scale, access to deep science, and isolation from the noise of the market. But it's up to the academy to rise to that challenge. Let Bitcoin inspire us, so the academy is the quarterback, not just the spectator, to the next radical innovation of our time.