

Ethical Concerns and Opportunities in Binding Intelligent Systems and Blockchain Technology

Davide Calvaresi¹[0000-0001-9816-7439], Jean-Gabriel Piguet¹, Jean-Paul Calbimonte¹[0000-0002-0364-6945], Timotheus Kampik²[0000-0002-6458-2252], Amro Najjar³[0000-0001-7784-6176], Guillaume Gadek⁴, and Michael Schumacher¹[0000-0002-5123-5075]

¹ University of Applied Sciences and Arts Western Switzerland, Switzerland

² Umeå University, Sweden.

³ University of Luxembourg, Luxembourg.

⁴ Airbus Defence and Space, France

{name.surname}@hevs.ch, tkampik@cs.umu.se, amro.najjar@uni.lu,
guillaume.gadek@airbus.com

Abstract. Intelligent systems are becoming increasingly complex and pervade a broad range of application domains, including safety-critical systems such as e-health, finance, and energy management. Traditional approaches are no longer capable of addressing the demand for trust and transparency in these applications. Hence, the current decade is demanding intelligent systems to be autonomous, and in particular explainable, transparent, and trustworthy. To satisfy such requirements, and therefore to comply with the recent EU regulations in the matter (e.g., GDPR), intelligent systems (e.g., Multi-Agent Systems - MAS) and technologies enabling tamper-proof and distributed consensus (e.g., Blockchain Technology - BCT) are conveying into reconciling solutions. Recently, the empowerment of MAS with BCT (and the use of BCT themselves) has gained considerable momentum, raising challenges, and unveiling opportunities. However, several ethical concerns have yet to be faced. This paper elaborates on the entanglement among ethical and technological challenges while proposing and discussing approaches that address these emerging research opportunities.

Keywords: Blockchain, Multi-agent Systems, Ethics, Transparency, Trust

1 Introduction

After years spent in trying to delegate *automatable* tasks (e.g., mass-production industry), our contemporary societies are experiencing machine delegation of a growing amount of *strategic/intelligent* tasks [9], for which transparency and trustworthiness of the intelligent entities (e.g., agents) and their behaviors must be ensured. For example, in the healthcare domain (eHealth in particular), any actor (e.g., caregiver, insurance, pharmacy, and smart-devices) can be modeled as an agent with specific behaviors [11, 12, 17]. Usually, actors such as medical

doctors and nurses are traditionally considered cooperative and trustworthy. Nevertheless, in many settings, they can have conflicting or competing interests (e.g., insurance companies and privately owned healthcare organizations). Thus, the demand for transparent reputation management to promote (and possibly ensure) *correct* behaviors is imperative.

To enforce basic coordination rules among agents, the usage of ledgers-based technologies (e.g., blockchain) has been proposed as a relevant solution [5]. At first glance, the systems reconciling BCT and MAS offer new perspectives for the empowerment of the individuals.

Currently, the users have to interact with Trusted Third Parties (TTP) unavoidably (e.g., platforms such as Airbnb and Uber), which are far from being egalitarian. Conversely, it is broadly accepted that BCT eradicates the monopoly on the information and ensures a certain degree of equality in control of the contracts' implementation. Thanks to the use of BCTs in MAS, the informational and executive asymmetry that underpins the business model of many platforms could come to an end. However, can the BCT actually deliver this massive societal promise? For systems that are based both on Multi-Agent Systems (MAS) and BlockChain Technology (BCT), and are operating in safety- and information-critical domains, a range of challenges and related opportunities is emerging. In particular, four fundamental questions should find an answer: *(i)* How can we prevent the ecosystem of technologies like BitCoins from being dominated by the wealthy few? *(ii)* In a MAS, when a “bad” decision is made, who is liable for it? *(iii)* Can we truly get rid of every intermediary, especially in the case of smart-contracts? and *(iv)* How to construct a fair reputation-building system through BCT and MAS, which could avoid setting a user’s undue reputation in stone?

This article *(i)* reviews recent trans-disciplinary work across computer science, legislation, and ethics; and *(ii)* proposes a structured synthesis of the current technological and ethical challenges of BCT and MAS, detailing the corresponding research opportunities. Solving these challenges may help BCT to comply with its promises and effectively empower multi-agent systems and their users.

The rest of the paper is organized as follow: Section 2 introduces background on blockchain and its intersection with MAS, Section 3 identifies challenges in this scope, and Section 4 discusses research opportunities to explore in this domain. Finally, Section 5 concludes the paper.

2 Blockchain and MAS: Transparency, Trust, and Ethics

To cope with the increasingly sophisticated AI systems deployed in our everyday lives, many recent initiatives called for explainable [21], transparent [3], and responsible AI systems [15]. Moreover, works such as [31] advocated the integration of implicit or explicit reasoning about ethics for such intelligent systems.

In MAS, which include complex, numerous, and fast negotiations, the complete visibility of the secured history of the transactions can be crucial (e.g.,

in the context of legal disputes). For these reasons, *trust* and *reputation* have been highlighted as two key factors in artificial agent societies [35, 36]. Nevertheless, very few works go beyond theoretical models to actually deploy trust and reputation-based mechanisms in agent societies. With the emergence of BCT, trust and reputation of different agents can be intertwined with distributed ledger technologies, thus accessible by all the stakeholders. Moreover, several works have recently underlined the need for **accountability** in MAS organizations [6]. Thanks to the use of BCTs in MAS, the informational and executive asymmetry that underpins the business model of many platforms can come to an end.

BCT eradicates the monopoly on the information and ensures a certain degree of equality on the control of the contracts' implementation. Nevertheless, recent studies allowed the conclusion that such features offered by BCT are "necessary" but not "satisfactory" conditions to meet the collective empowerment of users fully. Indeed, two limits to the empowerment strategy can be mentioned: *(i)* how should a user know that another would be ready to pass a contract with the need for a reminder (e.g., a virtual assistant)? An intermediary might still be needed to match supply and demand. Thus, the end of informational asymmetries and monopolies cannot be expected only from BCT. Possible inequalities can be overcome once users (people or agents) have achieved an agreement; *(ii)* although smart-contract-based technologies prevent malicious and erroneous agreements' implementations, they do not *eliminate the risks of mistakes*. Indeed, users are demanding that the decisions involving them or taken on their behalf (as a contracting party) have to be explainable. Henceforth, transparency without explainability is no longer acceptable [4, 13].

Among the various blockchain applications deployed in the market, their use as a monetary platform turned out to be problematic [14]. Focusing mainly on BitCoin, renowned economists are doubtful of the contributions of BCT to social good [27, 38]. From a practical point of view, money laundering and illegal transactions (drugs, murders, human trafficking) may use BCT to realize the payments [16]. From a more theoretical point of view, if BitCoin (or similar) were to be really used at a large scale, the role of the regulatory authorities could not be ensured anymore, which let a clear path to generalized fraud (and risk for the services provided by the states).

Finally, before the ethic stands the law: some popular uses of BCT represent an infraction. The most latent contravenes the European General Data Protection Regulation (GDPR), even though the complex nature of distributed fund transfer networks makes it hard to identify the actors, responsibilities, and acts⁵. Beyond the specific features of BitCoin, the unique identifiers of every wallet *is* to be considered as personal data; and so it should be stored and protected as such (which is incompatible with the nature of the network) [10].

⁵ <https://medium.com/@kkarnapp/what-does-the-gdpr-mean-for-bitcoin-9b57ebdd8766>

3 Challenges

Since its inception, blockchain has been typically associated with norms and values like democracy, transparency, and decentralization of power. However, the socio-technical properties of blockchain technologies as an economic system bring with them a set of ethical challenges that are diametrical to the idealistic moral propositions associated with blockchain [7].

C1 - *Erosion of the socio-cultural blockchain ecosystem.*

While blockchain was initially an ideology-driven technology, mostly supported by left-leaning or libertarian technology enthusiasts, the current domination of the ecosystem by “Bitcoin millionaires” has turned the blockchain community in a highly unequal society (Estimated Gini coefficients consider North Korea as the country with the most unequal distribution in the world⁶ [26]). The dominance of the ecosystem by a wealthy few, who can easily take control of emerging cryptocurrencies has rendered the vision of an egalitarian ecosystem impossible.

C2 - *Technical immutability of smart contracts in enterprise software contexts.*

Usually, the roll-out of new enterprise software solutions is associated with organizational stress that negatively influences the well-being of employees. However, the possibility of a controlled “roll-back” of updates and new implementations is typically possible and hence provides a safety net for worst-case scenarios. Moreover, in modern continuous integration scenarios, small improvements and ad-hoc bug fixing (so-called “hot fixing” can be deployed in production systems almost instantaneously. Given the technical immutability of smart contracts, and the socio-technical fact that finding cross-organizational consensus for new smart contract deployments is laborious. This becomes clear when looking at the prevalence of organizational conflicts and communication issues that hamper the *intra-organizational* implementation and adoption of IT systems [40], as well as at the slow pace at which multi-organizational technology-oriented negotiations in standardization processes move forward. These safety nets are weaker in BCT-based (typically consortium blockchain-based) enterprise system deployments, which can facilitate intra-organizational stress, inefficiencies, and operational mismanagement. New solution approaches are required to address this shortcoming of blockchain-based systems.

C3 - *Smart contract complexity as Blockchain Fraud.*

Fraud, overstated capabilities, and unfulfilled promises are obstacles hindering the evolution of public blockchains from being a niche phenomenon to a broadly applicable technology. Moreover, the advent of highly complex smart contracts (e.g., executing machine learning models and agent-based simulations) can exacerbate even more the already challenging path of BCT [24]. The deployment of such smart contracts will further increase the power imbalance between smart contract developers/providers and users. In particular, malicious smart contract providers can use complex models with *hard-to-foresee* emergent properties to

⁶ <https://bitcointalk.org/index.php?topic=51011.0>

perform behaviors that their users do not expect when they commit stakes (money or goods) to a blockchain-based system [23]. Finally, although the initial intention and functionality of a smart contract are beneficial for society, complex chain code can learn or evolves in unintended/unexpected directions (given the employment of BCT in fast-paced evolving environments). Unintended side-effects can facilitate system exploitation by malicious third parties (both at design- and run-time). Thus, to deliver an actual societal benefit, requirements elicitation, rules enforcement, norms definition, and laws revision have to face crucial changes.

C4 - *Theoretical and effective empowerment: from accessibility to workability.*

Merging BCT and MAS showed the possibility to implement *secure* behaviors and to provide the availability of *shared data* among the participating agents. Nevertheless, the dynamics in these environments are quite complex and sometimes relying on workarounds — thus, hindering an effective utilization of these data. Indeed, current prototypes and proofs-of-concept require extensive knowledge and support regarding legal aspects and data science. This limitation constitutes a barrier for the adoption of BCT-based agent systems as a viable means for user empowerment.

C5 - *The persistent need for intermediaries to implement smart contracts.*

The hope of getting rid of intermediaries through the use of smart contracts is associated with the specific challenge “to adapt the system and the contracting parties’ liability accordingly”. Such a challenge is not solely related to BCT. Indeed, it also extends to those MAS relying on smart contract technologies, where part of the implementation of a contract is delegated to agents or AI algorithms in general.

The possible risks connected with such a delegation are *(i)* lousy implementation of a contract through AI and *(ii)* even if no glaring mistake seems to have occurred, the terms of a contract may leave room for interpretations, i.e., requiring further consensus procedures between the contracting parties.

In both cases, there will be no platform liable for possible mistakes. Moreover, can we truly believe that the two contractual parties are liable for a decision that has been delegated?

C6 - *Liability in AI delegation.*

The advancements in AI-based decision-making engines allowed an increased degree of machine-delegation. However, in domains such as stock-negotiations and forensics analysis, to outline the liability boundaries amount the parties is increasingly complex. Assuming that human contracting parties let their virtual assistants (agents) interact on their behalf, how can *trust* and *liability* be unequivocally determined and associated with the actors? (e.g., single virtual agent, overall system, the system designer, and the human user).

C7 - Constructing a fair reputation-building system through BCT and MAS.

Several MAS rely on the ability to establish a mutual trust achieved via secure reputation building mechanisms and to ensure that no one has been over- or under-evaluated. Mutual trust demands at least for (i) the authenticity of the reputation’s origin, (ii) its traceability, and (iii) accuracy. Thanks to (i) and (ii), a given agent A should assume the correctness of B ’s remarks about C .

Certainly, BCT supports the satisfaction of (i) and (ii) [22]. The public storage of information ensures traceability and the possibility for each user to see if and about what precisely two actors agree or disagree with. The authenticity of an assessment is better warranted by the possibility of identifying the contributor (when using a non-anonymous permissioned BCT). In addition, it allows building a two-level assessment in a MAS: the evaluation of a given transaction, and a more general assessment of someone’s reputation, resulting in the previous local evaluations provided by his peers.

Ethical concerns arise concerning (iii). Does the advocated reputation management strengthen, in the absence of guarantees, the overall accuracy of the assessments? The epistemological principle that underlies the model is that transparency goes hand in hand with a deeper sense of responsibility for truth and accuracy. Unfair and malicious behaviors are supposed to be detected over time, leading to the exclusion of the non-objective member from the group.

Nevertheless, to what extent is this principle *fair*? Abramova et al. [1] proposed a study based on Airbnb showing that people read the comments and try to distinguish between fair and unfair criticisms, while Mayzlin et al. [30] proved that fraud is not a big matter in peer-to-peer platforms. Nevertheless, other studies mitigated the “collective wisdom” hypothesis. According to Origgi and Pais [33], whenever the cover of anonymity is removed, users might be reluctant to give negative evaluations. Such a behavior can be due to fear of retaliation [2, 34], to avoid conducts [20], and due to concerns regarding legal consequences [20]. Nevertheless, literature shows a certain number of biased reporting: herding behavior, self-selection [39], and strategic manipulation of reviewing [28]. Such social aspects can remarkably impact the accuracy and trustworthiness of the data stored on the shared ledgers.

Moreover, if an actor debuts in a community with a bad reputation rating (either if deserved, unfortunate, or undeserved), its near and possibly entire future might be affected (possibly creating unjust harms). Although these concerns are not solely and explicitly linked to bridging BCT and MAS, to elaborate on the ethics of rating dynamics is unavoidable.

C8 - Right to be forgotten

To pursue the principle of data minimization and limit potential abuses of personal data, GDPR aims at creating the right for the erasure of personal data that are “no longer necessary in relation to the purposes for which they were collected or otherwise processed” and if the data subject withdraws the consent for processing them (art. 17). In the case of BCT, compliance with article 17 raises at least two main issues. Above all, “since blockchain is a form of decentralized

transactions, it is questionable to whom the GDPR’s requirements of upholding a standard for personal data usage is addressed (...). Clearly, the GDPR has set a legal framework primarily intended for centralized personal data collectors. Thus, to scale it to distributed ledger technologies based on the concept of a completely decentralized environment (especially public BCT) is having unpredictable consequences and interpretations [41]. Moreover, data persistence is a fundamental feature of BCT. It seems almost impossible to guarantee the right to be forgotten without losing all the benefits inherited by this underlying technology (transparency, authenticity, and security). As indicated by the Open Data Institute, “the irreversibility and transparency of public BCT mean they are probably unsuitable for personal data” [37]. Finally, besides the general hype, the suitability of certain application domains for distributed ledger technologies should be adequately reassessed.

4 Research Opportunities

This section elaborates on the research opportunities standing beyond the challenges discussed above. Concerning **C1**, there is an emerging third era of blockchain that tangibly links the digital world with the physical world. Nevertheless, measuring the wealth distribution in cryptocurrency is quite challenging, mainly due to the anonymous nature of transactions (sender and receiver data are inaccessible). The New Bitcoin Distribution [8] models the sharing of Bitcoin Wealth in such a way that Wallet and Address data are entirely disregarded, assuming that the Power Law applies to Bitcoin Wealth Distribution and Bitcoin Wealth Distribution exactly mirrors the Global Wealth. Another solution is GoodDollar [32], designed with a mission to end economic inequality for developing a new cryptocurrency on a global scale and an open-source to distribute money through the principles of a Universal Basic Income (UBI).

Concerning **C2**, there are two possible approaches. On the one hand, in specific circumstances, being able to switch between blockchain-based and non-blockchain mechanism dynamically can improve the system flexibility. In particular, when organizational dynamicity becomes predominant (w.r.t. the need for following specific procedures specified in the chain code/smart contract), allowing a (semi)autonomous switch can be crucial. On the other hand, smart contracts can be designed to be autonomous, possibly adjusting and evolving their behavior over time. Kampik and Najjar [25] proposed a conceptual framework relying on agent-based simulations and/or machine learning algorithms to determine which and whether a process variant should be allowed to be executed under given circumstances. Nevertheless, more research (in particular, implementing and empirically evaluating the proposed approaches) is necessary.

Concerning **C3**, an essential approach would be to enable ethical and normative reasoning as an inherent capability of the blockchain protocol. Works in the literature distinguish between *implicit* and *explicit* ethical agency [31, 18]. On the one hand, current agents have no understanding of what is ethically “good” and what is “bad” since they implicitly respect their designer’s assessment of

the ethical implication of each action (if any). On the other hand, agents might be empowered with unbiased (possibly BCT-based) explicit moral knowledge. Thus decoupling artificial autonomous systems distinguish from the (possibly biased) moral perception of the designer, thus avoiding immoral behavior with or without human intervention [31, 18]. Nevertheless, in a multi-stakeholder scenario, integrating and reflecting the moral values and views of all stakeholders in the moral behavior of the autonomous system is more challenging. A possible solution is proposing an artificial moral agent architecture that uses normative systems’ techniques and formal argumentation to reach a moral agreement among stakeholders [29]. Using this technique, it is possible to abstract how a particular stakeholder can reach a particular decision concerning the morality of an action. Each stakeholder can be modeled as a source of arguments where an argument can be a statement about whether an action is moral or a reason for considering a particular action as moral. Therefore, the final decision will be the outcome of a consensus-making process involving the normative systems of all the stakeholders.

Concerning **C4**, a key aspect is to equip virtual or human agents with mechanisms and technologies to control both *data* and *behaviors* governed by BCT. To do so, we envision the development of new methodologies for co-constructing BCT solutions, in which MAS help bridging the gap between human interactions and the definition of smart contracts-based behaviors.

Concerning **C5**, intermediaries can play crucial roles in “interpreting” will and commands of its human user, detecting possible procedural or contextual mistakes, and in solving possible problems between the contracting parts. These scenarios demand that MAS leverage on smart contract technologies to enhance trust in their behaviors. Therefore, BCT-enabled MAS must provide internal procedures to assess and review possibly defective AI.

As described in **C6**, to assess delegation and correctness of the ledger requires considering the authenticity of information leading to reputation assessment, as well as its subjectivity. Possible approaches may include the study of decentralized regulation mechanisms where agents can use traceability analysis to determine the accuracy and/or correctness of the ledger’s contents. Moreover, the agents may be equipped with mechanisms for evaluating which data or interaction they are willing to make publicly visible through the ledger. The agents’ internal knowledge and their goals will play a fundamental role in the decision-making process that will evaluate potential trade-offs.

Concerning **C7**, the biases characterizing the human judgment are linked more to human psychology and nature rather than to BCT and MAS. For instance, the fact that prior ratings impact the evaluations of subsequent reviewers or that some users may artificially enhance the trustworthiness of others when writing reviews because they might be friends is not a technical matter.

However, the fact that the BCT “sets things in stone” raises the responsibility of the system’s designers to elaborate on the ethics of rating dynamics beyond the existing procedures of conciliation. Thus, we suggest undertaking two significant improvements that would enable to counter-balance the potential harm of an

under-rating. First, the rights of human agents would benefit from a culture of pseudonymization in MAS. Indeed, the systematic pseudonymization could help to mitigate the bias of under-reporting of negative ratings and to avoid that under-assessment have further consequences outside a MAS. Second, the designers might need to implement a “right to a second chance” either in MAS, BCT, or both. Differently from the “right to be forgotten”, the “right to a second chance” does not require deleting data at the end of a contractual commitment. Although it is ethically debatable, an individual can have the right to create a new account when his reputation is worsening. MAS, where the accounts are linked to transparent identities, would not be affected (even if a person wants to provide the same service that had been previously undermined).

Concerning **C8**, the right to be forgotten brings us back to the question of liability for which no concrete solution has been identified yet. Indeed, propositions such as removing the access to a given piece of data (e.g., impeding the execution of a smart contract or destroying its key — if encrypted) are still debated. Nevertheless, as Ward points out, the main data at hand when analyzing blockchains is transactional information and public keys, and only transactional data can be considered as personal data [41]. Henceforth, the only viable trend currently envisioned is to limit transactional information in a blockchain environment, by storing the personal data in a database without affecting the blockchain itself – off-chain – which would allow complying with data regulation while still keeping some advantages of the BCT [3, 19].

5 Conclusion

BCT-enabled MAS can boost next-generation systems yielding many societal advancements. Bridging MAS and BCT envisions stimulating opportunities and exciting challenges. Nevertheless, several ethical concerns (mostly implicated by the early stage of the involved technologies) will need to be carefully taken into consideration. This paper aimed at providing a first elicitation and mapping of the ethical challenges and research opportunities in the field of BCT-enabled MAS. In particular, we elaborated on the current ethical concerns entangled with the growing technological challenges. Although the systems’ complexity and the sophistication of their AI engines are growing at a fast pace, compliance with ethical standards will play a crucial role in future designs. However, the current research is still far from undertaking most of the challenges identified in this study. The solutions suggested in this paper are intended to provide credible paths for future research and to foster the recent initiatives for explainable, transparent, and responsible AI within full compliance with ethical regulations.

References

1. Abramova, O., Shavanova, T., Fuhrer, A., Krasnova, H., Buxmann, P., et al.: Understanding the sharing economy: The role of response to negative reviews in the peer-to-peer accommodation sharing network. In: ECIS 2015 Completed Research Papers. p. Paper 1. ECIS AISel, IL, USA (2015)

2. Adamic, L.A., Lauterbach, D., Teng, C.Y., Ackerman, M.: Rating friends without making enemies. In: *Fifth International AAAI Conference on Weblogs and Social Media* (2011)
3. Albanese, G., Calbimonte, J.P., Schumacher, M., Calvaresi, D.: Dynamic consent management for clinical trials via private blockchain technology. *Journal of Ambient Intelligence and Humanized Computing* pp. 1–18 (2020)
4. Anjomshoae, S., Najjar, A., Calvaresi, D., Främling, K.: Explainable agents and robots: Results from a systematic literature review. In: *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*. pp. 1078–1088. International Foundation for Autonomous Agents and Multiagent Systems (2019)
5. Aste, T., Tasca, P., Di Matteo, T.: Blockchain technologies: The foreseeable impact on society and industry. *Computer* **50**(9), 18–28 (2017)
6. Baldoni, M., Baroglio, C., Boissier, O., Micalizio, R., Tedeschi, S.: Engineering business processes through accountability and agents. In: *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*. pp. 1796–1798. International Foundation for Autonomous Agents and Multiagent Systems (2019)
7. Baldwin, T.: George edward moore. In: Zalta, E.N. (ed.) *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, summer 2010 edn. (2010)
8. BambouClub: Are you in the Bitcoin 1% ? A New Model of the Distribution of Bitcoin Wealth. <https://medium.com/@BambouClub/are-you-in-the-bitcoin-1-a-new-model-of-the-distribution-of-bitcoin-wealth-6adb0d4a6a95> (2017), [Online; accessed 04-November-2019]
9. Bruckner, M., LaFleur, M., Pitterle, I.: Frontier issues: The impact of the technological revolution on labour markets and income distribution. Department of Economic & Social Affairs, UN, accessed **24** (2017)
10. Buocz, T., Ehrke-Rabel, T., Hödl, E., Eisenberger, I.: Bitcoin and the gdpr: Allocating responsibility in distributed networks. *Computer Law & Security Review* **35**(2), 182–198 (2019)
11. Calvaresi, D., Cesarini, D., Sernani, P., Marinoni, M., Dragoni, A., Sturm, A.: Exploring the ambient assisted living domain: a systematic review. *Journal of Ambient Intelligence and Humanized Computing* pp. 1–19 (2016)
12. Calvaresi, D., Marinoni, M., Dragoni, A.F., Hilfiker, R., Schumacher, M.: Real-time multi-agent systems for telerehabilitation scenarios. *Artificial intelligence in medicine* **96**, 217–231 (2019)
13. Calvaresi, D., Mualla, Y., Najjar, A., Galland, S., Schumacher, M.: Explainable multi-agent systems through blockchain technology. In: *Explainable, Transparent Autonomous Agents and Multi-Agent Systems - First International Workshop, EXTRAAMAS 2019, Montreal, QC, Canada, May 13-14, 2019, Revised Selected Papers*. pp. 41–58. Springer, Berlin Heidelberg (2019). https://doi.org/10.1007/978-3-030-30391-4_3
14. Dierksmeier, C., Seele, P.: Cryptocurrencies and business ethics. *Journal of Business Ethics* **152**(1), 1–14 (2018)
15. Dignum, V.: Responsible autonomy. In: *Proceedings of the 26th International Joint Conference on Artificial Intelligence*. pp. 4698–4704. AAAI Press (2017)
16. Dostov, V., Shust, P.: Cryptocurrencies: an unconventional challenge to the aml/cft regulators? *Journal of Financial Crime* **21**(3), 249–263 (2014)

17. Dubovitskaya, A., Urovi, V., Barba, I., Aberer, K., Schumacher, M.I.: A multi-agent system for dynamic data aggregation in medical research. *BioMed Research International* (2016)
18. Dyrkolbotn, S., Pedersen, T., Slavkovik, M.: On the distinction between implicit and explicit ethical agency. In: *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*. pp. 74–80. ACM (2018)
19. Finck, M.: Blockchains and data protection in the european union. *Eur. Data Prot. L. Rev.* **4**, 17 (2018)
20. Fradkin, A., Grewal, E., Holtz, D., Pearson, M.: Bias and reciprocity in online reviews: Evidence from field experiments on airbnb. In: *Proceedings of the Sixteenth ACM Conference on Economics and Computation*. pp. 641–641. ACM (2015)
21. Gunning, D.: Explainable artificial intelligence (xai). *Defense Advanced Research Projects Agency (DARPA)*, nd Web **2** (2017)
22. Gürçan, Ö.: Multi-agent modelling of fairness for users and miners in blockchains. In: *International Conference on Practical Applications of Agents and Multi-Agent Systems*. pp. 92–99. Springer (2019)
23. Kampik, T., Najjar, A., Calvaresi, D.: Mas-aided approval for bypassing decentralized processes: an architecture. In: *2018 IEEE/WIC/ACM International Conference on Web Intelligence (WI)*. pp. 713–718 (Dec 2018). <https://doi.org/10.1109/WI.2018.000-6>
24. Kampik, T., Najjar, A.: Simulating, off-chain and on-chain: Agent-based simulations in cross-organizational business processes. *Information* **11**(1), 34 (2020)
25. Kampik, T., Najjar, A.: Simulating, off-chain and on-chain: Agent-based simulations in cross-organizational business processes. *Information* **11**(1), 34 (Jan 2020). <https://doi.org/10.3390/info11010034>, <http://dx.doi.org/10.3390/info11010034>
26. Kim, J.H., Kim, T.: Economic assimilation of north korean refugees in south korea: Survey evidence. *KDI School of Pub Policy & Management Paper (06-19)* (2006)
27. Krugman, P.: Bitcoin is evil. *The New York Times* **28**, 2013 (2013)
28. Lauterbach, D., Truong, H., Shah, T., Adamic, L.: Surfing a web of trust: Reputation and reciprocity on couchsurfing. com. In: *2009 International Conference on Computational Science and Engineering*. vol. 4, pp. 346–353. IEEE (2009)
29. Liao, B., Slavkovik, M., van der Torre, L.: Building jiminy cricket: An architecture for moral agreements among stakeholders. In: *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*. pp. 147–153. AIES '19, ACM, New York, NY, USA (2019). <https://doi.org/10.1145/3306618.3314257>, <http://doi.acm.org/10.1145/3306618.3314257>
30. Mayzlin, D., Dover, Y., Chevalier, J.: Promotional reviews: An empirical investigation of online review manipulation. *American Economic Review* **104**(8), 2421–55 (2014)
31. Moor, J.H.: The nature, importance, and difficulty of machine ethics. *IEEE intelligent systems* **21**(4), 18–21 (2006)
32. Moya, V.: GoodDollar: cryptocurrencies would end inequality. <https://latinamericanpost.com/24800-gooddollar-cryptocurrencies-would-end-inequality> (2019), [Online; accessed 04-November-2019]
33. Origgi, G., Pais, I.: Digital reputation in the mutual admiration society. *Studi di sociologia* **2**(2), 175–193 (2018)
34. Overgoor, J., Wulczyn, E., Potts, C.: Trust propagation with mixed-effects models. In: *Sixth International AAAI Conference on Weblogs and Social Media*. AAAI (2012)
35. Ramchurn, S.D., Huynh, D., Jennings, N.R.: Trust in multi-agent systems. *The Knowledge Engineering Review* **19**(1), 1–25 (2004)

36. Sabater-Mir, J., Vercouter, L.: Trust and reputation in multiagent systems. *Multiagent systems* p. 381 (2013)
37. Smith, J., Tennison, J., Wells, P., Fawcett, J., Harrison, S.: Applying blockchain technology in global data infrastructure. <https://theodi.org/article/applying-blockchain-technology-in-global-data-infrastructure/> (2016), open Data Institute
38. Stross, C.: Why i want bitcoin to die in a fire. <http://www.antipope.org/charlie/blog-static/2013/12/why-i-want-bitcoin-to-die-in-a.html> (2013)
39. Tussyadiah, I.P.: Strategic self-presentation in the sharing economy: Implications for host branding. In: *Information and Communication Technologies in Tourism 2016*. pp. 695–708. Springer (2016)
40. Ward, J., Hemingway, C., Daniel, E.: A framework for addressing the organisational issues of enterprise systems implementation. *The Journal of Strategic Information Systems* **14**(2), 97–119 (2005)
41. Ward, Y., et al.: An analysis of the applicability of gdpr to blockchain technologies. *L'Europe Unie* **13**(13), 86–90 (2018)