

Matthew E. Gladden

Georgetown University, Washington, DC

Polish Academy of Science, Institute of Computer Science, Warsaw

e-mail: matthew.e.gladden@gmail.com

Cryptocurrency with a Conscience: Using Artificial Intelligence to Develop Money that Advances Human Ethical Values

Abstract

Cryptocurrencies like Bitcoin are offering new avenues for economic empowerment to individuals around the world. However, they also provide a powerful tool that facilitates criminal activities such as human trafficking and illegal weapons sales that cause great harm to individuals and communities. Cryptocurrency advocates have argued that the ethical dimensions of cryptocurrency are not qualitatively new, insofar as money has always been understood as a passive instrument that lacks ethical values and can be used for good or ill purposes. In this paper, we challenge such a presumption that money must be ‘value-neutral.’ Building on advances in artificial intelligence, cryptography, and machine ethics, we argue that it is possible to design artificially intelligent cryptocurrencies that are not ethically neutral but which autonomously regulate their own use in a way that reflects the ethical values of particular human beings – or even entire human societies. We propose a technological framework for such cryptocurrencies and then analyse the legal, ethical, and economic implications of their use. Finally, we suggest that the development of cryptocurrencies possessing ethical as well as monetary value can provide human beings with a new economic means of positively influencing the ethos and values of their societies.

Keywords: cryptocurrency, socially responsible investing, payment systems, business ethics, economic cybernetics

JEL Classification: A13, D14, E22, G21, O33

1. Introduction: must our money always be value-free?

The speed, efficiency, anonymity, and global reach of cryptocurrencies like Bitcoin and Litecoin are offering creative new avenues for economic empowerment to entrepreneurs who might not otherwise be able to participate meaningfully in the global economy. However, the potential of such decentralized digital currencies to contribute to the common good has thus far been frequently overshadowed by cryptocurrency's destructive 'dark side,' including its use for facilitating human trafficking, paid assassinations, child exploitation, illicit weapons and drug sales, terrorism, identity theft, and other crimes that cause tremendous harm to some of the world's most vulnerable populations.¹ For cryptocurrency's proponents, one response has been to note that the potential for misuse is not unique to cryptocurrency: it has been true of all forms of currency throughout history that they could be used by in ways that directly harm others – or, perhaps more frequently, for purposes that are simply frivolous, wasteful, or selfish. Money neither realizes nor protests when it is put to such uses.

However, one might ask whether there is indeed some theoretical or practical necessity which requires that money serve as an empty cipher in such moral equations – or whether the moral inertness and neutrality that have traditionally been attributed to money are more of a historical accident than an essential characteristic inherent in the concept of money. Is it possible to imagine that money itself might be able to manifest its own set of ethical values – or even possess its own 'conscience'? While that might at first glance appear far-fetched, in this text we argue that ongoing and anticipated developments in fields such as artificial intelligence, cryptography, neurocybernetics, and quantum computing may provide humanity with a practical opportunity (and perhaps even imperative) for undertaking such a radical reconceptualization of the nature of money and its potential roles within human society. More sophisticated future cryptocurrency technologies will bring with themselves the threat of negative social and political impacts that are different from (and perhaps more far-reaching than) those produced by Bitcoin – but at the same time, it may be possible for humanity to harness such next-generation technologies to develop more advanced cryptocurrencies that not only possess some *monetary* value, but which literally embody the best shared *ethical* values of humankind.

¹ L.J. Trautman, *Virtual Currencies; Bitcoin & What Now after Liberty Reserve, Silk Road, and Mt. Gox?*, "Richmond Journal of Law and Technology" 2014, Vol. 20, No. 4, <http://jolt.richmond.edu/v20i4/article13.pdf> (accessed April 29, 2015).

2. Developing a framework for cryptocurrency that possesses ethical values

2.1 Converging technologies that lay the foundation for autonomous ethically guided cryptocurrencies (AEGCs)

As the starting point for our consideration of these issues, we can define an ‘autonomous ethically guided cryptocurrency’ (or AEGC) as a piece of software that combines two distinct features: (1) it possesses a type of artificial intelligence that senses its environment, gathers and analyses data, and autonomously makes decisions guided by particular ethical principles;² and (2) it functions as a digital cryptocurrency that can serve as a medium of exchange, unit of account, and store of value. In this text, we will focus primarily on the social, political, and economic implications of such technology; however, we will begin by reviewing the ongoing technological advances that will render the development of such cryptocurrencies possible, as well as the general technical specifications that such a cryptocurrency might possess. Scientific and technological advances which together provide the context for the development of AEGCs include:

- (1) **Expansion and refinement of cryptocurrency systems.** Building on the model of Bitcoin (and our understanding of its technical strengths and weaknesses that becomes more advanced as Bitcoin is tested through real-world use), computer science researchers and cryptocurrency advocates are developing next-generation cryptocurrencies that utilize more sophisticated technologies and are based on more effective economic models.³
- (2) **Development of AI capable of directing financial and economic activity.** While contemporary cryptocurrencies do not yet possess their own artificially intelligent software that could proactively guide their participation in economic exchange, such software is already used extensively in other financial systems and processes. For example, although the volume of computer-initiated transactions on American stock exchanges has declined since the “Flash Crash” of 2010, it has been estimated that as of 2013, roughly half of all stock transactions occurring on US exchanges were initiated by firms utilizing automated high-frequency trading

² It is possible for an agent to be both autonomous and constrained by certain ethical principles in its choice of actions; a human being is an example of such an agent. For a discussion of autonomy in the context of AI, see: R. Murphy, *Introduction to AI Robotics*, The MIT Press, Cambridge, Massachusetts 2000, pp. 31–34.

³ For a discussion of the strengths and limitations of Bitcoin, along with suggestions for technical improvements that could be incorporated into future cryptocurrencies, see: S. Barber, X. Boyen, E. Shi, E. Uzun, *Bitter to Better – How to Make Bitcoin a Better Currency* [in:] *Financial Cryptography and Data Security*, Lecture Notes in Computer Science 7397, ed. A.D. Keromytis, Springer, Berlin/Heidelberg 2012, pp. 399–414, doi:10.1007/978-3-642-32946-3_29.

(HFT),⁴ and automated trading systems continue to grow more sophisticated: many such systems are capable of teaching themselves and improving their investment strategies over time, without instruction from human beings.⁵ Moreover, as of July 2014, automated ‘robo-advisors’ were directly managing investment portfolios with \$19 billion in assets, and that amount was increasing at a rate of roughly 100% per year.⁶ On a more pernicious note, criminal enterprises are developing increasingly sophisticated ‘ransomware’ that infects a user’s computer, encrypts the user’s files to render them inaccessible, demands the payment of a ransom via Bitcoin, and then decrypts the files after receipt of the ransom payment. Although illegal and morally deplorable, from a purely technological perspective such software represents a notable development, insofar as its AI is capable of autonomously interacting with human beings in the real-world economy to generate revenue for itself and its owners. The (relatively simple) AI that guides such computer worms’ activities of replication and extortion is opposed by the AI of security software that attempts to detect and counteract such threats, disrupting the ransomware’s intended processes of economic exchange.⁷

- (3) **Machine ethics for AIs.** Computer scientists and ethicists are working to develop AI software that is capable of making judgments about complex real-world situations based on ethical principles. In some of these models, an AI gathers and analyzes data from its environment and makes decisions based on a fixed set of ethical principles that have been programmed into it; in more sophisticated metavolitional systems, an AI may be capable of learning new ethical principles and developing its own ‘conscience’ through interacting with the world.⁸
- (4) **Quantum computing, AI, and neural-network cryptography.** Research toward developing quantum computers is well underway, with ongoing advances expected in the field.⁹ If developed sufficiently to allow

⁴ M. Philips, *How the Robots Lost: High-Frequency Trading’s Rise and Fall*, “BloombergView” 2012, June 6, <http://www.bloomberg.com/bw/articles/2013-06-06/how-the-robots-lost-high-frequency-trading-rise-and-fall> (accessed April 29, 2015).

⁵ G. Scopino, *Do Automated Trading Systems Dream of Manipulating the Price of Futures Contracts? Policing Markets for Improper Trading Practices by Algorithmic Robots*, “Florida Law Review” 2015, Vol. 67, pp. 221–293.

⁶ S. Sharf, *Can Robo-Advisors Survive A Bear Market?*, “Forbes” 2015, January 28, <http://www.forbes.com/sites/samanthasharf/2015/01/28/can-robo-advisors-survive-a-bear-market/> (accessed April 29, 2015).

⁷ H.V. Nath, B.M. Mehre, *Static Malware Analysis Using Machine Learning Methods* [in:] *Recent Trends in Computer Networks and Distributed Systems Security*, Communications in Computer and Information Science 420, eds. G. Martínez Pérez, S.M. Thampi, R. Ko, L. Shu, Springer, Berlin Heidelberg 2014, pp. 440–50.

⁸ See: W. Wallach, *Robot Minds and Human Ethics: The Need for a Comprehensive Model of Moral Decision Making*, “Ethics and Information Technology” 2010, Vol. 12, No. 3, pp. 243–50, doi:10.1007/s10676-010-9232-8; D.J. Calverley, *Imagining a Non-Biological Machine as a Legal Person*, “AI & SOCIETY” 2008, Vol. 22, No. 4, pp. 523–37.

⁹ J.J. Pla, K.Y. Tan, J.P. Dehollain, W.H. Lim, J.J.L. Morton, D.N. Jamieson, A.S. Dzurak, A. Morello, *A Single-Atom Electron Spin Qubit in Silicon*, “Nature” 2012, Vol. 489, No. 7417, pp. 541–45, doi:10.1038/nature11449.

for the practical factoring of large prime numbers, quantum computing has the potential to render obsolete popular encryption systems such as RSA that are used to secure countless forms and quantities of information around the world.¹⁰ This possibility is already driving the development of new theoretical and technological bases for ‘post-quantum cryptography’ that could be used to secure digital data including cryptocurrency. One possible avenue for securing data that does not rely on the transmission of digital keys could be to embed the data within an artificially intelligent software program that relies on internally stored or generated criteria to decide when and to whom to make its data accessible. In particular, by embodying the AI as a physical artificial neural network, it might be possible to store such secured data and access criteria in a form that is available to the AI’s internal cognitive processes but which cannot be extracted or interpreted by external agents.¹¹

- (5) **Neuroprosthetics and cognitively based biometrics.** An alternative future approach to controlling access to digitized information (including units of cryptocurrency) might utilize biometric neuroprosthetic devices that are capable of recognizing cognitive processes or activity, such as thoughts or volitions, manifested within the brains of individual human beings who are authorized to access the data.¹² More particularly, a mnemonic cybernetic approach¹³ could utilize the memory-storage mechanisms of the human brain (rather than those of a physical artificial neural network) for biometric authentication.
- (6) **Artificial life-forms as economic actors.** Researchers are developing cybernetic frameworks for AIs that utilize the capacities described above to function as autonomous goal- and policy-driven agents that compete against human workers and businesses in the real-world economy. In the future, such agents may include artificial life-forms that dwell within physical-digital ecosystems in which they successfully secure environmental resources, produce goods and services, exchange these products

¹⁰ M. Heger, *Cryptographers Take on Quantum Computers*, “IEEE Spectrum” 2009, January 1, <http://spectrum.ieee.org/computing/software/cryptographers-take-on-quantum-computers> (accessed April 29, 2015).

¹¹ This approach differs from proposed models of neural cryptography that rely on multiple neural networks that synchronize with one another through mutual learning; for a discussion of those models, see: E. Volna, M. Kotyrba, V. Kocian, M. Janosek, *Cryptography Based On Neural Network* [in:] *Proceedings of the 26th European Conference on Modelling and Simulation*, eds. K.G. Troitzsch, M. Möhring, U. Lotzmann, European Council for Modelling and Simulation 2012, pp. 386–91. The model that we propose here is based instead on holographic and holonomic models of memory storage in the human brain (see: K.H. Pribram, *Prolegomenon for a Holonomic Brain Theory* [in:] *Synergetics of Cognition*, Springer Series in Synergetics 45, eds. H. Haken, M. Stadler, Springer, Berlin/Heidelberg 1990, pp. 150–184) and the inaccessibility to external agents of data stored in such systems.

¹² R. Palaniappan, *Two-Stage Biometric Authentication Method Using Thought Activity Brain Waves*, “International Journal of Neural Systems” 2008, Vol. 18, No. 1, pp. 59–66, doi:10.1142/S0129065708001373.

¹³ M.E. Gladden, *Tachikomatic Domains: Utopian Cyberspace as a “Contingent Heaven” for Humans, Robots, and Hybrid Intelligences*, [conference presentation at:] *His Master’s Voice: Utopias and Dystopias in Audiovisual Culture*, Ośrodek Badawczy Facta Ficta, Uniwersytet Jagielloński, March 24, 2015.

for money or other resources that generate a net profit, and invest these surplus resources in a manner that allows the life-forms to grow and reproduce.¹⁴ Such ‘synthetic organism-enterprises’ may incorporate and utilize cryptocurrency-like systems as a key means for participating in economic exchange.

By combining these elements, one can conceptualize a cryptocurrency that possesses specialized AI software that analyzes the circumstances of a particular financial transaction in which the human owner of the unit of cryptocurrency would like to use it (either to spend or invest it) and then makes the decision to allow or block the unit of cryptocurrency’s use in that transaction. When determining whether to allow itself to be utilized in a transaction, the cryptocurrency’s AI does not focus on the transaction’s financial aspects but rather identifies and assesses its ethical context and decides whether the transaction is permissible or impermissible according to a particular set of ethical principles that have been explicitly programmed into the cryptocurrency by its human designer. Such a technology would move us beyond the realm of merely ‘smart data’ and into a new world of ‘sapient data.’

In its simplest form, such a cryptocurrency could be built on the foundation of an existing platform such as Bitcoin. In order to spend a particular Bitcoin, its owner must possess and transmit the private key belonging to that coin. Some owners store their private keys online in ‘digital wallets’; others save them to a hard drive or in hard copy. Alternatively, an owner could transfer his or her private keys to a trusted third party who would manage and invest those Bitcoins on the owner’s behalf. In our model, this ‘third-party manager’ is not another human being but an artificially intelligent piece of software. In this text we can refer to this amalgam of a unit of cryptocurrency and the ethically oriented AI software that controls it as a ‘ConsCoin.’ The AI software gains and maintains exclusive access to the private key of any new ConsCoins that is mined by its human owner’s computer; thus, even the human being who ‘owns’ the coin never directly knows or controls its private key. After it has been mined, the AI software asks the coin’s human owner to define the specific purposes for which that coin should or should not be used and to describe his or her aspirations for the kinds of change that he or she would like to bring about in the world. In other words, the coin’s initial human owner *defines that particular coin’s ethical values*.

Because the AI software permanently controls access to the coin’s financial value, the coin and the software overseeing it become permanently fused into a single coin-AI unit that is autonomous and possesses both financial value and ethical values. From that moment forward, the AI software will invest and manage the money in accordance with the ethical guidelines given by the human being responsible for mining the coin. In this ConsCoin model, once the coin’s initial human owner has assigned a set of ethical values to the coin, it can never be changed.

¹⁴ M.E. Gladden, *The Artificial Life-Form as Entrepreneur: Synthetic Organism-Enterprises and the Reconceptualization of Business* [in:] *Proceedings of the Fourteenth International Conference on the Synthesis and Simulation of Living Systems*, eds. H. Sayama, J. Rieffel, S. Risi, R. Doursat, H. Lipson, The MIT Press, Cambridge, Massachusetts 2014, pp. 417–18.

A subsequent owner of the coin could ask its AI software to spend the money on merchandise or services or to transfer it as a donation or for any other purpose, as long as this does not violate the ethical conditions that the owner initially placed on the money; should a ConsCoin's human owner ask the AI to utilize the coin for some purpose incompatible with its ethical rules, the AI would refuse to comply.

2.2 Financial and legal aspects of AEGCs: building on existing models of government incentivization of philanthropic activity

The notion that a person who owns some piece of property could voluntarily restrict the ways in which it can be used – and that these restrictions would be permanently binding on any future owners of the property – is already well-established in the legal concept of an easement.¹⁵ In a sense, the initial owner of a ConsCoin who places binding future ethical restrictions on its use is creating a sort of 'easement' on the coin. With conservation and historical preservation easements, it is a government agency or non-profit organization that has the legal authority to enforce the easement by taking legal action against future owners of the property who try to use it in a way that violates the intent of the easement. With a ConsCoin, enforcement is accomplished not through legal or political means but through purely technological ones. While human legal and political institutions may sometimes fail to enforce easements as a result of corruption, incompetence, or lack of resources, the strength of a ConsCoin's ethical enforcement mechanism is limited only by the ingenuity and sophistication of the computer programmers who design it.

By limiting the ways in which property can be used, the creation of a conservation or historical preservation easement typically reduces the property's market value. A similar phenomenon could be expected to occur with ConsCoins: in open circulation, a single ConsCoin that can be spent in any way would be more valuable than a single ConsCoin that, for example, cannot be spent on products that were produced by companies utilizing animal testing or coal-burning power plants. In general, the more ethical constraints an owner places on his or her ConsCoin, the more its value will be reduced; this would tend to make the creation of ConsCoins financially unattractive. However, there are at least two possible countervailing phenomena that could encourage the creation of ConsCoins with significant ethical restrictions.

One means of encouraging the creation of ConsCoins with significant ethical restrictions is for governments to acknowledge and promote the non-financial value that such currency produces for society by offering financial incentives to citizens who ethically constrain their ConsCoins. This would build on the practices already implemented by some governments of offering tax benefits to those who create conservation or historical preservation easements or who donate a portion of their annual income to charitable organizations that have been recognized by the government. One approach would be for governments to offer tax benefits for any coin created within a ConsCoin system that is sponsored by an officially recognized charitable organiza-

¹⁵ J.A. Gustanski, R.H. Squires, *Protecting the Land: Conservation Easements Past, Present, and Future*, Island Press, Washington, DC 2000.

tion. Under this model, a government might promulgate general guidelines for calculating the amount of the tax benefit based on the type and degree of ethical constraints placed on the coin, but it would be up to the sponsoring organization to work out the details of their ConsCoin system. This would allow a government to broadly promote the creation of ethically conscious money that advances the common good, while leaving specific decisions about the values and aims to be promoted to other sectors of civil society.¹⁶

2.3 Sociopolitical aspects of AEGCs: creating communities of economic solidarity

A ConsCoin that possesses significant ethical restrictions might, in principle, be less ‘useful’ than one that possesses no ethical restrictions and which can thus be invested or spent in any way. However, in practice, human nature – as reflected in the psychological, social, and political behaviours of the human beings who would utilize this currency – might cause ConsCoins with certain kinds of ethical constraints to be subject to greater demand (and possess a higher market price) than those with no constraints at all. For example, imagine that an animal-welfare organization has produced a limited number of ConsCoins that will not allow themselves to be spent to purchase meat products or any cosmetics that were produced using animal testing. If such ‘ZooCoins’ were made available for purchase on an open market and the only relevant consideration were the currency’s flexibility for use in making purchases, the market might determine that a single ZooCoin possesses, say, only 96% of the value of a ConsCoin that has no ethical constraints. However, individuals who are deeply committed to animal welfare might conceivably prefer to own ZooCoins over ConsCoins with no ethical constraints, because either: (1) they sincerely believe that by adding to the economy money that can never be used to harm animals, they are working to bring about a better world; or (2) by possessing and spending currency that can never be used to harm animals, they reinforce their own self-understanding of themselves as ethical individuals and they publicly identify themselves with a particular lifestyle and ethical mindset. For those who consider themselves to be supporters of animal welfare, the use of ZooCoins could be both a public symbol of their commitment to that cause and a tangible means of advancing its goals. This heightened demand for ZooCoins on the part of particular groups could potentially raise its market price to match that of an unconstrained ConsCoin; even if this does not occur, an animal-welfare proponent might still voluntarily use an unconstrained ConsCoin to purchase a ZooCoin with the same face value but a lower market price, viewing the financial loss as a sort of ‘charitable contribution’ made to advance the cause of animal welfare.

As long as a particular ConsCoin is only exchanged among people who already share its ethical commitments, the coin’s ethical constraints do not represent a loss of utility; thus particular ConsCoins would likely possess the greatest financial

¹⁶ For examples of creative ways in which contemporary governments can incentivize and facilitate charitable activity within a capitalist economy, see L. McGoey, *The Philanthropic State: Market-state Hybrids in the Philanthrocapitalist Turn*, “Third World Quarterly” 2014, Vol. 35, No. 1, pp. 109–25, doi:10.1080/01436597.2014.868989.

value when exchanged among individuals who share the currency's ethical commitments. This phenomenon could spur the growth of new communities of economic solidarity, perhaps manifested through online marketplaces allowing individuals to automatically seek out producers and consumers who prefer to trade in ConsCoins expressing their shared ethical values. Such communities could potentially foster innovative new forms of microenterprise, building on the model of successful web-based microcredit providers like Zidisha and Kiva Microfunds.¹⁷

If the use of ethically infused ConsCoins were ever to reach a 'critical mass' of popularity within a particular society, those persons who continued to use ethically unconstrained currencies might come to be viewed with suspicion, as individuals who lack concern for the common good and who might be seeking to retain an ability to engage in harmful or illicit activities. Over time, such social pressures might lead ever more individuals throughout a society to assign their ConsCoins at least a certain minimum level of ethical constraints that reflect the society's most basic and widely shared ethical commitments.

2.4 AEGCs as the ultimate form of socially responsible investing

Taking the lead from socially conscious cryptocurrencies like MazaCoin¹⁸ and Dogecoin,¹⁹ ConsCoins could build on and accelerate the trend toward socially responsible investing (SRI) on the part of individual investors, advocacy groups, and pension funds, by combining SRI's techniques of negative screening and positive investing. For example, a person with a deep concern for the environment might instruct his or her newly mined ConsCoin that it: (1) should never allow itself to be invested in industries such as mining, oil refining, or cement production that have a disproportionately large negative environmental impact; and that (2) it should invest itself in environmentally beneficial industries such as renewable energy or eco-friendly construction, whenever it can do so in a way that meets specified criteria for financial risk and return.

2.5 Risks of corruption and abuse within the AEGC model

While the ConsCoin technology described above could be used to promote the global common good, it could also potentially be misused for destructive ends. For example, the corrupt authoritarian political leaders of a particular country could use ConsCoin technology to place permanent restrictions on new cryptocurrency issued

¹⁷ L. Gilpin, *How Zidisha Sidestepped Banks and Took Microfinancing Peer-to-Peer*, "TechRepublic" 2014, May 1, <http://www.techrepublic.com/article/how-zidisha-sidestepped-banks-and-took-microfinancing-peer-to-peer/> (accessed September 1, 2014).

¹⁸ L. Browning, *Oglala Sioux Hope Bitcoin Alternative, Mazacoin, Will Change Economic Woes*, "Newsweek" 2014, August 14, <http://www.newsweek.com/2014/08/22/tribe-brought-you-custers-last-stand-sitting-bulls-bitcoin-264440.html> (accessed September 1, 2014).

¹⁹ D. Gilbert, *"Most Valuable Tweet in History" Donates \$11,000 Worth of Dogecoin to Kenyan Water Charity*, "International Business Times UK" 2014, March 17, <http://www.ibtimes.co.uk/most-valuable-tweet-history-donates-11000-worth-dogecoin-kenyan-water-charity-1440565> (accessed September 1, 2014).

by their government, so that it could only be used to advance their own personal interests. If other nations were to outlaw the exchange of such currency (in a manner similar to the international bans on the sale of South African Krugerrands in the 1980s²⁰), this could make the issuance of such currency less attractive to its potential creators. However, it is unclear whether even the most concerted international political and economic pressure would be sufficient to eliminate all such occurrences. One means of avoiding such abuses could be to develop ConsCoin systems that draw their ethical values from broad segments of human society, as we shall describe in the following section.

3. Cryptocurrencies whose ethical values reflect the collective conscience of an entire society

In the AEGC model described above, a ConsCoin's ethical constraints are permanently and unalterably stamped on the coin at the time of its creation and are informed solely by the ethical commitments of the person who mined the coin. (We might describe such a cryptocurrency's ethical commitments as being 'individually informed' by the values of its original owner.) In such a case, the AI governing the coin's future use would simply attempt to interpret and apply the rigid set of instructions that it had been given. Such a ConsCoin's AI would be 'volitional' but not 'metavolitional'; it would be incapable of learning and growing and seeking to deepen its ethical understanding over time. A more sophisticated form of ConsCoin might be designed with a form of AI that allows it to expand and improve its ethical awareness beyond whatever limited set of instructions a single human being might be able to program into it. In principle, a ConsCoin's AI could be given complete freedom to develop its own ethical stances based on all that it learns from its observation of and interaction with the world. However, this raises the possibility – subject to ongoing discussion among ethicists of robotics and artificial intelligence – that an AI might adopt an ethical framework which, while appearing morally superlative in the AI's own judgment, would be unfathomable to human reason and perhaps even inimical to the continued existence of humanity.²¹

A 'safer' approach might be to create a ConsCoin whose AI obtains its set of ethical values not from a single human being but from an entire society. Such a ConsCoin would actively seek to learn and understand the values of a human society and then forge its own moral principles based on the best of these human insights; it would not attempt to 'surpass' human morality but to distill its most fundamental and exemplary aspects. (We could describe the ethical values of this sort of ConsCoin system as 'socially' rather than 'individually' informed.) Rather

²⁰ M.P. Malloy, *Human Rights and Unintended Consequences: Empirical Analysis of International Economic Sanctions in Contemporary Practice*, "Boston University International Law Journal" 2013, Vol. 31, pp. 79–129.

²¹ L. Muehlhauser, L. Helm, *The Singularity and Machine Ethics* [in:] *Singularity Hypotheses*, eds. A.H. Eden, J.H. Moor, J.H. Søraker, E. Steinhart, Springer, Berlin/Heidelberg 2012, pp. 101–26.

than accepting ethical values stamped onto it by its miner, such a ConsCoin would work to ascertain the ethical values of the human society in which it exists and then adopt those values as its own.²² Using AI and data-mining techniques, such a ConsCoin could probe, assess, and interpret a society's ethical aspirations by analyzing phenomena such as the contents of social media posts, levels of participation in different religious or philosophical movements, political activism, voting patterns, spending patterns, levels and kinds of charitable donations and volunteer service, and the levels and kinds of audience engagement with different publications, films, music, artwork, and other creative activities.²³ If intended primarily to promote the welfare of a particular social group or geographic area, the ConsCoin might adopt the ethical stances shared by that human community; a ConsCoin designed for broader impact might attempt to incorporate the ethical values of an entire nation, or even to discern and adopt those essential ethical values and aspirations that are shared by all human beings around the world.

Such a ConsCoin would possess the ability to adapt as human ethical priorities shift over time. While the most important ethical principles possess a permanent and objective validity, the ways in which human societies prioritize and reconcile competing ethical demands can take on different forms in different times and places. For example, a society that is experiencing a devastating war or famine might be willing to temporarily adopt industrial or agricultural policies which in the long term would be environmentally detrimental and unsustainable, in order to provide for the most desperate and immediate needs of its people. Facing such a situation, a ConsCoin that is able to recognize and analyze these circumstances might relax its normal restrictions against being spent in 'environmentally detrimental' ways.

4. A new way for the public to shape public policy

The development of ConsCoin systems could potentially offer a new non-political, economic means whereby citizens could shape and improve the character and public policies of their societies. For example, a nation's political leaders might think twice before launching a new war – and might instead be inclined to seek out peaceful solutions to its international disputes – if it knew that a significant portion of the nation's liquid assets consisted of currency that would refuse to allow itself to become involved in a war. By reflecting the moral sentiments of the nation, the economic dynamics of a country's autonomous ethically guided cryptocurrencies could become a sort of democratic 'para-government' that acts as a check on (and conscience for) its political government. Moreover, if the individuals holding political,

²² Importantly, the ConsCoin would not attempt to replicate the ways in which human beings actually behave – with our frequent flaws and failings – but rather to discern our collective conscience, the ways in which we *wished* we behaved.

²³ M. Hilbert, *The Maturing Concept of E-Democracy: From E-Voting and Online Consultations to Democratic Value Out of Jumbled Online Chatter*, "Journal of Information Technology & Politics" 2009, Vol. 6, No. 2, pp. 87–110, doi:10.1080/19331680802715242.

military, and economic power within a country knew that much of the nation's wealth would cease to work for the government if the current peaceful, democratically elected government were overthrown in a coup – and that there would be no way to force this wealth to cooperate with the new regime, either through threats or extortion – this might conceivably discourage the overthrow of legitimate, peaceful governments and even enhance the world's long-term geopolitical stability.

5. Conclusion

Many of the individual technological components needed to create autonomous ethically guided cryptocurrency systems whose behaviours are informed by human ethical values are already being developed, although combining these technologies to create a successfully functioning 'cryptocurrency with a conscience' will require further focused innovation. Those who are rightly troubled by the negative social impacts generated by the rise of Bitcoin may be inclined to argue that the development of autonomous cryptocurrencies should not be pursued. However, we would suggest that further research into such possibilities on the part of responsible computer scientists, economists, ethicists, and regulators is indeed merited, for two reasons. First, such research will allow us to better understand and prepare to counteract the negative impacts that such autonomous cryptocurrencies could have if implemented by lone programmers or corrupt state actors who would attempt to utilize such cryptocurrencies for purposes of self-interest rather than the global common good. Second, and perhaps more significantly, autonomous ethically guided cryptocurrencies represent a powerful new tool that possesses the potential to aid humanity in achieving our most deeply held moral and ethical aspirations. As human beings participate in economic exchanges of goods, services, and money with one another, they would also be participating in an exchange of information about their ethical commitments and even propagating those commitments in themselves. Our hope is that through the development of such autonomous ethically guided cryptocurrencies, advances in fields like artificial intelligence can be harnessed in a way that does not undermine human ethical agency but instead aids us to strengthen the role of ethical values in our shared economic life.

References

- Barber S., X. Boyen, E. Shi, E. Uzun, *Bitter to Better – How to Make Bitcoin a Better Currency* [in:] *Financial Cryptography and Data Security*, Lecture Notes in Computer Science 7397, ed. A.D. Keromytis, Springer, Berlin/Heidelberg 2012, pp. 399–414, doi:10.1007/978-3-642-32946-3_29.
- Browning L., *Oglala Sioux Hope Bitcoin Alternative, Mazacoin, Will Change Economic Woes*, "Newsweek" 2014, August 14, <http://www.newsweek.com/2014/08/22/tribe-brought-you-custers-last-stand-sitting-bulls-bitcoin-264440.html>.

- Calverley D.J., *Imagining a Non-Biological Machine as a Legal Person*, "AI & SOCIETY" 2008, Vol. 22, No. 4, pp. 523–37.
- Gilbert D., "Most Valuable Tweet in History" Donates \$11,000 Worth of Dogecoin to Kenyan Water Charity, "International Business Times UK" 2014, March 17, <http://www.ibtimes.co.uk/most-valuable-tweet-history-donates-11000-worth-dogecoin-kenyan-water-charity-1440565>.
- Gilpin L., *How Zidisha Sidestepped Banks and Took Microfinancing Peer-to-Peer*, "TechRepublic" 2014, May 1, <http://www.techrepublic.com/article/how-zidisha-sidestepped-banks-and-took-microfinancing-peer-to-peer/>.
- Gladden M.E., *Tachikomantic Domains: Utopian Cyberspace as a "Contingent Heaven" for Humans, Robots, and Hybrid Intelligences*, [conference presentation at:] *His Master's Voice: Utopias and Dystopias in Audiovisual Culture*, Ósrodek Badawczy Facta Ficta, Uniwersytet Jagielloński, March 24, 2015.
- Gladden M.E., *The Artificial Life-Form as Entrepreneur: Synthetic Organism-Enterprises and the Reconceptualization of Business* [in:] Proceedings of the Fourteenth International Conference on the Synthesis and Simulation of Living Systems, eds. H. Sayama, J. Rieffel, S. Risi, R. Doursat, H. Lipson, The MIT Press, Cambridge, Massachusetts 2014, pp. 417–18.
- Gustanski J.A., R.H. Squires, *Protecting the Land: Conservation Easements Past, Present, and Future*, Island Press, Washington, DC 2000.
- Heger M., *Cryptographers Take on Quantum Computers*, "IEEE Spectrum" 2009, January 1, <http://spectrum.ieee.org/computing/software/cryptographers-take-on-quantum-computers>.
- Hilbert M., *The Maturing Concept of E-Democracy: From E-Voting and Online Consultations to Democratic Value Out of Jumbled Online Chatter*, "Journal of Information Technology & Politics" 2009, Vol. 6, No. 2, pp. 87–110, doi:10.1080/19331680802715242.
- Malloy M.P., *Human Rights and Unintended Consequences: Empirical Analysis of International Economic Sanctions in Contemporary Practice*, "Boston University International Law Journal" 2013, Vol. 31, pp. 79–129.
- McGoey L., *The Philanthropic State: Market-state Hybrids in the Philanthrocapitalist Turn*, "Third World Quarterly" 2014, Vol. 35, No. 1, pp. 109–25, doi:10.1080/01436597.2014.868989.
- Muehlhauser L., L. Helm, *The Singularity and Machine Ethics* [in:] *Singularity Hypotheses*, eds. A.H. Eden, J.H. Moor, J.H. Søraker, E. Steinhart, Springer, Berlin/Heidelberg 2012, pp. 101–26.
- Murphy R., *Introduction to AI Robotics*, The MIT Press, Cambridge, Massachusetts 2000, pp. 31–34.
- Nath H.V., B.M. Mehtre, *Static Malware Analysis Using Machine Learning Methods* [in:] *Recent Trends in Computer Networks and Distributed Systems Security, Communications in Computer and Information Science* 420, eds. G. Martínez Pérez, S.M. Thampi, R. Ko, L. Shu, Springer, Berlin Heidelberg 2014, pp. 440–50.
- Palaniappan R., *Two-Stage Biometric Authentication Method Using Thought Activity Brain Waves*, "International Journal of Neural Systems" 2008, Vol. 18, No. 1, pp. 59–66, doi:10.1142/S0129065708001373.

- Philips M., *How the Robots Lost: High-Frequency Trading's Rise and Fall*, "BloombergView" 2012, June 6, <http://www.bloomberg.com/bw/articles/2013-06-06/how-the-robots-lost-high-frequency-tradings-rise-and-fall>.
- Pla J.J., K.Y. Tan, J.P. Dehollain, W.H. Lim, J.J.L. Morton, D.N. Jamieson, A.S. Dzurak, A. Morello, *A Single-Atom Electron Spin Qubit in Silicon*, "Nature" 2012, Vol. 489, No. 7417, pp. 541–45, doi:10.1038/nature11449.
- Pribra K.H., *Prolegomenon for a Holonomic Brain Theory* [in:] *Synergetics of Cognition*, Springer Series in Synergetics 45, eds. H. Haken, M. Stadler, Springer, Berlin/Heidelberg 1990, pp. 150–84.
- Scopin G., *Do Automated Trading Systems Dream of Manipulating the Price of Futures Contracts? Policing Markets for Improper Trading Practices by Algorithmic Robots*, "Florida Law Review" 2015, Vol. 67, pp. 221–293.
- Sharf S., *Can Robo-Advisors Survive A Bear Market?*, "Forbes" 2015, January 28, <http://www.forbes.com/sites/samanthasharf/2015/01/28/can-robo-advisors-survive-a-bear-market/>.
- Trautan L.J., *Virtual Currencies; Bitcoin & What Now after Liberty Reserve, Silk Road, and Mt. Gox?*, "Richmond Journal of Law and Technology" 2014, Vol. 20, No. 4, <http://jolt.richmond.edu/v20i4/article13.pdf>.
- Volna E., M. Kotyrba, V. Kocian, M. Janosek, *Cryptography Based On Neural Network* [in:] *Proceedings of the 26th European Conference on Modelling and Simulation*, ed. K.G. Troitzsch, M. Möhring, U. Lotzmann, European Council for Modelling and Simulation 2012, pp. 386–91.
- Wallach W., *Robot Minds and Human Ethics: The Need for a Comprehensive Model of Moral Decision Making*, "Ethics and Information Technology" 2010, Vol. 12, No. 3, pp. 243–50, doi:10.1007/s10676-010-9232-8.