

Consent as a Foundation for Responsible Autonomy

Munindar P. Singh

Department of Computer Science
North Carolina State University
Raleigh, NC 27695-8206, USA
singh@ncsu.edu

Abstract

This paper focuses on a dynamic aspect of responsible autonomy, namely, to make intelligent agents be responsible at run time. That is, it considers settings where decision making by agents impinges upon the outcomes perceived by other agents. For an agent to act responsibly, it must accommodate the desires and other attitudes of its users and, through other agents, of their users.

The contribution of this paper is twofold. First, it provides a conceptual analysis of consent, its benefits and misuses, and how understanding consent can help achieve responsible autonomy. Second, it outlines challenges for AI (in particular, for agents and multiagent systems) that merit investigation to form as a basis for modeling consent in multiagent systems and applying consent to achieve responsible autonomy.

1 Introduction

Recent advances in the capabilities of AI have rightly brought the concerns of responsible AI to the forefront. Mainstream AI efforts consider topics such as algorithmic accountability (Diakopoulos 2016) and fairness, referring to statistical properties of AI algorithms (Leben 2020). These approaches consider the application of AI in different settings, such as assistance in judicial sentencing guidelines or loan application processing, and bring out the broader societal context. However, in those cases the context is largely fixed—we would not be able to easily change the judicial or financial systems, and get rid of the disparities and inequities entrenched in those systems.

Research into responsibility in AI focuses on static aspects such as design methodologies and development practices for AI (Dignum 2017, 2019), e.g., via checklists for agent developers. In contrast, although such design-time aspects of responsible AI are undoubtedly valuable, we focus here on *responsible autonomy*, which we define as including the challenges of ensuring that an autonomous agent exercises its autonomy responsibly, i.e., at runtime.

There are increasing prospects of people contesting decisions made by (or using) AI, especially when an AI agent's decisions affect people without being mediated by a human. Consent would be a crucial element of any justification of agent decisions on ethical and legal grounds

Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

1.1 Sociotechnical Systems

In our formulation of responsible autonomy, we consider agents embedded in microsocieties, wherein the context is modeled computationally and can be reasoned about and potentially controlled by the agents. We model such a micro-society as a *sociotechnical system (STS)* comprising autonomous social entities (people and organizations or *principals*) and technical entities (software or *agents*, who help principals) (Kafali, Ajmeri, and Singh 2020; Singh 2013). Simply put, we place agents in a social setting. Designers of STSs and agents and the agents themselves at runtime represent and reason about the social setting provided by an STS. A classical means to model the relevant parts of the social setting is through deontic constructs called *norms*. Norms in this sense encompass both social norms of the ordinary vernacular and legal constructs (Von Wright 1963, 1999).

We posit that understanding responsible autonomy from a dynamic perspective will facilitate responsible AI in static settings by making explicit what is otherwise hidden.

1.2 Contributions

We formulate consent from an STS standpoint to address responsible autonomy. We identify crucial criteria for consent based on a brief review of the relevant literature. We articulate a way of understanding consent in AI that accommodates both the social and the technical architecture of an STS. Our overarching claim is that consent provides a new foundation for approaching responsibility that highlights autonomy, social interaction, and accountability.

Incorporating consent in AI is not only practically important as AI becomes more capable (and open to contestation), but also exposes important research questions. Specifically, consent lies at the heart of autonomy (Alexander 1996; Hurd 1996) and our STS formulation ties it to responsible autonomy in individual and social action. The ensuing research questions and require participation by AI research communities in multiagent systems, human-AI systems, agent communication, and agent architectures with additional opportunities in dialogue understanding and active learning.

Organization Section 2 introduces governance. Section 3 describes how consent is used in computing, law, and philosophy. Section 4 summarizes key challenges for consent. Section 5 discusses our vision for future research.

2 Responsible Autonomy and Governance

We contrast consent-based governance for autonomy.

2.1 Moral Quandaries

Besides statistical properties of algorithms, a second dominant theme in research into AI ethics concerns agent decision making, focusing on moral quandaries an agent may face. These quandaries are formulated in contexts where an agent's decision has outcomes (on third parties) with ethical import. The canonical such quandary is where a (train) trolley is running on some tracks and the agent has the option to switch it to an alternative pair of tracks (Foot 1967). People of differing numbers and attributes are tied to each pair of tracks. Thus, the agent's decision would lead to saving some and killing some. The term "trolley problem" is often used generically for all such moral dilemmas, many of which do not involve vehicles (Wood 2011).

Fried (2012) characterizes trolley problems as posing hypothetical dilemmas where (1) the agent is an individual, not an institution; (2) the agent faces a one-off decision; (3) the causal chain between decision and outcome is short; (4) the consequences of the decision are known *a priori* with certainty; and (5) there are no opportunities for scaling up the moral reasoning or its principles to a larger number and variety of cases. In essence, the trolley problems are rigged to be context-free situations seeking to highlight and promote a non-consequential viewpoint, downplaying any aggregative (utilitarian) approach to reasoning about ethics and risk.

But realistic AI approaches, just like human decision making (Fried 2012, p. 506), must accommodate uncertainty and trade off risks in the long run. Thus, we posit that focusing on moral quandaries has little to offer in the way of valuable research questions for responsible autonomy.

2.2 Governance

Responsible autonomy means not only promoting one's own values and preferences but also refraining from violating the values and preferences of others. In addition to producing agents who are individually responsible, we want to provide system-level guidance for responsibility. This is the province of *governance* in STS (Baldoni et al. 2021; Singh 2013), not to be confused with offline designs or administrative processes through which agents may be coordinated. That is, governance concerns how agents operate and interact to achieve system-level objectives while satisfying their users' needs. A motivation for this formulation of governance is that in complex contexts, it is not viable to produce constraints that are inflexible and effective; therefore, the agents much govern themselves (i.e., each other). Such flexibility is especially valuable when an agent's decisions affect the outcomes for others, that is, when responsibility matters.

In an STS setting, the challenge of achieving responsible autonomy splits into two parts, corresponding respectively to macro and micro ethics (Chopra and Singh 2018):

- Specifying an STS based on the requirements that minimizes unethical outcomes as judged by its stakeholders.
- Specifying agents who behave in a way compatible with the values of their stakeholders and the norms of the STS.

2.3 Consent-Based Governance

Both of the above components of responsible autonomy rely upon an approach for governance based on a proper understanding of *consent*. Specifically, our motivation for placing consent at the center of responsible autonomy is that consent is a foundational construct in autonomy, both in terms of exercising one's autonomy and in recognizing the autonomy of others, the latter being an essential element of Kantian ethics (Hill Jr. 1980, p. 90).

As an illustration of responsible autonomy via consent, suppose an agent assists its user, Alice, in taking actions that promote Alice's goals. Consider a mobile social application that supports sharing one's picture or location. The problem is social because one user's action affects another's privacy (Kurtan and Yolum 2021; Mosca, Such, and McBurney 2020), such as when a joint picture is shared. When Alice's agent acts on behalf of her, it affects outcomes for her and for other people. For example, when posting a picture of Alice with her friends, the agent must act responsibly regarding the wellbeing both of Alice and of the friends. The agent may have a formal fiduciary duty toward Alice and at least an informal fiduciary duty toward her friends, the latter reflecting her moral duty not to exploit her friends.

Consent is a natural abstraction here. In informal terms, the agent should ensure that the people whose picture is being shared consent to the sharing. The agent may have previously obtained consent from Alice and must explicitly or otherwise obtain consent from the others. More generally, for ethical behavior, ideally, any party affected adversely, or potentially so, must consent to the first party's actions. How can we specify the sharing micro-society (possibly realized via the app) so that the values of the users are respected? How can we build such an agent to apply the norms of the corresponding STS?

Surprisingly, however, consent has not been studied in AI. That is, consent is applied in a purely uninterpreted manner without any way to represent and reason about it.

3 The Many Uses of Consent

We review some key applications of consent.

3.1 Moral and Legal Legitimacy

Understanding consent is crucial in assessing the legitimacy of an action, which would determine whether a crime took place. For example, consent is the difference between borrowing and stealing and between lovemaking and rape. The *volenti maxim* is that an explicit consent (or request implying consent) overrides ordinary prohibitions (Dempsey 2018). Consent characterizes when some action by one autonomous party gains legitimacy despite potentially infringing upon the autonomy or authority of another party—that is its "moral magic" (Alexander 1996; Hurd 1996).

3.2 Consent of the Governed

Political philosophy has one of the oldest traditions in consent, going back to John Locke, and using the "consent of the governed" as the basis for the legitimacy of government. This topic is relevant to AI because the membership of an

agent in an STS subjects the agent to the norms of that STS (relative to its role in the STS) and, therefore, the agent must consent to play a specific role in that STS.

Pitkin (1965, 1966) describes how consent of the governed relates to their political obligation to obey the government. She brings forth challenges in how the traditional (Lockean) notion of consent may be applied. Specifically, Pitkin makes a case that the concept proves to be vacuous in human societies. On the one hand, Locke argues consent by an individual is essential for that individual to be subject to the laws of society. On the other hand, he postulates tacit consent derived merely from living in a jurisdiction: that is, nonconsenting individuals (competent adults, for the sake of simplicity) are governed the same as consenting individuals.

Pitkin advocates for the alignment of values as the crucial point in that an agent ought not to merely think of their consent as conferring legitimacy but of an evaluation of the moral nature of the government as conferring legitimacy. In this regard, she also calls out an emphasis on the opposite of obeying under consent, namely, a duty to resist tyranny.

3.3 Business Contracts

Consent turns up as a basis for modeling business contracts to capture the idea that the contracting parties enter into contracts freely and to avoid some of the challenges arising in accounts based on (1) intent or expectation; (2) efficiency or fairness; and (3) enforcement processes (Barnett 1986). Consent provides a way to model the social context and capture what each party's entitlements are in that context and how they are assigned through a contract. An illustration would be in arbitration clauses through which the contracting parties agree to waive a jury trial in case of a dispute.

Consent accords well with *relational contracts* (Bernstein 1993, 2015), which focus on social relationships between contracting parties. By modeling the social context, relational contracts accommodate renegotiation on the fly to avoid disputes that invariably arise because no contract can specify all possible eventualities.

3.4 Consent in Computing

Established practice in computing regarding consent goes back to work on privacy early in the information age. Westin (1967) established the influential doctrine of *notice and choice*, under which all you need to do to respect an individual's privacy rights is to (1) disclose what information of theirs you are obtaining, storing, using, or sharing, and (2) ask them to consent to that action.

Notice and choice spread because it is easy to implement. One usage is in click-wrap licenses for a software product where a user is provided a license and must accept its terms before accessing the product. Likewise, many websites demand that users consent to being tracked. And, social media apps have users consent to their information being analyzed and shared with third parties.

But users are ill-equipped to figure out the ways their information may be shared, analyzed, and combined with other information. Therefore, notice and choice has been criticized by privacy scholars (Schermer, Custers, and van der Hof 2014; Sloan and Warner 2014). Nissenbaum

(2004) makes a case for understanding the contexts in which information is shared and developing privacy norms accordingly. Notice and choice hides the contexts and blind sides the user, leading to their information propagating across contexts that are decoupled in their mind.

Notice and choice naively assumes that people are rational (Hoofnagle and Urban 2002). But people consent to whatever terms are offered because of complexity and the power differential (the terms being non-negotiable) (Fassl, Gröber, and Krombholz 2021). Recent work in human-computer interaction seeks improved models of and user interfaces for consent (Im et al. 2021; Lindegren et al. 2021), but this work doesn't offer much for responsible autonomy.

All considered, consent in computing is often misguided or ill-intentioned, and does not shed light on responsibility.

4 Challenges to Understanding Consent

Intuitions about consent are far from established (Schnüriger 2018). First, consent reflects a mental action of the consenting party—indicating that it is the exercise of an internal choice. Second, consent reflects a communicative act or performative by the consenting party conferring powers on the recipient—indicating that it is the exercise of a normative power (Hohfeld 1919; Hurd 2018; Koch 2018). The mental approach doesn't explain how a normative power arises from an internal action without a communication. The communicative approach doesn't explain the treatment of mistakes in performing a communication that grants consent. Alexander (2014) distinguishes *wrongdoing* (causing harm by acting without there being true mental consent on part of the party whose consent was necessary) from *culpability* (acting without belief that the requisite consent exists). These distinctions are important for ascribing blame. Table 1 summarizes important criteria, using a grouping explained next.

Criterion	Example or Explanation
Visibility	Consent is observable, i.e., a communication
Free will	Consenter acts without being coerced
Truth	Consenter's beliefs are true and complete
Capacity	Consenter is mentally fit
Cognition	Consenter believes and intends to
Attention	Consenter exercises mental faculties
Statutes	Consenter meets statutory criteria, e.g., age
Power	Consenter is not subjugated by consentee
Honesty	Consentee does not mislead consenter

Table 1: Important criteria in consent grouped as Habermas's (1984) objective, subjective, and practical validity claims.

5 Vision: Research on Consent in AI

Understanding consent is not only a prerequisite for achieving responsible autonomy, it is also a subtle concept that melds intuitions regarding ethics, law, usability, and decision making. Our sociotechnical stance reveals important opportunities for AI research. We first outline ideas for a semantics of consent and then some promising research directions.

5.1 Toward a Social Semantics of Consent

Traditional disputes, as between the objective and subjective elements of consent, arise because of a confusion of meanings and meaning standards. We propose to apply Habermas's (1984) framework to reconcile these ideas by building on a public semantics for agent communication (Singh 2000), which was an adaptation of Habermas. Specifically, we enhance the original nonmentalistic conception of communication (Austin 1962; Sbisà 2007, 2018). We develop validity criteria for consent from the perspectives of the consenter, the consentee, and a third party (potentially the STS in which consenter and consentee interact).

Habermas's (1984) theory of communication in the public sphere associates three *validity claims* (i.e., distinct standards of meaning) with each communication: *objective* (true); *subjective* (appropriate beliefs and intentions); *practical* (justified in the social context). We relate these claims to the criteria in Table 1. Objectively, granting consent is a social action and its meaning is for the consenter to empower the consentee, e.g., by forgoing any moral objections to the consentee acting as specified. It is valid under visibility and free will. Subjectively, consent is a mental object and granting consent an intentional action. Its meaning is the corresponding intention. It is valid provided it is performed with a capacity to reason about consent, holds the right cognitive state under full attention, and the consenter's beliefs pertinent to the consent are true and include relevant facts. Practically, granting consent is valid if the consent is not prohibited by statutes (i.e., norms of the STS), the consenter is not subjugated by the consentee or misled by the consentee.

5.2 Bridging AI Ethics and Law

Consent is key in distinguishing right from wrong and in legitimizing actions and making them legal. Thus, it opens research challenges on bridging the gap between ethics and law to develop responsible agents.

Legal positivism is the doctrine that the law is as it is posited. Variants of positivism take stronger or weaker stances on these key theses (Green and Adams 2019; Himma 2021): (1) *pedigree*, that the law's existence and validity rely upon social facts, i.e., that it is declared a law; (2) *separability*, that though the law and morality may align, it is not necessary that they do; (3) *fallibility*, that the law may be intended to be moral and yet be deficient in that regard; (4) *neutrality*, that even though the law is not value neutral, it should be described and argued about in value-free terms; and (5) *discretion*, that judges may exercise discretion where the law is not clear and in doing so, they extend the law.

Natural law is an older doctrine that the law derives its legitimacy from being *natural*, i.e., granted by nature or by divine right. In modern versions, it is the view that ethics be incorporated into the law (Gavison 1982). Modern judicial practice focuses on "applying the law" and avoiding justifications based on ethics. That is, natural law has largely been supplanted by legal positivism, whose main theses summarized above sit largely in contrast with natural law.

Consent today is conceived of legally positively, i.e., based on a consenter's utterances. But its shortcomings arise

mainly because this thinking disregards ethics: is consent right if obtained from someone who is desperate? The discretion thesis of legal positivism with practical validity from Habermas provides a potential opening in how we might formulate consent so that it bridges the gap with ethics.

5.3 Verifying Agents and STSs for Consent

Responsible autonomy presupposes that an STS would provide social and technical controls to promote responsible actions by its members and to limit harm in case of malfunction or malfeasance. How can we verify an STS and its member agents to ensure that they provide consent where needed, refrain from doing so where it is not needed, and respect the absence of consent from others in their decisions? For an STS, in addition, we need ways to (1) minimize the risk of an agent being placed in an ethical quandary when providing or receiving consent; and (2) ensuring that social and technical controls on agent behavior balance any propensity to violate another's consent.

The literature on consent takes a *retrospective* view, as to adjudicate on a violation in a court of law. For AI ethics, the *prospective* view of consent and responsibility (Van de Poel 2011) is no less important, to assess an agent's decisions about when to consent and when to obtain and act based on another's consent. This direction would lead to formal reasoning for verifying and certifying (Dennis et al. 2016; Fisher et al. 2021) agents and STSs with respect to consent.

5.4 Consent-Based Design

This paper expands the ontology of requirements beyond goals and dependencies (Yu et al. 2011), legal norms such as social commitments (Chopra et al. 2014) and powers (Singh 2013), and values (Cranefield et al. 2017) to give first-class status to consent. How would we produce suitably equipped design methodologies that use consent to express requirements? Such a methodology would accommodate the criteria of Table 1. For example, using these criteria, a suitable methodology would help model *consentability*—the power accorded within an STS to its participants regarding the kinds of consent they can issue (Kim 2019) that has the appropriate pragmatic consequences in enabling further action by other participants or their agents.

5.5 Learning and Interaction about Consent

An important direction is to model agent-user dialogue so that an agent can elicit its user's consent, obtain valid consent (as described above), and explain its decisions in light of consent from its user as well as from other agents (on behalf of their users). In this light, how might we extend research on values (Liscio et al. 2021) to develop methods by which an agent understands its user's consent? To evade the criticisms of traditional methods as discussed in Section 3.4, a desirable approach must be explicit yet not tedious, even when the consent is nuanced and contextual.

Acknowledgments

Thanks to the NSF (grant IIS-2116751) and to the DoD (Science of Security Lablet) for support.

References

- Alexander, L. 1996. The Moral Magic of Consent (II). *Legal Theory*, 2(3): 165–174.
- Alexander, L. 2014. The Ontology of Consent. *Analytic Philosophy*, 55(1): 102–113.
- Austin, J. L. 1962. *How to Do Things with Words*. Oxford: Clarendon Press. ISBN 978-0-19-824553-7.
- Baldoni, M.; Baroglio, C.; Micalizio, R.; and Tedeschi, S. 2021. Robustness Based on Accountability in Multiagent Organizations. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, 142–150. Online: IFAAMAS.
- Barnett, R. E. 1986. A Consent Theory of Contract. *Columbia Law Review*, 86(2): 269–321.
- Bernstein, L. 1993. Social Norms and Default Rules Analysis. *Southern California Interdisciplinary Law Journal (ILJ)*, 3(1): 59–90.
- Bernstein, L. 2015. Beyond Relational Contracts: Social Capital and Network Governance in Procurement Contracts. *Journal of Legal Analysis*, 7(2): 561–621.
- Chopra, A. K.; Dalpiaz, F.; Aydemir, F. B.; Giorgini, P.; Mylopoulos, J.; and Singh, M. P. 2014. Protos: Foundations for Engineering Innovative Sociotechnical Systems. In *Proceedings of the 22nd IEEE International Requirements Engineering Conference (RE)*, 53–62. Karlskrona, Sweden: IEEE Computer Society.
- Chopra, A. K.; and Singh, M. P. 2018. Sociotechnical Systems and Ethics in the Large. In *Proceedings of the AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society (AIES)*, 48–53. New Orleans: ACM.
- Cranefield, S.; Winikoff, M.; Dignum, V.; and Dignum, F. 2017. No Pizza for You: Value-based Plan Selection in BDI Agents. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, 178–184. Melbourne: IJCAI.
- Dempsey, M. M. 2018. The Volenti Maxim. In Schaber, P.; and Müller, A., eds., *The Routledge Handbook of the Ethics of Consent*, chapter 17, 187–198. London: Routledge.
- Dennis, L. A.; Fisher, M.; Slavkovik, M.; and Webster, M. 2016. Formal Verification of Ethical Choices in Autonomous Systems. *Robotics and Autonomous Systems*, 77: 1–14.
- Diakopoulos, N. 2016. Accountability in Algorithmic Decision Making. *Communications of the ACM (CACM)*, 59(2): 56–62.
- Dignum, V. 2017. Responsible Autonomy. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, 4698–4704. Melbourne: IJCAI.
- Dignum, V. 2019. *Responsible Artificial Intelligence: How to Develop and Use AI in a Responsible Way*. Number 8 in Artificial Intelligence: Foundations, Theory, and Algorithms. Cham, Switzerland: Springer.
- Fassl, M.; Gröber, L. T.; and Krombholz, K. 2021. Stop the Consent Theater. In *Extended Abstracts of the Conference on Human Factors in Computing Systems (CHI)*, 1–7. Yokohama, Japan: ACM.
- Fisher, M.; Mascardi, V.; Rozier, K. Y.; Schlingloff, B.-H.; Winikoff, M.; and Yorke-Smith, N. 2021. Towards a Framework for Certification of Reliable Autonomous Systems. *Autonomous Agents and Multi-Agent Systems*, 35(1): 8:1–8:65.
- Foot, P. 1967. The Problem of Abortion and the Doctrine of Double Effect. *Oxford Review*, 5: 5–15.
- Fried, B. H. 2012. What Does Matter? The Case for Killing the Trolley Problem (or Letting it Die). *The Philosophical Quarterly*, 62(248): 505–529.
- Gavison, R. 1982. Natural Law, Positivism, and the Limits of Jurisprudence: A Modern Round. *Yale Law Journal*, 91(2): 1250–1286.
- Green, L.; and Adams, T. 2019. Legal Positivism. In Zalta, E. N., ed., *The Stanford Encyclopedia of Philosophy*. Stanford: Metaphysics Research Lab, Stanford University, Winter 2019 edition.
- Habermas, J. 1984. *The Theory of Communicative Action, volumes 1 and 2*. Cambridge, United Kingdom: Polity Press.
- Hill Jr., T. E. 1980. Humanity as an End in Itself. *Ethics*, 91(1): 84–99.
- Himma, K. E. 2021. Legal Positivism. In *The Internet Encyclopedia of Philosophy*. Martin, Tennessee: IEP. Accessed: 2021-08-11.
- Hohfeld, W. N. 1919. *Fundamental Legal Conceptions as Applied in Judicial Reasoning and other Legal Essays*. New Haven, Connecticut: Yale University Press. A 1919 printing of articles from 1913.
- Hoofnagle, C. J.; and Urban, J. M. 2002. Alan Westin’s Privacy homo economicus. *Wake Forest Law Review*, 49(4): 261–317.
- Hurd, H. M. 1996. The Moral Magic of Consent. *Legal Theory*, 2(2): 121–146.
- Hurd, H. M. 2018. The Normative Force of Consent. In Schaber, P.; and Müller, A., eds., *The Routledge Handbook of the Ethics of Consent*, chapter 4, 44–54. London: Routledge.
- Im, J.; Dimond, J.; Berton, M.; Lee, U.; Mustelier, K. W.; Ackerman, M. S.; and Gilbert, E. 2021. Yes: Affirmative Consent as a Theoretical Framework for Understanding and Imagining Social Platforms. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI)*, 403:1–403:18. Yokohama, Japan: ACM.
- Kafalı, Ö.; Ajmeri, N.; and Singh, M. P. 2020. DESEN: Specification of Sociotechnical Systems via Patterns of Regulation and Control. *ACM Transactions on Software Engineering and Methodology (TOSEM)*, 29(1): 7:1–7:50.
- Kim, N. S. 2019. *Consentability: Consent and its Limits*. Cambridge, United Kingdom: Cambridge University Press.
- Koch, F. 2018. Consent as a Normative Power. In Schaber, P.; and Müller, A., eds., *The Routledge Handbook of the Ethics of Consent*, chapter 3, 32–43. London: Routledge.
- Kurtan, A. C.; and Yolum, P. 2021. Assisting Humans in Privacy Management: An Agent-Based Approach. *Autonomous Agents and Multi-Agent Systems*, 35(1): 7.

- Leben, D. 2020. Normative Principles for Evaluating Fairness in Machine Learning. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society (AIES)*, 86–92. New York: ACM.
- Lindgren, D.; Karegar, F.; Kane, B.; and Pettersson, J. S. 2021. An Evaluation of Three Designs to Engage Users when Providing Their Consent on Smartphones. *Behaviour and Information Technology*, 40(4): 398–414.
- Liscio, E.; van der Meer, M.; Siebert, L. C.; Jonker, C. M.; and Murukannaiah, P. K. 2021. Axes: Identifying and Evaluating Context-Specific Values. In *Proceedings of the 20th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 1–9. London: IFAAMAS.
- Mosca, F.; Such, J. M.; and McBurney, P. 2020. Towards a Value-driven Explainable Agent for Collective Privacy. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*, 1937–1939. Auckland: IFAAMAS.
- Nissenbaum, H. 2004. Privacy as Contextual Integrity. *Washington Law Review*, 79(1): 119–157.
- Pitkin, H. F. 1965. Obligation and Consent—I. *American Political Science Review*, 59(4): 990–999.
- Pitkin, H. F. 1966. Obligation and Consent—II. *American Political Science Review*, 60(1): 39–52.
- Sbisà, M. 2007. How to read Austin. *Pragmatics*, 17(3): 461–473.
- Sbisà, M. 2018. Varieties of Speech Act Norms. In Witek, M.; and Witczak-Plisiecka, I., eds., *Normativity and Variety of Speech Actions*, volume 112 of *Poznań Studies in the Philosophy of the Sciences and the Humanities*, 23–50. Leiden, Netherlands: Brill Rodopi.
- Schermer, B. W.; Custers, B.; and van der Hof, S. 2014. The Crisis of Consent: How Stronger Legal Protection May Lead to Weaker Consent in Data Protection. *Ethics and Information Technology*, 16(2): 171–182.
- Schnüriger, H. 2018. What is Consent? In Schaber, P.; and Müller, A., eds., *The Routledge Handbook of the Ethics of Consent*, chapter 2, 21–31. London: Routledge.
- Singh, M. P. 2000. A Social Semantics for Agent Communication Languages. In *Proceedings of the 1999 IJCAI Workshop on Agent Communication Languages*, number 1916 in *Lecture Notes in Artificial Intelligence*, 31–45. Berlin: Springer.
- Singh, M. P. 2013. Norms as a Basis for Governing Sociotechnical Systems. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 5(1): 21:1–21:23.
- Sloan, R. H.; and Warner, R. 2014. Beyond Notice And Choice: Privacy, Norms, And Consent. *Journal of High Technology Law*, 14(2): 370–412.
- Van de Poel, I. 2011. The Relation between Forward-Looking and Backward-Looking Responsibility. In Vincent, N. A.; de Poel, I. V.; and den Hoven, J. V., eds., *Moral Responsibility: Beyond Free Will and Determinism*, number 27 in *Library of Ethics and Applied Philosophy*, chapter 3, 37–52. Berlin: Springer.
- Von Wright, G. H. 1963. *Norm and Action: A Logical Enquiry*. International Library of Philosophy and Scientific Method. New York: Humanities Press.
- Von Wright, G. H. 1999. Deontic Logic: A Personal View. *Ratio Juris*, 12(1): 26–38.
- Westin, A. F. 1967. *Privacy and Freedom*. New York: Atheneum.
- Wood, A. 2011. Humanity as End in Itself. In Parfit, D., ed., *On What Matters*, volume 2, 58–82. Oxford University Press.
- Yu, E.; Giorgini, P.; Maiden, N.; and Mylopoulos, J., eds. 2011. *Social Modeling for Requirements Engineering*. Cambridge, Massachusetts: MIT Press. ISBN 9780262240550.