LEDGER
ledgerjournal.org

ISSN 2379-5980 (online)
DOI 10.5195/LEDGER.2016.37

RESEARCH ARTICLE

Autonocoin: A Proof-of-Belief Cryptocurrency

Michael Abramowicz*†

Abstract. This paper proposes a self-governing cryptocurrency, dubbed Autonocoin. Cryptocurrency owners play formal tacit coordination games by making investments recorded on the blockchain. 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 blockchain. Autonocoin completes a circle by making transactions on the blockchain that in turn define those conventions and the expected software behavior. The distributed consensus mechanism embodied by formal tacit coordination games, meanwhile, can make other types of decisions, including which of competing blockchains 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.

1. Introduction

Cryptocurrency software establishes conventions that determine whether particular cryptocurrency transactions are valid and which collection of valid transactions (generally known as a "blockchain") is authoritative. Each instance of the client software maintains a copy of the blockchain 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 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 or private firms creating new cryptocurrencies.

This paper proposes a new cryptocurrency, which it calls "Autonocoin" to emphasize that the coin is autonomous. The coin would govern itself in the sense that all relevant decisions about the cryptocurrency would be determined by cryptocurrency transactions recorded in the blockchain. 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 blockchain, 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 otherwise). The awards would be proportional to the community's perception of contributions. This avoids problems with two

†Michael Abramowicz (abramowicz@law.gwu.edu) is Professor of Law at George Washington University, USA.
*1MioFQjewAb69Kw9rgtCDEQVscbFNLQ8DN

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? Autonocoin is built on the concept of a *formal tacit coordination game*, which can identify a community consensus resolution of a normative issue. 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 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 resolution of a particular issue is. Each participant's incentive is to answer a particular question as the participant expects hypothetical *later* participants to answer the question. Participants thus search for answers that are salient, and normative arguments in response to the questions posed create salience.

Among the tasks that a formal tacit coordination game perform is the fundamental one of confirming the validity of blocks and of determining which of multiple blockchains containing valid blocks is itself authoritative. Bitcoin accomplishes these tasks through "proof of work." Miners can create a new block only by solving a computationally difficult puzzle, and the authoritative blockchain is the one that required the most work to create. Autonocoin works through "proof of belief." Autonocoin holders can earn coins by successfully challenging a consensus as to whether particular blocks are in fact valid. In ordinary operation, this will generally be clear and mechanical based on an established convention for which transactions should be included, but the mechanism allows occasional policing of bad actors who intentionally omit transactions in the hope of achieving double-spends.

Part II 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.[1] Part III explains how formal tacit coordination games can perform the central challenge of a cryptocurrency, determining which of multiple blockchains (each internally consistent and containing only appropriately signed transactions) is authoritative.

Part IV concludes, highlighting that Autonocoin is the first 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 make it a promising candidate to become a relatively focal cryptocurrency in the already existing tacit coordination game that determines cryptocurrency value.

2. 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. This section summarizes the notion of tacit coordination, and also explains how cryptocurrencies already reflect tacit coordination. Later in this section I explain the notion of a formal tacit coordination game, and then describe specific formal tacit coordination games that could be used to make binary or quantity decisions.

Tacit Coordination for Cryptocurrencies—Thomas Schelling developed the idea of tacit coordination games.[2] 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 cannot

ISSN 2379-5980 (online)
DOI 10.5195/LEDGER.2016.37

communicate. In a survey, 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. In a tacit coordination game, multiple equilibria exist but participants have incentives to choose one that social cognition makes focal.

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. Second, a tacit coordination game can be normative. Consider members of Congress listening to the President's State of the Union address. Each member may wish to clap at the same time as others in the member's party. This requires an analysis of the normative content of the President's statement.

Existing cryptocurrencies already depend on tacit coordination. As Kroll et al. note,[3] 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." Existing tacit coordination can lead to change in cryptocurrency conventions, when the community concludes not only that the change would be beneficial, but also that others (anticipating what others will think, again in infinite regress) will think so as well. If a particular version of the cryptocurrency code (such as a Github site originally created by the founders) is seen as focal, then tacit coordination around relatively noncontroversial changes to the software conventions will likely occur.

Decisions about the evolution of cryptocurrencies are decentralized, but they are not peer-to-peer. Even if there are only some individuals authorized to make a change to an open-source repository, anyone can fork the code and make a new repository. This power constrains the open source developers.[4] But there is an inherent status quo bias. Majority or even supermajority views in favor of changes are not likely to control given sufficient opposition.

The difficulty in tacitly coordinating on controversial issues has been illustrated in an ongoing debate about whether to increase Bitcoin block size, which has led to forked versions of the cryptocurrency.[5] Existing cryptocurrencies have no objective metric for determining whether a proposal has sufficient support. Cryptocurrencies achieve a type of peer-to-peer coordination around (2), determining which transactions are authoritative, but cryptocurrencies do not include a mechanism to coordinate on (1), determining which changes to the protocol should be made. Legislatures, by contrast, can resolve controversies because there is strong tacit coordination around voting rules and thus around what rule changes are authoritative.

Many familiar mechanisms, such as allowing designated individuals to vote, could be used to provide unambiguous resolution of cryptocurrency controversies. But these mechanisms are not peer-to-peer; they necessarily privilege some decisionmakers over others. The goal of this paper is to describe how to produce unambiguous answers to controversial questions in a peer-to-peer way. The strategy is to build on tacit coordination by formalizing it, providing a specific protocol for individuals to coordinate on normative questions and determine a specific answer.

The existence of formal tacit coordination would not remove the need for informal tacit coordination. Informal tacit coordination would be applied to issue (3), determining which cryptocurrency is valuable. Cryptocurrencies with peer-to-peer decisionmaking features might through tacit coordination become relatively valuable, for the same reason that perception of cryptocurrency features affects value. But informal tacit coordination that a particular cryptocurrency using formal tacit coordination has value would eliminate the need to use informal tacit coordination to make further evolutionary decisions about the cryptocurrency.

The ability to resolve controversies may be an especially valuable feature for a cryptocurrency because this adaptability can make it easier to adopt other innovations without

ISSN 2379-5980 (online)
DOI 10.5195/LEDGER.2016.37

risking cryptocurrency forks. With a cryptocurrency that depends on informal tacit coordination to resolve disputes about the cryptocurrency protocol, there is always the danger of a hard fork. This paper thus imagines a cryptocurrency built on formal tacit coordination, using formal tacit coordination even as an alternative to "proof of work" or "proof of stake" mechanisms for identifying the authentic blockchain The formal tacit coordination mechanism, however, also could be applied to other cryptocurrencies that continue to use proof of work or stake.

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,[6] 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.

ISSN 2379-5980 (online)
DOI 10.5195/LEDGER.2016.37

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 points out,[7] 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, relying on them not only to determine the evolution of the cryptocurrency but also, as described in Part III, to validate individual blocks and blockchains. 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 blockchain), 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 blockchain.

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 the proposed changes, for example the address of a fork of a 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 blockchain), 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

ISSN 2379-5980 (online)
DOI 10.5195/LEDGER.2016.37

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 how 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 simple decisionmaking mechanism described here thus could become more elaborate over time, with code support for more complex scenarios and evolution of textual guidance.

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.

ISSN 2379-5980 (online)
DOI 10.5195/LEDGER.2016.37

(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 serves two purposes. First, with the exception of the initial proposal, each 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. Second, each investment constitutes an offer to bet with the next participant. Thus, each participant after the first must put down sufficient funds to consummate a bet plus at least a certain minimum as an offer.

For example, suppose the proposal fee and the amount one must offer as a bet are 1 Autonocoin. Then, suppose A initiates by placing 1 Autonocoin on the number 50, and B responds by placing 2 Autonocoin (i.e., 1 to challenge A plus 1 more to be subject to challenges) on the number 75, and C in turn places 3 Autonocoin (i.e., 1 to challenge B plus 2 more on offer) on 25. If there is no further activity, the resolution value is 25. Thus, 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 while also receiving back the 3 Autonocoin that it invested.

3. 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 blockchain to make decisions (including a decision to update the software in a particular way) based on a community consensus. Existing distributed consensus mechanisms, such as PAXOS or Raft, can allow distributed systems to achieve consensus about the correct answer that an algorithm should produce when individual nodes implementing an algorithm may fail and thus not implement it correctly.[8] By contrast, the formal tacit coordination game can identify normative consensus, and it may conclude with a result different from that favored by a majority, if they are insufficiently willing to back their "votes" with cryptocurrency payments.

Formal tacit coordination game capabilities could be integrated into cryptocurrencies based on various mechanisms, including proof of work (like Bitcoin) or proof of stake (like Peercoin[9]). 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 blockchain is valid.

Existing Blockchain Validation Mechanisms—The central problem that a cryptocurrency must solve is not the problem of ensuring that transactions are authorized. Cryptography makes that trivial. Rather, it is the problem of ensuring that all authorized transactions are included and ordered in an authoritative ledger. The challenge is that because the system is peer-to-peer, a participant may try to exclude a transaction, thus enabling a double-spend transaction.

ISSN 2379-5980 (online)
DOI 10.5195/LEDGER.2016.37

A cryptocurrency must thus meet two challenges: First, it must provide a convention for determining whether a proposed additional block of transactions should be added to the blockchain. Second, it must provide a convention for determining which purported record of all transactions (the blockchain) is in fact the authoritative one.

The proof-of-work approach accomplishes the first 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. A miner could exclude a transaction (and thus forfeit transaction fees) from a block, but this will not work to keep a transaction permanently off the blockchain, because another miner will include the transaction. Proof of work also leads directly to the mechanism for determining which of two alleged blockchains is authoritative: The authoritative blockchain is the one that required more work (generally, but not necessarily, the longest blockchain). One can easily generate a fake blockchain that incorporates only selected transactions, but it will be clear that the real blockchain took much more computational power to create. This protects Bitcoin against "Sybil" attacks where a large number of nodes disseminate false information about the blockchain. If one connects to even a single node with the correct information, one can verify the correctness of the information.

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 to time since they were last spent. In such a system, the valid blockchain is the one that uses the greatest coin age.

The Proof-of-Belief Distributed Consensus Mechanism—For Autonocoin, the two tasks can be accomplished as follows: First, a formal tacit coordination game can be used to determine the validity of a block. So long as some convention is set up in advance for what transactions should be included, identification of the correct block should be simple for any cooperating software client. The challenge occurs only if a noncooperator ignores the convention, for example by omitting a transaction that should be included or by including a transaction that was not announced publicly (though there is no incentive to do the latter). The formal tacit coordination game can thus choose among different blocks, each of which meets the formal requirement of containing a set of authorized signed transactions. A block will be deemed valid if approved by the formal tacit coordination game. Second, a blockchain is valid if it consists of cryptographically valid blocks, each including a hash to the prior block. The authoritative blockchain is considered to be the valid blockchain with the highest proof of belief.

The formal tacit coordination games provides this measure for each particular block and for each purported blockchain. The phrase "proof of belief" 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. Note that a transaction in support of or in opposition to a block is not included on *that* block, but on subsequent blocks to be added, and such transactions will be added for the same reason that any other transactions will be added. The measure of proof of belief in a valid blockchain (that is, one in which the hash for each block refers to the previous block) is the sum of the proofs of belief. Even if one block temporarily has a negative proof of belief, perhaps as a result

ISSN 2379-5980 (online)
DOI 10.5195/LEDGER.2016.37

of manipulation, a blockchain containing it and later blocks could still be the authoritative blockchain. In the long run, though, the normative convention for the formal tacit coordination games is that a block will have a positive proof of belief if and only if it should be included in the authoritative blockchain.

An Authoritative Block—Autonocoin could sort out which transactions should be included for a block to be valid over time, but it may be desirable to have some more specific convention established at least initially. This convention could be in writing, even if not part of the software code that assembles and validates the blockchain. For example, one convention might be as follows: A block should be added to the blockchain 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 approach should be contrasted with Bitcoin's. Bitcoin has no rule requiring specific transactions to be included in a block. It simply relies on the incentives of miners eventually to include transactions. The drawback of proof of work, of course, is that it is energy- and resource-intensive. Proof of belief can give much lower incentives to participants, who bet on whether blocks are valid or not. Ordinarily, this will be boring, straightforward, and low stakes. Client software can straightforwardly determine whether the requirements of the convention are met. If a manipulator purported to invest belief in an invalid block, bots programmed to detect this would bet against that block. Anticipating this, manipulation would rarely occur.

It is possible, of course, that esoteric forms of manipulation might arise. For example, a timestamper might issue false early timestamps, in an effort to claim that some block other than the one that was approved should have been approved. This would not allow achievement of a double-spend, but it could cause some confusion, and if it worked might allow someone to earn money in the formal tacit coordination game itself. But if someone detected such behavior, then humans would be alerted, and they could use judgment to decide how much to invest against the manipulator in the formal tacit coordination game. Programmers of software clients might then have incentives to identify such behavior and to modify their algorithms to detect it in the future. Thus, the vast majority of blocks would be approved entirely by software, but the possibility that humans can intervene means that any hypothetical clever manipulations not anticipated could be counteracted as soon as they are detected.

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 the cryptocurrency XRP for determining whether a block is valid. XRP is the native cryptocurrency that is part of the broader Ripple project, which is designed to provide

ISSN 2379-5980 (online)
DOI 10.5195/LEDGER.2016.37

means for easy exchange of currencies, especially fiat currencies.[10] 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.

So far, the Ripple consensus protocol has avoided inadvertent forks, but another cryptocurrency based on Ripple experienced a consensus fork that some argued put into question the ability of Ripple to maintain consensus consistently.[11] A mechanism for judging whether blocks should be included thus may be useful. It eliminates the need to determine whether particular nodes are trusted or not. 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.

The Authoritative Blockchain—The convention that Autonocoin establishes is that the authoritative blockchain 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 blockchain. It is this convention for determining the authoritative blockchain that protects against Sybil attacks. As with Bitcoin, even if many nodes advertise the wrong blockchain, one need connect to only one node advertising the correct blockchain to verify that this blockchain indeed is the authoritative one.

In determining which of two competing blockchains is authoritative, three clarifications are necessary. First, a client should take into account all cryptographically signed transactions on both blockchains. This ensures that one cannot create a blockchain that has an artificially high proof-of-belief score simply by omitting transactions challenging the legitimacy of one or more

ISSN 2379-5980 (online)
DOI 10.5195/LEDGER.2016.37

blocks. If, for example, a blockchain were presented with a 1 million Autonocoin verified transaction on one blockchain attesting to the validity of the blockchain, but the competing blockchain 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 blockchain any transactions that are invalid according to the other blockchain, as well as any transactions that are descendants of this transaction. Thus, if someone has spent the same Autonocoin in different transactions on two blockchains, 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 blockchain in which the spending had not yet taken place.

Third, if a purported blockchain 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 blockchain 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 blockchain in which the person has been awarded many Autonocoin and then uses some of those Autonocoin to bolster the blockchain's proof of work.

With this set of rules defining the identification convention, the authoritative blockchain could change, if someone were to sign a transaction manifesting sufficient belief in one or more blocks in a proposed new blockchain or a transaction expressing sufficient doubt about one or more of the blocks in the blockchain. The longer the blockchain, however, the more expensive it would be to change the authoritative blockchain 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. The investments manifesting proof of belief will eventually be added to the blockchain in subsequent blocks, and the resolution of the tacit coordination game stemming from these investments will determine who profits and who loses.

Making more radical changes to the blockchain would be even more difficult. To dislodge a block 10 blocks from the end of the blockchain, 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 blockchain 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 blockchain than newer blocks. In this sense, Autonocoin is like Bitcoin. A particular block in the Bitcoin blockchain may eventually be removed from it, for example because another proposed block is added for the same spot in the blockchain at around the same time. (In this case, the authoritative blockchain 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 blockchain 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 blockchain and what to leave off and can even remove some blocks from the blockchain, replacing them over time with more blocks than everyone else can

ISSN 2379-5980 (online)
DOI 10.5195/LEDGER.2016.37

add to the valid blockchain 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 an attack by a government whose goal is to destroy the cryptocurrency rather than to profit from it.

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 blockchain 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, the interests of miners may have a significant and perhaps even dominant effect on Bitcoin policy, 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 blockchain, place the transactions performing this spending on another blockchain and present this as the valid one. According to the rules above, the purported double-spends 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 blockchain, possibly allowing an attack with less than 50% ownership of cryptocurrency potentially available for proof-of-belief transactions. 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 blockchain. 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.

ISSN 2379-5980 (online)
DOI 10.5195/LEDGER.2016.37

4. Rewarding Activities Benefiting the Cryptocurrency

An additional benefit 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.[12] One might defend this decision on the ground that 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 the originating 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 blockchain 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 very large businesses that adopted the cryptocurrency might receive larger rewards than smaller businesses. 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 blockchain, 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.[13] This

ISSN 2379-5980 (online)
DOI 10.5195/LEDGER.2016.37

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 1,000 Autonocoins. 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 1,000 Autonocoins 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.

5. Conclusion

Open-source cryptocurrencies to date have been governed by consensus. Changes can be made to the source code defining the cryptocurrency protocol only when there is sufficient confidence that others will also accept the change. The danger that disagreement can lead to a fork can thus discourage change. It is possible, however, to build governance mechanisms directly into a cryptocurrency protocol, permitting even controversial issues to be resolved in a way that is generally viewed as authoritative. The hypothetical Autonocoin cryptocurrency described here would performance governance through formal tacit coordination games. This governance capability could enable the cryptocurrency to include a built-in system for rewarding contributors to the cryptocurrency in proportion to their contributions.

Acknowledgment

The author is indebted to three anonymous referees. All errors are the author's own.

Notes and References

[1] Abramowicz, M. "Cryptocurrency-Based Law." *Arizona Law Review* **58.2** 359-420 (2016)

[2] Schelling, T. C. *The Strategy of Conflict*. Cambridge: Harvard Univ. Press 55-56 (1981)

[3] Kroll, J. A., Davey, I. C., Felten, E. W. "The Economics of Bitcoin Mining or, Bitcoin in the Presence of Adversaries." *Proceedings of WEIS* 1-21 (2013)

[4] Nyman, L., Lindman, J. "Code Forking, Governance, and Sustainability in Open Source Software." *Technology Innovation Management Review* (Jan. 2013)

[5] Hertig, A. "A Controversial Bitcoin Alternative is Seeking a Comeback." *Coindesk* (accessed 26 September 2016) `http://www.coindesk.com/controversial-bitcoin-alternative-seeking-comeback/`

[6] Keynes, J. M. *General Theory of Employment, Interest, and Money*. London: Palgrave Macmillan (1936).

[7] Hosni, H. "Interpretation, Coordination and Conformity." In O. Majer (Ed.), *Games: Unifying Logic, Language and Philosophy* **15** 37-55 (2009)

[8] Ongaro, D., Ousterhout, J. "In Search of an Understandable Consensus Algorithm." No Publisher (May 20, 2014) `https://raft.github.io/raft.pdf`

[9] King, S., Nadal, S. "PPCoin: Peer-to-Peer Crypto-Currency with Proof-of-Stake." No Publisher (Aug. 19, 2012) `https://peercoin.net/assets/paper/peercoin-paper.pdf`

ISSN 2379-5980 (online)
DOI 10.5195/LEDGER.2016.37

[10] Shwartz, D., Youngs, N., Britto, A. "The Ripple Protocol Consensus Algorithm." No Publisher (2014) https://ripple.com/consensus-whitepaper/

[11] Higgins, S. "Stellar Network Fork Prompts Concerns over Ripple Consensus Protocol." *Coindesk* (accessed 8 December 2016) http://www.coindesk.com/stability-questions-dog-ripple-protocol-stellar-fork/

[12] No Author. "Ripple (payment protocol)." *Wikipedia* (accessed 8 December 2016) https://en.wikipedia.org/wiki/Ripple_(payment_protocol)#Reception

[13] No Author. "Auroracoin." *Wikipedia* (accessed 8 December 2016) https://en.wikipedia.org/wiki/Auroracoin

ISSN 2379-5980 (online)
DOI 10.5195/LEDGER.2016.37