2015

Autonocoin: A Proof-of-Belief Cryptocurrency

Michael B. Abramowicz
George Washington University Law School, abramowicz@law.gwu.edu

Autonocoin: A Proof-of-Belief Cryptocurrency

Michael Abramowicz, *Professor of Law, George Washington University*

***Abstract*—This paper proposes a self-governing cryptocurrency, dubbed Autonocoin. Cryptocurrency owners play formal tacit coordination games by making investments recorded on the block chain. Such investments represent bets about the focal point resolution of normative issues, such as whether a proposed change to Autonocoin should occur. The game produces a result that resolves the issue. With a typical cryptocurrency, the client software establishes conventions that ultimately lead to the identification of the authoritative block chain. Autonocoin completes a circle by making transactions on the block chain determine the authoritative client software. The distributed consensus mechanism embodied by formal tacit coordination games, meanwhile, can make other types of decisions, including which of competing block chains is authoritative and whether new Autonocoins should be rewarded to benefit those who have taken actions to benefit Autonocoin. This establishes a unique funding model for a cryptocurrency, and it addresses objections to cryptocurrencies issued predominantly to the initial founders, as well as to those that encourage wasteful mining activities.**

***Index Terms*—cryptocurrency, Autonocoin, Bitcoin, tacit coordination, focal point, Schelling point.**

I. Introduction

Cryptocurrency software establishes conventions that determine whether particular cryptocurrency transactions are valid and which collection of valid transactions (generally known as a "block chain") is authoritative. Each instance of the client software maintains a copy of the block chain and updates it based on these conventions. But the client software itself does not establish conventions for making decisions about the future direction of the cryptocurrency, such as whether the conventions or other aspects of the software should be changed or whether particular individuals or firms should be rewarded for taking actions supporting the cryptocurrency. Those decisions are exogenous to the cryptocurrency, made in open-source projects and/or private firms creating new cryptocurrencies.

This paper proposes a new cryptocurrency, which it calls "Autonocoin" to emphasize that the coin is autonomous.[1] The coin would govern itself in the sense that all relevant decisions about the cryptocurrency would be determined by cryptocurrency transactions recorded in the block chain. For example, a computer running the client software could determine whether to download and launch a proposed change to the client software by consulting the block chain, applying an algorithm to determine whether the proposed change had been approved by the Autonocoin community.

A more important feature enabled by self-governance is the ability for Autonocoin to reward coins to individuals who have promoted the cryptocurrency (whether by coding, providing liquidity, adopting the currency on Internet storefronts, or by contributing in some other way). The awards would be proportional to the community's perception of the size of the contribution. This avoids problems with two alternative systems for distributing cryptocurrency – where founders allocate the initial cryptocurrency (often to themselves) and where cryptocurrency is automatically distributed to those who "mine" it – and can provide better incentives for fostering continued contributions to a cryptocurrency's development.

But how does Autonocoin govern itself, and how does it work more generally? The short answer is that Autonocoin is built on the concept of a *formal tacit coordination game*, which can be used to identify a community consensus resolution of an issue about which there might be disagreement. This is different from simple vote counting (which is infeasible for a cryptocurrency in which individuals are unknown) or auctioning decisions to the highest bidder. Participants in these games will have economic incentives to consider not their own preferences, but what they think others will think (about what others will think, and so on in infinite regress) the best thing to do about a particular issue is.

Part II of the paper explains how formal tacit coordination games work and how they can be applied to the tasks of determining rewards and determining whether a cryptocurrency should be improved. This part develops ideas that are explored in more detail in an earlier paper, Abramowicz (2015), but it presents them in a different light to highlight issues in the next part. The reader need not have read that paper, but a familiarity with the basic mechanics of Bitcoin and cryptocurrencies generally is assumed in the remainder of this paper.

Part III extends the earlier work by explaining how formal tacit coordination games can be used to perform the central task of a cryptocurrency, in particular determining which of multiple block chains is authoritative. The system described is called "proof of belief," and the central idea is that if a controversy develops as to which block chain is authoritative, this can be resolved through a tacit coordination game. Thus, the block chain that cryptocurrency owners *believe* is authoritative will be recognized as such in a formal tacit coordination game.

Part IV concludes, highlighting that Autonocoin is the first

January 2015. Prof. Abramowicz is with the George Washington University Law School, 2000 H St. NW, Washington, DC 20052 (e-mail: abramowicz@law.gwu.edu).

[1] The name, broken down as "auto no coin," also highlights that the cryptocurrency works in an automatic way and is virtual.

proposed cryptocurrency that would not merely be a peer-to-peer cryptocurrency, but a peer-to-peer institution, one whose decisions are made without any central intervention. Its adaptability and potential for rewarding contributors makes it a promising candidate to become a relatively focal cryptocurrency in the already existing tacit coordination game that determines cryptocurrency value.

II. Tacit Coordination

Inherent in the proposal for Autonocoin is the recognition that decisions can be made without dictatorship or voting. They can be made based on *tacit coordination*. Section II.A summarizes the notion of tacit coordination, and Section II.B explains how cryptocurrencies already reflect tacit coordination. Section II.C explains the notion of a *formal tacit coordination game*, and Section II.D describes specific formal tacit coordination games that could be used to make binary decisions (for example, whether to approve a proposed change to software) or quantity decisions (for example, how large a reward someone should receive).

A. Schelling Points

The idea of tacit coordination games was developed by the game theorist Thomas Schelling (1980, pp. 55-56). He imagined the following dilemma: Two people are to meet in New York on a particular day, but they have not established exactly where or when to meet and have no means of communication. Where should each go? In a survey he posed, a majority of respondents gave the same answer: the information booth at Grand Central Terminal at noon. That location and time are in some way "focal" or salient, standing out from other alternatives.

Schelling was a game theorist and recognized that this was a game that could not be solved in the traditional way. There are multiple equilibria. One will do whatever one expects one's partner to do, and if that is to meet at Broadway and 73rd St. at 2:17 am, so be it. Perhaps there could be some situation in which that place and time would stand out, for example for the reunion of two lovers who met there and then for the first time. Even absent unusual circumstances, there could be disagreement about the appropriate focal point. Some individuals would undoubtedly end up disappointed on top of the Empire State Building. But what Schelling's example establishes is that sometimes, as a result of social cognition, some answers are likely to appear more salient than others.

Two extensions to Schelling's observation will be helpful to the proposal to be developed later. First, a tacit coordination game could occur across time. Suppose *A* must leave a package for *B* one day, and *B* must pick it up the next day. This is a hard problem absent further facts in New York (*A* would not want to cause a terrorist alert), but under some circumstances there might be a solution to it, and the problem in any event is a problem in tacit coordination.

Second, a tacit coordination game can be *normative*. That is, two or more people may wish to coordinate in a way that depends on their shared normative conception of a particular issue. Take the dilemma of a member of Congress listening to the President's State of the Union address. Each member may wish to clap when others in the member's party do so, yet not wish to observe others before starting. This requires an analysis of the normative content of a statement by the President and an assessment of whether the party approves or disapproves of that statement, as well as whether it is an important enough statement and an appropriate enough time for applause.

B. Tacit Coordination for Cryptocurrencies

Existing cryptocurrencies already depend inherently on tacit coordination. As Kroll et al. (2013, p. 2) note, those who transact with or mine cryptocurrency "must maintain consensus (1) on the rules to determine validity of transactions, (2) on which transactions have occurred in the system, and (3) that the currency has value."

To be sure, (1) is more akin to finding someone in New York after one has already agreed to meet someplace than it is to coordinating without communication. After all, the rules are already embodied in the client software. Nonetheless, this highlights that human activity often reflects tacit coordination based on social facts perceptions of which are widely shared. Someone could propose that Bitcoin suddenly switch to a different set of rules, or even create software that would do so. The switch will occur only if there is some process by which the community concludes not only that this would be beneficial, but also that others (anticipating what others will think, again in infinite regress) will think the benefits sufficient to justify the change at some particular moment.

Tacit coordination of this sort is extremely difficult. Imagine a very large group of New Yorkers who meet every day at Grand Central. Perhaps many individually think that the Empire State Building roof would be a better meeting place, but no one is likely to make the change absent some reason to think others not only agree, but agree at the same time. This usually requires some form of centralized coordination, such as a vote or at least a group discussion including explicit commitments by some people that they will leave. This type of coordination is especially difficult if there are people with strong opinions on both sides and if the costs of a coordination failure are high. In those circumstances, the status quo will provide a very strong focal point.

The reason that Bitcoin is able to evolve is that the status quo is one in which the Bitcoin community (particularly the miners) accept whatever changes are implemented by those who control the official repository of the software code. Moreover, those individuals have explicitly committed to conservative development principles, to not making changes where there is substantial disagreement about their value. The combination of these factors allows the community to maintain consensus on the rules for determining transaction validity and the software code generally. But it also means that there is a strong status quo bias.

This highlights an important point: Bitcoin and other cryptocurrencies provide peer-to-peer means of validating currency transactions but decisions about how those cryptocurrencies should evolve are not peer-to-peer. There may

be a temptation to lump open-source software development (used for Bitcoin and many other cryptocurrencies) and peer-to-peer mechanisms. The observation is that even if there are only some individuals authorized to make a change to an open-source repository, the ability of anyone to fork the code and make a new repository constrains the open source developers. (See Nyman & Lindman 2013). Thus in neither case, the argument might go, do authoritative decision makers exist.

The difference, though, is that peer-to-peer mechanisms (like the Bitcoin protocol or various file-sharing protocols) operate pursuant to clearly identified rules, encapsulated in algorithms and computer programs. It is easy to tacitly coordinate once there is agreement on such rules. The question of which open source software repository should be seen as authoritative (and thus of whether those protocol rules should be) is far less clearly defined. The result of this is that majority or even supermajority views about how protocols should change are not likely to control. Informal tacit coordination will have a strong status quo bias (which may or may not involve strong deference to a central decisionmaker) because of the absence of a mechanism for peers to resolve disputes among themselves. Traditional mechanisms for resolving disputes, such as voting and use of representatives, are centralized (and thus not peer-to-peer), but the absence of one of these mechanisms does not mean that decisionmaking is peer-to-peer in any meaningful way.

The difference can be highlighted by focusing on (2), the coordination about which transactions are authoritative. This *is* accomplished peer-to-peer in Bitcoin and most other cryptocurrencies. There are clear rules for determining which transactions are valid, and so the tacit coordination is around these rules rather than around particular transactions. This highlights the challenge of this paper. The goal is to build a cryptocurrency in which tacit coordination is used to make decisions about which of competing block chains is valid. But to make this possible, we will first need to develop tacit coordination around rules for formalizing tacit coordination more generally, so that subjective and potentially controversial questions can be resolved.

Perhaps the most interesting tacit coordination is (3), coordination around the notion that a particular cryptocurrency is valuable. To some extent this is no different from other securities, whose value derives largely from the expectation of value in the future. But companies produce goods and services, and their securities' values are derivative of this. The tacit coordination is accentuated with cryptocurrencies, because a cryptocurrency has value *only* because others will value it in the future. Of course, this is true of other currencies as well, but with fiat currencies, the perception of central government power makes tacit coordination easy. Sure, U.S. citizens can transact in any currency but the fact that dollars are legal tender and thus *must* be accepted as a means of payment makes the focal point obvious, at least where the government's power is unchallenged.

What makes (3) particularly difficult is that there are many cryptocurrencies (over 500). Thus, even if there is a consensus that cryptocurrencies as a group have value (because transacting in them is cheaper, or more private, or more independent of governmental authority), that does not resolve the question of *which* cryptocurrency has value. This is a task entirely for tacit coordination. Each person will assess the relative value of a cryptocurrency based on how he or she expects others to expect others to value it. But what makes one cryptocurrency more salient than another?

The salience of a cryptocurrency can in part be normative. For some users of cryptocurrency, the appeal of a cryptocurrency is inherently ideological, generally stemming from libertarian premises. The relative appeal of individual cryptocurrencies may then depend on their ideological appeal within the cryptocurrency community, as well as their anticipated normative appeal within the broader group of consumers who represent potential users of cryptocurrencies. There is a perception among some that cryptocurrencies dependent on mining are more fair than those in which cryptocurrency is distributed to founders. (See, e.g., https://twitter.com/dwr/status/496578410856984576) Later, we will argue that Autonocoin may be perceived as fairer than non-mining cryptocurrencies. This perception boosts Bitcoin and other mined currencies over others.

On the other hand, the claim that Bitcoin mining is wasteful may decrease the normative appeal of Bitcoin. And if Bitcoin mining is in fact wasteful, that will of course also decrease Bitcoin's strength in the tacit coordination game competition among different cryptocurrencies. Mining is costly, because it deflates the value of the currency, and subsidizing mining through transaction costs above marginal cost faces the same problem. The costs of Bitcoin mining may be too great if mining provides a level of security against manipulation that is either excessive or that could be obtained more cheaply with some other mechanism. The case for Autonocoin will similarly depend in part on whether its mechanism for identifying the authoritative block chain provides sufficient security. Of course, Bitcoin and other cryptocurrencies may have other flaws, real and perceived, and their relative status will depend in part on participants in the informal tacit coordination game's recognition of these flaws.

Absence of flaws in a cryptocurrency's design, however, is not sufficient. An obvious aspect of salience is whether a cryptocurrency is in common use, and another is whether the cryptocurrency was innovative. This explains much of Bitcoin's high market capitalization relative to other cryptocurrencies. After all, anyone else could take the Bitcoin reference software and make some new coin (MeTooCoin) just like Bitcoin but with a different block chain root. Whether or not the person mines a lot of MeTooCoins before releasing information about MeTooCoin to the public, MeTooCoin is not likely to be successful. It has low salience relative to Bitcoin. If MeTooCoin could easily take away market share from Bitcoin, then there would be many coins indistinguishable in market share.

A cryptocurrency that innovates, introducing features absent in Bitcoin, can become salient, and thus tacitly coordinating individuals can agree that the cryptocurrency has value. This tacit coordination may reflect a judgment that in the long term, many cryptocurrencies can coexist, with relative values

determined by tacit coordination. Or it may reflect a judgment that the upstart has some chance of becoming the top cryptocurrency of the future. If there is a perception that Bitcoin has some limitation – say, that mining is inherently inefficient, or that it lacks some feature that some other cryptocurrency has – then there might be a gradual move from Bitcoin to some other cryptocurrency. That could lead to a tipping point in which its value increases abruptly and Bitcoin's decreases and it becomes far more valuable than Bitcoin. The dynamic could be similar to a familiar one in the corporate world, where an upstart like Facebook can dominate an earlier entrant like MySpace because of a gradually increasing perception that the upstart's technology is superior.

It might seem that the Bitcoin developers should respond by mimicking the features of other cryptocurrencies. The follower in a yachting race will often angle sails in a way that is likely to increase the lead but may give the follower a chance if the wind changes direction; the leader then may mimic the follower's move to ensure that it stays ahead of the follower. With many followers, however, this strategy is less sensible. Moving in the *wrong* direction to prevent one follower from overtaking it would enable some other follower to catch up.

Nonetheless, it might seem at least that Bitcoin should incorporate all features widely regarded as beneficial. This could even involve radical change, such as from a proof-of-work to a proof-of-stake currency. But the bias toward the status quo would prevent that, even if most felt that the latter was superior. To ensure tacit coordination around Bitcoin and within Bitcoin (that is, avoiding hard forks of the currency), the developers are wisely hesitant to add new features without very strong levels of agreement. Even though the tacit coordination status quo is to allow the developers on the central repository to make decisions, making controversial ones could endanger that, so the developers themselves respect the status quo on controversial matters. The downside of this is that this conservativeness potentially helps competitors, and though any new cryptocurrency starts with the disadvantage of not being in common use, one may someday find itself in the dominant position that Bitcoin enjoys today.

That competitor, of course, would then have similar incentives to be conservative and might be vulnerable to competitors in turn. The strongest competitor to Bitcoin today is XRP, in part because it is part of a network called Ripple that facilitates currency exchange. But if it surpasses Bitcoin, then it may face competition from some other rival with some other feature that represents a significant technical departure. The developers of Ripple will likewise hesitate to add that feature if there is significant controversy about it, because of the risk of a hard fork should a large number of those operating client software refuse to integrate it.

This suggests that perhaps the most desirable feature for a cryptocurrency would be adaptability. This is an ability to introduce changes generally thought beneficial even where they are somewhat controversial, without causing a schism within the community surrounding that cryptocurrency. As long as the authoritative version of particular cryptocurrency depends on informal tacit coordination, that will be hard to achieve. There is always the danger of a hard fork, and that reduces innovation, especially for a leader. For a cryptocurrency to be adaptable, there would need to be very strong informal tacit coordination on some formal means of determining whether changes are legitimate, in much the same way as a stable mature republican government is very stable even in the face of controversy about the wisdom of particular laws. The core principle of Autonocoin (and thus one that could attract strong informal tacit coordination) is that decisions are based on a particular formal means of tacit coordination.

C. Formal Tacit Coordination

By a "formal game," I mean simply a game in which each player must make one or more discrete decisions and receives a payoff based entirely on the decisions made by the players as a whole. In a formal tacit coordination game, each player's incentive is to make a decision that will be the same as the one that the player expects other players to make. Formal game theory cannot predict the equilibrium of a tacit coordination game, but we can predict what a player will do on the assumption that there is some focal point that all players can estimate.

Consider the following simple example of a formal game: Suppose an observer of a beauty contest is told to rate the beauty of a competitor on a scale of 1 to 10. The observer of the contest is further told that there is some probability (say, 50%) that another observer will afterward be asked the same question. The first observer is promised payment according to a schedule that ensures that the amount will always be greater, the closer the observer is to the observation of the second observer. If by chance there is no second observer, then the first observer will receive nothing. If there is a second observer, however, that observer will face exactly the same incentives, asked to anticipate the answer of a hypothetical third observer, who would be asked to anticipate a fourth, and so on.

This example is reminiscent of a famous one by Keynes (1936), who described a game with "the prize being awarded to the competitor whose choice most nearly corresponds to the average preferences of the competitors as a whole." Keynes notes that such a game requires each player to anticipate "what average opinion expects the average opinion to be." But Keynes described such a game because he was skeptical that the game would reach the correct result, and indeed his broader point was to criticize the rationality of stock markets. Though Schelling had not yet written about focal points, Keynes recognized the possibility that someone might give a different answer when asked what the average participant would think than when asked simply for an opinion.

Keynes's implicit critique is that irrelevant factors might enter into the evaluations by participants in such a game. This might be especially true in Keynes's game, where the winner-take-all nature of the prize might complicate decisionmaking; one might search for a number that would be an average but that people wouldn't recognize as such. But even with our version of the game, it is possible that there might be some additional focal points that could affect decisionmaking – the middle of the spectrum (5.5), a round number (5), a lucky number (7), the

height of the candidate in feet, the order of the candidate in a group, the score of the previous contestant, and so on. Usually, though, these will tend to cancel out, and the best strategy will be to consider what one thinks the general view actually would be.

The question is ultimately empirical, but it seems highly unlikely that a participant in the simple 1-to-10 beauty game above would give a middling score to a contestant who, by conventional standards, appeared to be one of the most strikingly beautiful women in the world. At least this is so if the financial incentive is sufficiently great to matter more than considerations such as making a social statement. As long as it seems more likely that the next participant will focus on beauty than on any other approach, it will make sense for the first participant to do so as well.

Formal tacit coordination games are especially likely to elicit opinions about normative questions when there is some collective benefit from following this approach, as opposed to seeking some other focal point. As Hosni (2009, p. 37) points out, players in coordination games will often coordinate around the solution that produces the highest total payoffs to the players. It is in any one player's interest simply to anticipate the actions of the other players, but if there is a particular approach that is best for all the players (including hypothetical players), that approach then becomes a powerful focal point.

Autonocoin is built on formal tacit coordination games. If for early decisions made by Autonocoin, players were to latch onto some other focal point (such as the midpoint of the permissible range of values), then Autonocoin would fail. The players would thus lose their opportunity for earning further profits by continuing to make decisions for Autonocoin. That makes it all the more likely that each player will anticipate that the next player will make moves consistent with the normative questions posed. Meanwhile, once a norm of doing so is established, it will become entrenched. That is, once there has occurred repeated informal tacit coordination that the "right" way to play the formal tacit coordination game is to resolve the relevant normative questions, that affects what any participant thinks the next participant will do in any such game. Absent some compelling reason to expect sudden deviation from the established strategy in such games, players are likely to continue playing as before.

As long as Autonocoin's formal games are structured to produce tacit coordination, there is thus a strong chance that the results of Autonocoin's self-governance will be consistent with community views about how particular questions should be resolved. Not any formal tacit coordination game will work, however. The simple beauty pageant game above would require some mechanism for choosing random participants. While it is possible to imagine such an approach (for example, by generating pseudo-random numbers from the block chain), the particular holders of Autonocoin chosen might not want to participate. An alternative is to develop a formal game in which anyone can participate by making an investment of Autonocoin, and the result of the game can be conclusively determined by assessing the sequence of investments in the block chain.

1) Binary Decisions

Consider first binary decisions. A particularly important binary decision would be to approve or disapprove a proposed set of changes to the reference software code. Anyone could initiate a proposal to adopt a particular change to the code by creating a transaction with metadata referring to a hash of the proposed changes to the reference code, as well as perhaps other information, such as the address of the git repository with the proposed changes. (If the proposed changes did not exist or were not publicly accessibly, then other decisionmakers could reject the proposal in the formal tacit coordination game.) Others could then either support or oppose the proposal. The initial proposal would require a fixed fee, while others could allocate any number of Autonocoins in favor or against the proposal. All spending decisions would be made by sending Autonocoins to addresses based on the hash of the metadata (for example, addresses that can be generated by a hash of the original metadata hash plus "Yes" or "No").

The game would end after two conditions are met: first, at least a specified amount of time has passed (or number of blocks have been added to the block chain), perhaps a week; and second, there has been no change in what the final resolution would be over some shorter amount of time, perhaps an hour. The winning position would be the position attracting more total investment, counting the proposal fee as being in favor. Any money then spent on the losing position would then be distributed to supporters of the winning position, up to the amount spent in order of investment. Any money spent on the winning position is refunded. So, if *A* initiates a proposal by spending 5, *B* opposes by spending 6, and *C* supports by spending 5, then the proposal is passed, and *A* receives 10 (a refund of its own 5 plus 5 from *B*), and *C* receives 6 (a refund of its own 5 plus the last 1 from *B*).

This produces the coordination dynamics of a formal tacit coordination game. If the current position favors position *X*, then one will have an incentive to add enough support for *Y* to put *Y* in the lead if one thinks that either no one else will participate or else that any subsequent participants will in total be more likely to favor *Y* than *X*. Continuing the example above, after *C* moves, there are 10 Autonocoins in favor of the proposal and 6 in opposition. *D* (who could be the same as *B*) will have an incentive to add at least a little more than 4 Autonocoins if *D* expects that *E*, *F*, etc. will in total be more likely to be opposed than in favor.

D might be especially likely to do this if *D* develops an argument or some analysis supporting this position and shares that with other participants (for example, by posting on the Internet using the hash so that the argument can be found by others). Participants do not merely have incentives passively to identify the focal point, but also to try to persuade one another about what the focal point is. Formal tacit coordination games are thus a deliberative process, providing incentives to produce arguments and analysis that may change others' views about where the focal point is.

D might add more than 4 Autonocoins as a signal that it is willing to support its position aggressively, but others might do the same in the opposite direction, and if the investments

become large, that will provide incentives for more participants to enter into the competition. This is why it is important for participants to consider not only who is likely to play the game, but also who might play the game in the next round if it turns out that there is a large amount of disagreement, keeping in mind that those playing then would be thinking ahead to the possibility of even higher stakes drawing in even more participants.

Binary decisions are, of course, simple, but they can be aggregated to make more complicated decisions. Indeed, the possibility of using binary decisions to approve changes to the reference software code establishes that series of binary decisions can also be used to make decisions about evolution of a text. This text might also be a set of rules or norms concerning formal tacit coordination games should be resolved. For example, Autonocoin could be used to approve rules governing what type of support must be provided before a proposed change to the code can be considered. This might, or might not, include provision of unit and integration tests to establish that the change will work successfully, expert opinions, explanation of why the change should be made now rather than later, etc. The point is that Autonocoin can be used both to make decisions about whether to change the reference software code *and* to develop principles about how those decisions should be made. The relatively simple decisionmaking mechanism described here thus could become more elaborate over time, with code support for more complex scenarios and evolution of textual guidance.

2) Quantity Decisions

Autonocoin could also be used to make quantity decisions. This is particularly important for Autonocoin to be able to serve its role of providing rewards for those who engage in activities beneficial to Autonocoin. Someone could propose a reward for his or her own account by paying the proposal fee, and others would then determine whether a reward should be given at all (a binary decision), and if so, how high the reward should be. One important use of a reward might be to provide compensation for those who have initiated proposals that ended up receiving support. This answers a potential objection to the binary decision approach, that the first participant has little incentive to pay the proposal fee. (Subsequent participants have incentives to participate if they believe that the amounts paid so far are supporting the wrong answer.)

One way of making quantity decisions would be to simply combine binary decisions, allowing each to serve in effect as a bit in a number. But this may be unnecessarily complex. An alternative is to allow each participant (including the initial proposer) to specify the value that they believe is appropriate in metadata for the investment. That is, one might pay 10 Autonocoin in favor of the position that someone should receive 3 Autonocoin as a reward for some support of Autonocoin.

Each new participant's investment establishes a bet with the prior investor that the new participant's proposed value will be closer to the final proposed value than the previous participant's. The new participant must bet as much as the immediately prior participant, plus some additional amount that a subsequent participant can challenge in a bet. So, suppose *A* initiates by placing 1 Autonocoin on the number 50, and *B* responds by placing 2 Autonocoin on the number 75, and *C* in turn places 4 Autonocoin on 25. If there is no further activity, then *A* wins its 1 Autonocoin bet with *B* (and also receives back its own 1 Autonocoin investment), and *C* wins its 1 Autonocoin bet with *B* and also receives back the 4 Autonocoin that it invested. Note that *C*'s bet with *B* is for only 1 Autonocoin, because the first 1 Autonocoin that *B* invests is for its bet with *A*.

III. Proof of Belief

The preceding sections have shown that it is possible to design a cryptocurrency to provide the capability of playing formal tacit coordination games, thus allowing the software managing the cryptocurrency's block chain to make decisions (including a decision to update the software in a particular way) based on a community consensus. Such capabilities could be integrated into cryptocurrencies based on various distributed consensus mechanisms, including proof of work (like Bitcoin) or proof of stake (like NXT). But formal tacit coordination games are an alternative means of determining distributed consensus. This thus raises the question whether formal tacit coordination games can be used to furnish the distributed consensus mechanism underlying a cryptocurrency. This section argues that they can do so and explains how Autonocoin could be used to identify what block chain is valid.

A. The Proof-of-Belief Distributed Consensus Mechanism

Existing distributed consensus mechanisms for cryptocurrencies must accomplish three distinct but related tasks: First, a mechanism must provide for determining the validity of a transaction. This is accomplished by using digital signatures and thus does not vary across cryptocurrencies in ways relevant here. Second, a mechanism must provide a convention for determining whether the record of all transactions (the block chain) is in fact the authoritative one. The central problem that a cryptocurrency must address is the danger of double-spends, and so this mechanism must ensure that this record is sufficiently comprehensive. Third, a mechanism must provide a convention for determining whether a proposed additional block of transactions should be added to the block chain.

The proof-of-work approach accomplishes the last of these tasks by awarding new currency to the first to solve a puzzle. The puzzle is to generate a new block, consisting of valid new transactions, an arbitrary nonce, and a link to the previous block, in a way that produces a sufficiently low hash score. A miner must try an enormous number of nonce values and transaction permutations to solve the puzzle correctly. This leads directly to the mechanism for determining which of two alleged block chains is authoritative: The authoritative block chain is the one that required more work (generally, but not necessarily, the longest block chain).

Proof-of-stake approaches vary in their precise implementation, but the general idea is similar. A valid block in some versions is a block generated by a user whose "turn" it is to mine new currency; thus, each user has an incentive to

participate in the mining process, but need not solve difficult problems. In other versions, a valid block is a block consisting of transactions with sufficient "coin age," which is proportional time since they were last spent. In such a system, the valid block chain is the one that uses the greatest coin age.

For Autonocoin, let us take the challenges in the reverse order, addressing first the question of how to determine which block chain is authoritative among multiple competing block chains. This can be determined by allowing decisions on the block chain as to whether any particular block is a valid block that should be on the block chain. This is simply a binary version of the formal tacit coordination game, and any player can initiate a decisionmaking process to approve of a particular block as authoritative. If otherwise there would be insufficient incentives to initiate this decisionmaking process, rewards could be proposed for those who successfully initiated designations of blocks as valid.

That allows a measure of *proof of belief* in a particular block and in a particular block chain. The phrase follows from the recognition that any payment made in a formal tacit coordination game represents a bonded signal that the participant making it believes that others will agree with the participant's recommended decision. The measure of proof of belief in a particular block is the difference between payments made in support of a block's authenticity and payments made in opposition to a block's authenticity. The measure of proof of belief in a valid block chain (that is, one in which the hash for each block refers to the previous block) is the sum of the proofs of belief for each block. Note that a block currently viewed as invalid could still be part of the authoritative block chain; such a situation could endure in the long term as an indication that validation of a particular block was a mistake but that the mistake is too far in the past to correct.

1) The Authoritative Block Chain

The convention that Autonocoin establishes is that the authoritative block chain is the one with the highest proof-of-belief score. Critically, a user who endorses a block (or who does the reverse) places the currency invested in the formal tacit coordination game at risk. Regardless of the final determination of whether this existing block is authoritative, the transaction will be broadcast and thus may count as a spend on some later block of the authoritative block chain. (We will explain in the next section what provides incentives for individual blocks to include all applicable transactions.)

In determining which of two competing block chains is authoritative, three clarifications are necessary. First, a client should take into account all cryptographically signed transactions on both block chains. This ensures that one cannot create a block chain that has an artificially high proof-of-belief score simply by omitting transactions challenging the legitimacy of one or more blocks. If, for example, a block chain were presented with a 1 million Autonocoin verified transaction on one block chain attesting to the validity of the block chain, but the competing block chain also included a 1 million Autonocoin transaction opposing validity, the net effect would be zero proof of belief.

Second, a client should exclude from its analysis of the proof-of-belief of one block chain any transactions that are invalid according to the other block chain, as well as any transactions that are descendants of this transaction. Thus, if someone has spent the same Autonocoin in different transactions on two block chains, both of these spends will be disregarded in measuring proof-of-belief. This prevents someone from obtaining power by remembering private keys for already spent Autonocoins and then using these private keys to respend the money to bolster some other block chain in which the spending had not yet taken place.

Third, if a purported block chain indicates that someone has received currency as a result of a conclusion of a formal tacit coordination game, either in the form of winnings or in the form of a reward, but the other block chain has not resolved that game or not resolved it in the same way, then it will be disregarded. This prevents an attacker from making up a block chain in which the person has been awarded many Autonocoin and then uses some of those Autonocoin to bolster the block chain's proof of work.

With this set of rules defining the identification convention, the authoritative block chain could change, if someone were to sign a transaction manifesting sufficient belief in one or more blocks in a proposed new block chain or a transaction expressing sufficient doubt about one or more of the blocks in the block chain. The longer the block chain, however, the more expensive it would be to change the authoritative block chain significantly even for a short time. One might delete a single block by initiating a formal tacit coordination game to recognize some alternative last block (linking to the penultimate block) as authoritative. But this would require an investment greater than the proof of belief of the existing last block. Moreover, there would be little reason to do this if the existing block would be viewed as the valid one by the community. Someone would have an immediate incentive to win the investment of money in the alternative block by opposing it.

Making more radical changes to the block chain would be even more difficult. To dislodge a block 10 blocks from the end of the block chain, one would need to dislodge all of the last 10 blocks, since the question is not just whether the block is valid but whether the block chain as a whole is valid. This would require the challenger to establish a proof-of-belief score against that block at least as great as the net proof-of-belief of the last 10 blocks combined. Thus, older blocks are more secure in the block chain than newer blocks. In this sense, Autonocoin is like Bitcoin. A particular block in the Bitcoin block chain may eventually be removed from it, for example because another proposed block is added for the same spot in the block chain at around the same time. (In this case, the authoritative block chain becomes the one onto which the next valid block will be added.) But it is unlikely that many blocks will be removed.

Is Autonocoin susceptible to more serious attacks that could destabilize the block chain for long periods? The danger of attack is similar to the danger for Bitcoin. With Bitcoin, someone who obtains more mining power than everyone else combined can out-mine everyone else. That person can then

ignore even valid blocks produced by others. The attacker can then determine what transactions to put on the block chain and what to leave off and can even remove some blocks from the block chain, replacing them over time with more blocks than everyone else can add to the valid block chain and eventually becoming recognized by the authoritative software. This is the essence of a so-called 51% attack.

Similarly, in Autonocoin, someone who has a credible threat to be able to put up 51% of Autonocoins in formal tacit coordination games to achieve desired results will be able to control the currency. But it is not enough to have more Autonocoins than others who are actually participating in the formal tacit coordination games; the question is always what the resolution would be if a particular game became sufficiently controversial. So, one might need to own 51% of Autonocoins, or at least 51% of the Autonocoins owned by those who might participate in a formal tacit coordination game if the stakes became sufficiently large. This would be quite a bit to accumulate, especially since the consequence of success would be to destroy Autonocoin and any value accumulated. The most likely scenario, as with a Bitcoin 51% attack, might involve a "Goldfinger" attack by a government whose goal is to destroy the cryptocurrency rather than to profit fromit.

It is difficult to determine whether a 51% attack would be easier to mount against Bitcoin or Autonocoin, though it would likely be impossible with either. A potential weakness of Bitcoin, however, is the possibility of collusion by miners. There are already mining pools, and one pool recently came close to 50% market share. Moreover, at any time, multiple pools could theoretically decide to work together. There is little danger that they would do so to destroy the block chain or to execute a double-spend. But they might do so to change the Bitcoin protocol, either by increasing the schedule at which Bitcoins are issued or increasing mining fees. After all, miners have large irreversible investments in computers dedicated to mining and thus have an incentive to avoid bankruptcy. In the long run, Bitcoin may be determined by the interests of miners, while Autonocoin will evolve based on the perceived interests of a broader community.

At least one other type of attack, however, must be considered. An attacker might spend many already spent Autonocoins in the block chain, place the transactions performing this spending on another block chain and present this as the valid one. According to the rules above, these transactions would not count, but any descendants of the original transactions also would be discounted. The goal would be to ensure that the Autonocoins previously believed valid could not be spent in support of the true valid block chain, possibly allowing an attack with less than 51% of total power. Especially once ownership becomes widely dispersed, however, there will still be many Autonocoins that could be used to counter the attack, and the owners of such coins will have sufficient incentives to counter.

Addition of even occasional checkpointing could virtually eliminate momentary destabilization from such an attack, because coins spent before a checkpoint could not be respent on a fake block chain. A checkpoint could be added using binary formal tacit coordination games. An Autonocoin owner could initiate such a game to identify a particular block (through its hash) as a block that should be in every blockchain. After some period of time, the software would enforce a checkpoint authoritatively chosen, thus always preferring a block with known checkpoints to one without them. Checkpointing is not essential to the proof-of-belief system of distributed consensus, but it could be useful as a mechanism for stabilizing the blockchain. It could also help fight denial-of-service attacks, as is the case with Bitcoin.

2) An Authoritative Block

The inclusion of a mechanism to define the authoritative block chain is sufficient but perhaps question-begging. The formal tacit coordination games described here are designed to answer normative questions. It may sometimes be clear that an attempt to change the block chain would be normatively undesirable – for example, if valid transactions were being eliminated – but there might also be multiple proposed new blocks that seem equally normatively justifiable. They might differ in relatively minor ways, such as the order of the transactions in the block, or whether a transaction offered just as a block was being created should be included.

Autonocoin could sort out these questions over time, but it may be desirable to have some convention established at least initially. This convention could be in writing, even if not part of the software code that assembles and validates the block chain. For example, one convention might be as follows: A block should be added to the block chain every five seconds. A valid block could be defined as incorporating every transaction broadcast by a reputable third party, using a digital signature with a timestamp, ordered by timestamp (and, in the event of a tie, by hash). So, transactions timestamped between 12:00:00 and 12:00:05 would be placed in the same block. A convention could also provide that such a block would not be submitted for approval until 12:00:10, to give sufficient time to ensure that all valid transactions would be included.

This convention would still leave room for some normative questions. Most obviously, participants in the proof-of-belief formal tacit coordination games might need to assess whether a timestamper is reputable. Usually, this is not a hard question. A timestamper who issued old timestamps or delayed broadcasting transactions would be quickly spotted, and so the formal tacit coordination game would likely lead to elimination of such timestampers. As in all normative matters, there might be close cases, but once such a question were resolved for one block, that resolution would be given some weight in determining the resolution for the next block.

Another difficult case that might arise would be a timestamped transaction that somehow was not broadcast properly because of a protocol error. But again, this could likely be sorted out either in writing or by convention. For example, a convention might be that if a transaction from 12:00:03 did not become known until after 12:00:10, it would be deemed invalid. There might still be difficult cases – did the transaction become known at 12:00:09 or 12:00:11 – but the formal tacit coordination game to approve a particular block should be able to resolve such rare cases. Sometimes, the resolution might

need to be so fast that the result would determined by automated betting algorithms, but prospective players of the tacit coordination games could write bots to invest on their behalf. The nature of the tacit coordination game would produce strong incentives for such bots to converge. If the principles to which they converged came to be seen as problematic, any bad results could be revisited by humans revising the relevant conventions in writing.

This highlights the central virtues of a self-governing cryptocurrency: It can use judgment, and it can adapt. Not every contingency needs to be worked out in advance, because judgment can be applied to individual cases and issues, as well as to broader issues. As long as the underlying proof-of-belief system is foolproof, many other imperfections can exist for a time in the software, because those imperfections can be addressed both on a case-by-case basis and with new policies. With Autonocoin, one could imagine even cancelling individual transactions if they were proven to represent theft of Autonocoins. Perhaps this is advisable, or perhaps the danger of manipulation is too great, but at least with Autonocoin, such a reversal is conceivable, and the case for allowing reversals can be considered on the merits as a general matter and in specific cases if the community consensus is to allow such consideration.

It is quite possible in the end that the evolution of the cryptocurrency might lead it in some ways to be similar to existing cryptocurrencies. For example, Autonocoin could well develop a convention similar to that of XRP for determining whether a block is valid. XRP is a cryptocurrency that is part of the broader Ripple project, which is designed to provide means for easy exchange of currencies, especially fiat currencies. The decentralized nodes that maintain the XRP ledger, known as validation nodes, come to consensus about the set of transactions to be included in a new block to be approved every few seconds. They do this through a voting protocol in which each node drops transactions that do not maintain support of 50% of the nodes trusted by that node (including itself). A node will support any transaction that it knows about that is not included on or inconsistent with the ledger, but it will stop supporting a transaction if the 50% threshold is not met. This threshold is gradually raised to 60% and then still higher levels so that any transactions that may be close calls are eliminated. These transactions will generally be those that were very close to the end of the time window for that block, and they will then be added to the next block. This system leads to consistent development of consensus.

The key to XRP is the trust mechanism. This is decentralized, and there is no protocol determining which nodes a node should trust (though some recommended nodes are referenced by the official version of the client software). Recall that the only way to manipulate a ledger is to keep valid transactions off the ledger. Any single server that attempted to keep one or more transactions off the ledger would fail, because 50% or more of servers would still approve the transaction. An attacker would thus need to create many servers and allow them to become trusted over time before using them to keep transactions off the ledger. Even that, XRP advocates argue, would ultimately fail, because leaving broadcast transactions off the ledger for more than one period is a transparent form of manipulation that other clients would recognize. Those clients thus would stop trusting the manipulative clients, and the non-manipulative clients would thus agree eventually to add the transactions that had been omitted.

It would be straightforward to use the XRP mechanism or a variant as the convention for determining whether to add a block to the Autonocoin block chain. Combining this convention with the convention described above for determining the authoritative block chain, there would be no need for XRP nodes to choose which other nodes to trust, which some might see as being comparable to the designation of one or more servers as authoritative and thus undermining the peer-to-peer nature of the project. XRP nodes must know whom to trust because they identify the correct block chain by starting with a known trusted block chain and then identifying each authentic block as it is processed.

An Autonocoin software node could simply initialize by checking all ledgers available online (most, of course, would have the same hash and thus require no comparison), rather than by starting with a particular location or locations. Before participating itself in transaction approval, it must decide which other nodes to trust by monitoring nodes and eliminating any that had failed to include any transactions that were broadcast sufficiently early that they clearly should have been included. When it participated, it would count votes only from trustworthy nodes in determining whether to drop particular transactions, with longer periods of trustworthiness required for higher approval thresholds.

XRP defenders might argue that there is no need for authoritative ledger identification given a sufficiently good mechanism for identifying the newest block to be added to the block chain. Indeed, XRP has proven to be stable so far. But a mechanism for identifying the authoritative ledger may be useful for reasons other than reducing reliance on specifically designated servers. It provides some degree of insurance in case a fork does occur, whether as a result of attempted manipulation or as a result of a natural disaster that prevents synchronization of various groups of clients. Normative judgment should not commonly be needed to resolve discrepancies, but it is useful for a cryptocurrency to include a means of relying on such judgment without requiring tacit coordination on some new version of the client software. Moreover, such a mechanism eliminates the danger that the system could be attacked by a simultaneous cyberattack shutting down many of the servers, allowing a small percentage of servers to assume control.

IV. Conclusion

The most significant benefit, however, of building a cryptocurrency from the bottom-up with a system based on normative judgment is that it would highlight the cryptocurrency's ability to rely on such judgment, which then could be used for other purposes, such as providing rewards to those who act to benefit the cryptocurrency. XRP has been criticized because its currency was allocated to the founders of Ripple and their friends and associates. Economists might

prefer this to the Bitcoin approach, because it avoids wasteful rent-seeking, or at least it channels the rent-seeking to the stage of creating useful new cryptocurrencies. But Bitcoin may have ideological appeal precisely because of the absence of an entity getting rich through an IPO.

Self-governance enables a funding mechanism for Autonocoin that is different from the funding mechanisms for other cryptocurrencies. Existing cryptocurrencies issue currency in one or both of the ways exemplified by XRP and Bitcoin. First, the cryptocurrency founders may issue themselves or others (who may or may not have supported the initial development of the cryptocurrency) cryptocurrency units. This can lead to perceptions of unfairness. Second, cryptocurrency units can be issued to individuals engaged in mining or similar activities that serve to protect the integrity of the block chain and reduce the danger of double-spend transactions. This can lead to concerns, particularly for proof-of-work cryptocurrencies such as Bitcoin, that the cryptocurrency is encouraging wasteful activity. Issuance of any cryptocurrency after the cryptocurrency is initially created, meanwhile, necessarily dilutes the value of existing cryptocurrency holdings.

The result is that those who contribute to the success of a cryptocurrency do not necessarily share proportionately to their contributions. Founders of a cryptocurrency will have incentives to create a strong product, because they intend to issue cryptocurrency to themselves. But this does not produce strong incentives for them and others to improve the cryptocurrency or to engage in activities such as marketing and regulatory compliance. Such incentives exist only so long as the founding entity holds cryptocurrency, and so the incentives will decline as their stock of cryptocurrency declines. Meanwhile, founders may issue to themselves what others perceive as too much cryptocurrency.

Autonocoin can solve this problem, because the cryptocurrency community can decide whether to issue rewards, and in what amounts, to individuals or firms who have made particular contributions to the cryptocurrency. For example, an early adopter merchant might apply for a reward; presumably, earlier adopters would receive larger rewards than later adopters, and larger early adopters would receive larger rewards than smaller early adopters. Someone who develops an improvement to Autonocoin (or comes up with the idea of Autonocoin in the first place!) would likely be entitled to a reward. The cryptocurrency community would determine, by investments on the block chain, how high the reward for each of these should be.

This still leaves one question, that of who receives the initial award of Autonocoin. One possibility is to give the initial coins to some initial developers. This might be based on an agreement among them as to their initial contributions. Or the amount might be given to them as reward for a portion of their contributions, and the remainder of their contributions might be rewarded through Autonocoin itself. Subsequent rewards to others would then be proportional to the perceived value of those later rewards.

Another possibility is for the initial distribution of the currency to be distributed through an airdrop. The Aurora cryptocurrency was designed in this way, taking advantage of unique identification numbers in Iceland to stake the entire country in the cryptocurrency. This approach led to great interest in the cryptocurrency, which ended up failing for other technical reasons. A similar approach for Autonocoin might be to allow anyone with a mobile number to claim 1 Autonocoin. Some who signed up might be included in the genesis transaction. Anyone afterward could then operate a service that would certify public keys as belonging to particular mobile numbers, and the Autonocoin community could then confirm or deny individual applications for 1 Autonocoin with a formal tacit coordination game, based presumably on the reputation of the certifier. This approach would provide an egalitarian approach for initial distribution of Autonocoins, while still allowing rewards for those who did work before and after. Other cryptocurrencies, lacking a mechanism for making judgments about mobile numbers and certifiers, might not be able to perform an airdrop as easily.

REFERENCES

[1] Abramowicz. 2015. “Peer-to-Peer Law, Built on Bitcoin.” (unpublished manuscript)

[2] Hykel Hosni. 2009. “Interpretation, Coordination and Conformity.” Games: Unifying Logic, Language and Philosophy (Ondrej Majer eds.).

[3] John Maynard Keynes. 1936. “General Theory of Employment, Interest, and Money.”

[4] Joshua A. Kroll et al. 2013. “The Economics of Bitcoin Mining or, Bitcoin in the Presence of Adversaries.” http://www.weis2013.econinfosec.org/papers/KrollDaveyFeltenWEIS2013.pdf.

[5] Linus Nyman & Juho Lindman. 2013. “Code Forking, Governance, and Sustainability in Open Source Software,” Tech. Innovation Mgmt. Rev., Jan. 2013, at 7.

[6] Thomas C. Schelling. 1980. “The Strategy of Conflict.” 55-56.