

Schrödinger's Robot: Privacy in Uncertain States

*Ian Kerr**

Can robots or AIs operating independently of human intervention or oversight diminish our privacy? There are two equal and opposite reactions to this issue. On the robot side, machines are starting to outperform human experts in an increasing array of narrow tasks, including driving, surgery, and medical diagnostics. This is fueling a growing optimism that robots and AIs will exceed humans more generally and spectacularly; some think, to the point where we will have to consider their moral and legal status. On the privacy side, one sees the very opposite: robots and AIs are, in a legal sense, nothing.

* Canada Research Chair in Ethics, Law and Technology, University of Ottawa, Faculty of Law, iankerr@uottawa.ca. I would like to thank the Social Sciences and Humanities Research Council and the Canada Research Chairs program for their generous support. Special thanks to Carys Craig for teaching me about and inspiring me to undertake a relational account of privacy. Thank you to David Matheson for being my epistemological guardian angel; and to Joelle Pineau, Laurel Reik, Bill Smart, and Jodi Forlizzi for lending precision to some of my technical assertions. I am also grateful to the participants of *The Problem of Theorizing Privacy* conference, organized by Michael Birnhack, Julie Cohen and Mireille Hildebrandt. This event generated very thoughtful commentary from Eldar Haber, and useful feedback from Tal Zarsky, Mireille Hildebrandt, Helen Nissenbaum, Eran Toch, Ruth Gavison, Anita Allen, Michael Bar-Sinai, Alon Jasper, Lisa Austin, and Mauricio Figueroa Torres. This article also benefitted from a second presentation at the University of Surrey's *Workshop on the Regulation of AI* organized by Ryan Abbott and Alex Sarch, with excellent commentary from Steven Bero. Saving the best for last, my extreme gratitude goes out to Ida Mahmoudi for the outstanding research assistance that she so regularly and reliably provides and to Katie Szilagyi — engineer, lawyer, doctoral candidate *par excellence* and proud owner of these fine footnotes — for grace under pressure, her tireless enthusiasm, her ability to find anything under the sun, her insatiable intellectual curiosity, and her deep-seated disposition for *arête* ... which she has not only cultivated for herself but, through collaboration, inspires in others. Cite as: Ian Kerr, *Schrödinger's Robot: Privacy in Uncertain States*, 20 THEORETICAL INQUIRIES L. 123 (2019).

The received view is that since robots and AIs are neither sentient nor capable of human-level cognition, they are of no consequence to privacy law. This article argues that robots and AIs operating independently of human intervention can and, in some cases, already do diminish our privacy. Epistemic privacy offers a useful analytic framework for understanding the kind of cognizance that gives rise to diminished privacy. Because machines can actuate on the basis of the beliefs they form in ways that affect people's life chances and opportunities, I argue that they demonstrate the kind of cognizance that definitively implicates privacy. Consequently, I conclude that legal theory and doctrine will have to expand their understanding of privacy relationships to include robots and AIs that meet these epistemic conditions. An increasing number of machines possess epistemic qualities that force us to rethink our understanding of privacy relationships with robots and AIs.

PROLEGOMENA

One can even set up quite ridiculous cases. A robot is penned up in a steel chamber, programmed with a machine learning algorithm (giving it the potential to cognize personal data). Deep in a series of nested commands there is also a logic bomb, so designed that perhaps in the course of the hour any cognizing detected by the robot would satisfy the if/then Boolean for the logic bomb's activation, but also, with equal probability, perhaps not; if the robot cognizes personal data, the logic bomb detonates and, through a series of commands, deactivates the robot. If one has left this entire system to itself for an hour, one would say that the robot still functions if meanwhile no personal data have been cognized. The psi-function of the entire system would express this by having in it the living and dead robot (pardon the expression) mixed or smeared out in equal parts.¹

¹ This passage is inspired by the famous thought experiment known as "Schrödinger's Cat." Erwin Schrödinger, *The Present Situation in Quantum Mechanics: A Translation of Schrödinger's "Cat Paradox Paper,"* 124 PROCEEDINGS AM. PHIL. SOC'Y. 323 (1935) (John D. Trimmer trans., 1980). Its use and the title of this paper were inspired by an extremely interesting conversation with the brilliant Ashkan Soltani at *WeRobot 2016*, Ashkan Soltani, Discussant at *We Robot 2016 Conference: Legal & Policy Issues Relating to Privacy*, the University of Miami Newman Alumni Center (Apr. 1 2016).

INTRODUCTION

This Article responds to two equal and opposite reactions tugging at the intersection of robots and privacy.

On the robot side, too many technologists, government decision-makers, and captains of industry have become preoccupied with what they see as an inevitable shift from today's artificial narrow intelligence (ANI) to tomorrow's artificial general intelligence (AGI).² The fact that machines are starting to outperform human experts in an increasing array of narrow tasks³ fuels a growing optimism that AIs will exceed humans more generally and spectacularly. Seduced by the rapture of singularity⁴ and superintelligence,⁵ many credible⁶

-
- 2 AI is often divided into these two categories: Artificial Narrow Intelligence (ANI), in which systems can learn specific, defined tasks, and Artificial General Intelligence (AGI), in which systems would possess human-like intelligence, thereby becoming capable of completing general, undefined tasks. While current ANI research has made great strides in the former, AGI remains the stuff of science fiction. For more on this distinction, see Ryan Calo, *Artificial Intelligence Policy: A Primer and Roadmap*, 51 U.C. DAVIS L. REV. 399 (2017).
 - 3 Perhaps ANI's greatest advantage is its ability to perform repeatable tasks with precision and consistency. For example, the Soft Tissue Autonomous Robot (STAR) can perform surgery: stitching soft tissue together with a needle and thread more adeptly than the best doctor. See Will Knight, *Nimble-Fingered Robot Outperforms the Best Human Surgeons*, MIT TECH. REV. (May 4, 2016), <https://www.technologyreview.com/s/601378/nimble-fingered-robot-outperforms-the-best-human-surgeons/>. Relatedly, ANI is also making great strides in medical diagnostics, see A. Michael Froomkin, Ian R. Kerr & Joelle Pineau, *When AIs Outperform Doctors: The Dangers of a Tort-Induced Over-Reliance on Machine Learning and What (Not) to Do About It* (Univ. Miami Legal Studies, Working Paper No. 18-3, 2018), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3114347.
 - 4 RAY KURZWEIL, *THE SINGULARITY IS NEAR: WHEN HUMANS TRANSCEND BIOLOGY* (2005).
 - 5 NICK BOSTROM, *SUPERINTELLIGENCE: PATHS, DANGERS, STRATEGIES* (2014).
 - 6 Credible experts known to espouse this view include Stewart Russell and Sam Harris. See, e.g., Natalie Wolchover, *This Artificial Intelligence Pioneer has a Few Concerns*, WIRED (May 23, 2015), <https://www.wired.com/2015/05/artificial-intelligence-pioneer-concerns/> (interviewing Stewart Russell on developing AI that is provably aligned with human values). See also Sam Harris, *Can We Avoid A Digital Apocalypse?: A Response to the 2015 Edge Question*, SAM HARRIS (Jan. 16, 2015), <https://samharris.org/can-we-avoid-a-digital-apocalypse/>.

(and incredible⁷) experts and governmental bodies⁸ are pressing us to look beyond the social implications of today's AIs and robots. Instead of the focus being squarely on how human rights like equality and privacy are affected by rapid technological advance, undue attention is being paid to fantastical ideas. These include an all-out robot apocalypse and — less traumatic but still highly problematic — granting legal status to robots. Many policymakers seem fascinated by the idea of robot rights, or other protections and entitlements to incentivize and facilitate an increasing population of robots and AIs.⁹

On the privacy side, one sees the opposite: on their own, robots and AIs are nothing. According to the received view, there can be no loss of privacy without human sentience or cognition. Since robots and AIs are neither sentient nor capable of human-level cognition, they are seen to be of no consequence to privacy law. Robots and AIs can collect, use, disclose, make

7 Incredible experts (*i.e.*, famous experts in other fields) known to espouse this view include Stephen Hawking, Elon Musk, and Bill Gates. *See, e.g.*, Quincy Larson, *A Warning from Bill Gates, Elon Musk, and Stephen Hawking*, MEDIUM (Feb. 18, 2017), <https://medium.freecodecamp.org/bill-gates-and-elon-musk-just-warned-us-about-the-one-thing-politicians-are-too-scared-to-talk-8db9815fd398>.

8 Should We Fear Artificial Intelligence?: In-depth Analysis, EUR. PARL. DOC. PE 614.547 (2018), [http://www.europarl.europa.eu/RegData/etudes/IDAN/2018/614547/EPRS_IDA\(2018\)614547_EN.pdf](http://www.europarl.europa.eu/RegData/etudes/IDAN/2018/614547/EPRS_IDA(2018)614547_EN.pdf). *See also* At a Glance: Civil Law Rules on Robotics, EUR. PARL. DOC. PE 599.250 (2017), [http://www.europarl.europa.eu/RegData/etudes/ATAG/2017/599250/EPRS_ATA\(2017\)599250_EN.pdf](http://www.europarl.europa.eu/RegData/etudes/ATAG/2017/599250/EPRS_ATA(2017)599250_EN.pdf).

9 Saudi Arabia granted the world's first robot citizenship to "Sophia," a robot produced by Hanson Robotics. While a publicity stunt, Sophia's "citizenship" was met with criticism from global women's rights advocates and many Saudis alike, aghast that a "female" robot, who does not wear a head covering and need not be accompanied in public by a male relative, arguably has more rights than human Saudi Arabian women. *See* Dom Galeon, *Saudi Arabia Made a Robot a Citizen. Now, She's Calling for Women's Rights*, FUTURISM (Dec. 15, 2017), <https://futurism.com/saudi-arabia-made-robot-citizen-calling-womens-rights/>. Meanwhile, in Tokyo, Japan, official residency was granted to "Shibuya Mirai," an AI chatbot that exists only on the popular "Line" messaging app. *See* Patrick Caughill, *An Artificial Intelligence Has Officially Been Granted Residency*, FUTURISM (Nov. 6, 2017), <https://futurism.com/artificial-intelligence-officially-granted-residency/>. For more on the discussion of whether robots need their own rights regime, see generally Mark Coeckelbergh, *Robot rights? Towards a social-relational justification of moral consideration*, 12 ETHICS INFO. TECH. 209 (2010). *See also* Ryan Abbott, *I Think, Therefore I Invent: Creative Computers and the Future of Patent Law*, 57 B.C. L. REV. 1079 (2016) (on the topic of whether computers that invent things should be listed on patents).

decisions about and act upon exabytes of personal information but, from a doctrinal perspective, none of that matters, not one single bit, as long as no human has laid eyes on the data.

Without invoking the Copenhagen Interpretation,¹⁰ this Article offers a refutation of dead-or-alive, all-or-nothing accounts of robots and privacy. It is my contention that current robots and AIs can diminish our privacy without sentience, consciousness or cognition, and without human intervention, oversight, knowledge, or awareness. Building on an epistemic theory of privacy, I demonstrate that today's robots and AIs are capable of truth-promoting belief formation processes, thereby allowing them to form reliable beliefs and observational knowledge about people without human intervention, oversight, knowledge, or awareness. Because machines can actuate on the basis of the beliefs they form, they can affect people's life chances and opportunities in ways that definitively implicate privacy.

To be clear, the rather modest claim I am advancing in this short Article is that non-sentient robots and AIs can diminish our privacy. The Article is meant to say very little about how people perceive privacy violations by robots. It says even less about the normative elements of human-robot privacy relationships and the violations, infringements, wrongs, or harms that could be occasioned or avoided by robots. Although my argument is a necessary precondition for such discussions, my focus here is limited to the epistemological conditions giving rise to privacy, and my narrow claim is that some robots and AIs are already capable of epistemological states that can reduce our privacy. A proper account of the deeper normative elements would require a full-blown relational theory of robots, which this Article seeks to encourage, but does not strive to accomplish.

The Article proceeds as follows. In Part I, I argue that privacy is relational and briefly examine several key theories in order to establish privacy's relational core, namely: a person loses privacy just in case some "other" gains some form of epistemic access to her. Part II offers a closer examination of the "other"¹¹ in a privacy relationship — historically conceived as the person who comes to know personal facts about a data subject. I demonstrate how

10 In quantum theory, the Copenhagen Interpretation is the notion that a quantum particle doesn't exist in one state or another, but instead exists in all of its possible states at once. First posited by Niels Bohr in 1920, Schrödinger's thought experiment provided theoretical support for this idea. *See generally* Jan Faye, *Copenhagen Interpretation of Quantum Mechanics*, in *STANFORD ENCYCLOPEDIA OF PHILOSOPHY* 10 (Edward N. Zalta ed., 2014), <https://plato.stanford.edu/entries/qm-copenhagen/>.

11 These days usually referred to as the "data recipient."

robots and AIs are replacing the human “other” and that the delegation of informational transactions to robots and AIs therefore puts the traditional privacy relationship in an uncertain state. The uncertainty rests on whether an AI has the epistemic qualities necessary to diminish privacy in cases where there is no human intervention, oversight, knowledge, or awareness. Part III responds to the doctrinal view that individuals whose information is exposed only to automated systems incur no cognizable loss of privacy. To do so, I borrow from epistemic privacy — a theory that understands a subject’s state of privacy as a function of another’s state of cognizance regarding the subject’s personal facts. The theory of epistemic privacy offers a useful analytic framework for understanding the kind of *cognizance* that implicates privacy. In Part IV, I apply the theory of epistemic privacy in order to determine whether *artificial* cognizers are truly ignorant in the way that legal doctrine suggests. To the contrary, I conclude that artificial cognizers can be said to form truth-promoting beliefs that are justified. In Part V, I examine how today’s navigational robots form beliefs and argue that the observational knowledge they acquire through reliable belief formation processes easily meets the epistemic conditions necessary for diminished privacy. I suggest that, because the beliefs generated by artificial cognizers can also be programmed to actuate automatically, not only can they diminish a person’s state of privacy, they also have the potential to violate it. Having shown that today’s robots are by no means ignorant, I propose in Part VI the need to develop a theory of relational privacy that counts robots and AIs as integral to the configuration of what Julie Cohen has called the “networked self,”¹² not to mention what I am calling the “networked other.” In Part VII, I conclude with the observation that when we view epistemic privacy’s notion of “a duty of ignorance” through a relational lens, we are led toward a useful heuristic for our increasingly complex web of human-robot relationships: a presumption of ignorance.

This Article is meant to demonstrate how robots and AIs disturb the presumption of ignorance in epistemologically significant ways, undermining the presumption’s core aim of providing fair and equal treatment to all by setting boundaries around the kinds of assumptions and beliefs that can and cannot be made about people. Consequently, it is my contention that legal theory and doctrine will have to expand their understandings of privacy relationships to include robots and AIs that meet these epistemic conditions. An increasing number of machines possess epistemic qualities that force us to rethink our understanding of privacy relationships.

12 JULIE E. COHEN, CONFIGURING THE NETWORKED SELF: LAW, CODE, AND THE PLAY OF EVERYDAY PRACTICE (2012).

I. PRIVACY IS RELATIONAL

Privacy is protean.¹³ Constituted of various “family resemblances”, wise counsel tells us that we are best not to try to locate privacy in any one theory.¹⁴ But, if we were in search of a common denominator, certainly it would be that privacy is relational. As David Matheson has put it, “[a]n individual’s informational privacy is possessed or lacked only in relation to other individuals and to personal facts about her.”¹⁵

A relational core can be found — implicitly or explicitly — at the very foundation of privacy’s best-known theories. For example, one finds it in the classical *control theory* proposed by Alan Westin,¹⁶ Charles Fried,¹⁷ and others. Westin famously defined privacy in relational terms: “the claim of individuals, groups, or institutions to determine for themselves when, how, and to what extent information about themselves is communicated to *others*.”¹⁸ If one further considers Westin’s four states of privacy, *i.e.*, solitude, intimacy, anonymity, and reserve, one sees in each of them the centrality of the relation between the self and others. Solitude represents the physical separation of an individual from *others*. Intimacy likewise involves the seclusion of couples or groups from *others*. Anonymity permits individuals to participate in public

-
- 13 R. v. Tessling, [2004] 3 S.C.R. 432, ¶ 25 (Can.). The term “protean” is derived from the marine deity Proteus: in Greek mythology, Proteus was the son and servant of the sea god, Poseidon. At his father’s request, he acted as shepherd of the sea creatures and was rewarded with the gift of prophecy. Proteus therefore knew all about the past, present, and future, but disliked sharing his knowledge. When asked probing questions, he would change forms into different animals or dragons to avoid interrogation. *See generally Proteus: Greek Mythology*, ENCYCLOPAEDIA BRITANNICA, <https://www.britannica.com/topic/Proteus-Greek-mythology>.
 - 14 Solove writes: “Wittgenstein uses the term ‘family resemblances,’ analogizing to the overlapping and crisscrossing characteristics that exist between members of a family, such as ‘build, features, colour of eyes, gait, temperament, etc.’” *See* Daniel J. Solove, *Conceptualizing Privacy*, 90 CALIF. L. REV. 1087 (2002), citing LUDWIG WITTGENSTEIN, *PHILOSOPHICAL INVESTIGATIONS* §§66-67 (Gertrude Elizabeth Margaret Anscombe trans., 1958).
 - 15 David Matheson, *Unknowableness and Informational Privacy*, 32 J. PHIL. RES. 251 (2007).
 - 16 ALAN F. WESTIN, *PRIVACY AND FREEDOM* (1967).
 - 17 Charles Fried, *Privacy*, in *PHILOSOPHICAL DIMENSIONS OF PRIVACY: AN ANTHOLOGY* 203-22 (Ferdinand David Schoeman ed., 1984).
 - 18 WESTIN, *supra* note 16, at 7 (emphasis added).

participation without identification by *others*. And, reserve enables the creation of a psychological barrier against the unwanted intrusion of *others*.¹⁹

The relational aspect of privacy is also front and center in *limited access theory*, as articulated by proponents such as Ruth Gavison²⁰ and Anita Allen.²¹ As Gavison puts it, “our interest in privacy . . . is related to our concern over our accessibility to *others*: the extent to which we are known to *others*, the extent to which *others* have physical access to us, and the extent to which we are the subject of *others*’ attention.”²² Here, one loses privacy just in case another gains some form of epistemic access to her. Irwin Altman similarly depicted privacy as “a central regulatory process by which a person (or group) makes himself more or less accessible and open to *others* . . .”²³

The recognition that privacy is relational — though practically indisputable — does little on its own to illuminate privacy theory. Some of the most interesting privacy scholarship in recent years has therefore sought to develop a more fulsome relational account of privacy. Although the moniker of *relational privacy* has been used to mean different things,²⁴ I use it here to refer broadly to any account of privacy that coincides, borrows, or builds upon broader relational theory in the tradition of Jennifer Nedelsky and other feminist thinkers who “insist on the centrality of relationships in human lives” and recognize that privacy values “require structures of relationships that support them — that allow people the opportunities to retreat from others in various ways.”²⁵ On this view, “people are not self-made.”²⁶ Instead, they are constituted by networks of nested relationships. Their autonomy “can thrive or wither . . . depending on the structures of relationships they are embedded in.”²⁷

19 *Id.* at 31 (emphasis added).

20 Ruth Gavison, *Privacy and the Limits of Law*, 89 YALE L.J. 421 (1980).

21 ANITA ALLEN, *UNEASY ACCESS: PRIVACY FOR WOMEN IN A FREE SOCIETY* (1988).

22 Gavison, *supra* note 20, at 423 (emphasis added).

23 IRWIN ALTMAN, *THE ENVIRONMENT AND SOCIAL BEHAVIOR* 3 (1975) (emphasis added).

24 Louis Sacharoff, *The Relational Nature of Privacy*, 16 LEWIS & CLARK L. REV. 1250 (2012); *See also* Jean L. Cohen, *Rethinking Privacy: The Abortion Controversy*, in *PUBLIC AND PRIVATE IN THOUGHT AND PRACTICE: PERSPECTIVES ON A GRAND DICHOTOMY* 143, 149 (Jeff Alan Weintraub & Krishan Kumar eds., 1997).

25 JENNIFER NEDELSKY, *LAW’S RELATIONS: A RELATIONAL THEORY OF SELF, AUTONOMY, AND LAW* 35 (2011).

26 *Id.* at 8. Referring, of course, to the so-called ‘self-made man’ — the gendered instantiation of the autonomous liberal individual.

27 *Id.* at 39.

Stuart Hargreaves uses relational theory as described above to articulate an intellectually rich account of what he refers to as “relational privacy.”²⁸ Following Nedelsky and others, Hargreaves believes that a true state of privacy can be meaningfully achieved only within a dense network of relationships. Consequently, Hargreaves believes that a privacy loss is something that lessens our ability to engage others or modulate our exposure within the network of relationships, and that the most serious privacy losses are those that seriously impinge upon the ability of individuals to function as members of a community in ways they see fit.²⁹ Departing from James Rachels,³⁰ Hargreaves’ relational privacy does not entail the creation of a ‘zone’ in which the individual is protected from intrusion in order that they can develop certain kinds of relationships. Rather, Hargreaves proposes that a proper relational account will attempt to measure the seriousness of a privacy violation in terms of the harm it does to the web of relationships within which any individual finds herself.

Relational privacy in this sense borrows quite heavily — consciously and unconsciously — from Julie Cohen’s *Configuring the Networked Self*, which was written around the same time as Nedelsky’s *Law’s Relations*. In her work, Cohen proposes a decentered model of subjectivity and recognizes “the situated subject’s perspective, the collective dimension of subjectivity, and the play that overlapping social and cultural networks afford.”³¹ With liberalism’s autonomous individual no longer center stage, Cohen quite intentionally avoids classical references to the self and the other, preferring instead the “self-society relation”³² or the “interaction between self and culture.”³³ By shifting the relational lens away from self and other in the classical liberal sense, Cohen depicts the networked self as a (partial) function of many others. Each of those others is itself a part of an evolving network of subjectivity. In this sense, Cohen offers an important lever to a deeper understanding of the privacy relationship not previously articulated. Although Cohen does not state it explicitly, the tradition makes precisely the same error when it comes

28 Stuart Hargreaves, ‘Relational Privacy’ & Tort, 23 WM. & MARY J. WOMEN & L. 433 (2017).

29 *Id.* at 437.

30 James Rachels, *Why Privacy is Important*, in PHILOSOPHICAL DIMENSIONS OF PRIVACY: AN ANTHOLOGY, *supra* note 17, at 292.

31 JULIE E. COHEN, CONFIGURING THE NETWORKED SELF: LAW, CODE, AND THE PLAY OF EVERYDAY PRACTICE 127 (2012).

32 *Id.* at 118. Hargreaves similarly understands “privacy as an on-going, negotiable relationship between the Self and the Other (whether that is between two individuals, or the individual and the community, or between groups and the community).” Hargreaves, *supra* note 28, at 463.

33 COHEN, *supra* note 31, at 128.

to understanding privacy's other as it did in its conceptualization of the self. Privacy's other is no more or less a romantic, autonomous individual than is the self. Privacy's other — it is at least implied — must also be understood within a dense network of relationships.

Some scholars have paid careful attention to what is happening at the other end of a privacy relationship. Helen Nissenbaum is one of them. Her theory of privacy as contextual integrity is inextricably linked to the spiderweb-like relations of our lives. For example, "according to the theory of contextual integrity it is crucial to know the context — who is gathering the information, who is analyzing it, who is disseminating it and to whom, the nature of the information, the relationships among the various parties, and even larger institutional and social circumstances."³⁴

II. PRIVACY'S OTHER

That privacy entails a certain kind of relationship between self and other has been well acknowledged, while many of Nissenbaum's finer points about the nature and context of privacy's other have perhaps gone undertheorized until very recently.³⁵ With much of privacy theory and doctrine attendant to the data subject or to the data itself,³⁶ less has been said about the necessary and sufficient conditions of the privacy relationship *vis-à-vis* the other.³⁷

34 HELEN NISSENBAUM, *PRIVACY IN CONTEXT: TECHNOLOGY, POLICY AND THE INTEGRITY OF SOCIAL LIFE* 154-55 (2010).

35 Andrew Selbst, *Contextual Expectations of Privacy*, 35 *CARDOZO L. REV.* 643 (2013).

36 Whether it is 'personal' or 'sensitive' data; how that data has been collected, used, or disclosed, and so forth.

37 I sometimes refer to privacy's other as "the data recipient" and, later in the Article, as the "cognizer." Modern privacy legislation recognizes the importance of the data recipient, distinguishing between situations where it is an individual, corporation, organization, or government. Canada's federal private-sector privacy legislation applies only to "the organization [that] collects, uses or discloses in the course of commercial activities," or "is about an employee of, or an applicant for employment with, the organization and that the organization collects, uses or discloses in connection with the operation of a federal work, undertaking or business." This excludes regulation of individuals, corporations or governments, and those who collect, use, or disclose personal information for "journalistic, artistic or literary purposes and does not collect, use or disclose for any other purpose." See *Personal Information Protection and Electronic Documents Act*, S.C. 2000, c 5, § 4(1), 4(2) (Can.).

To some extent, this is because modern privacy theory — like contract before it — was constructed on a paradigm of a face-to-face (or, at least a person-to-person) interaction.³⁸ Under this paradigm, the other (party) was usually known, or knowable, and the interpersonal relationship itself was generally comprehended. Historically, as Matthew Tokson has observed, “our conception of a loss of privacy is bound up with the presence of a human observer.”³⁹

But the old road is rapidly aging. These days, the privacy relationship is primarily understood as the result of informational transactions wherein a data subject often has no real clue about the other or how she/he/they/it are using the data. Indeed, following in the footsteps of contract law, it is getting harder and harder to understand privacy as an interpersonal relationship at all.⁴⁰

Starting with contract, since at least the advent of the standard form agreement, a proper understanding of the contractual relationship as interpersonal has become strained.⁴¹ Standard forms — especially alongside the invention of the unilateral offer⁴² — have made it difficult if not impossible to identify *who*, if anyone, actually agreed to or accepted the other side of the contract, or *where* exactly to locate the so-called *consensus ad idem* upon which the contractual relationship is said to be founded. Although automated standard forms are easily understood and recognized as enforceable, the justification for

38 The law of contract is classically defined as the “manifestation of a mutual concordance between [two] parties as to the existence, nature and scope of their rights and duties.” See G.H.L. FRIDMAN, *THE LAW OF CONTRACT* 5 (5th ed. 2006). From an American perspective, Arthur Corbin defines the law of contract similarly: “The law of contract deals with those legal relations that arise because of mutual expressions of assent. The parties have expressed their intentions in words, or in other conduct that can be translated into words.” See Arthur L. Corbin, *Conditions in the Law of Contract*, 28 YALE L.J. 739, 740 (1919).

39 Matthew Tokson, *Automation and the Fourth Amendment*, 96 IOWA L. REV. 581, 611 (2010).

40 At least where “interpersonal” is meant to designate two individual human beings engaged in a single, definable transaction.

41 FRIDMAN, *supra* note 38, at 5. See also Todd D. Rakoff, *Contracts of Adhesion: An Essay in Reconstruction*, 96 HARV. L. REV. 1173, 1284 (1983).

42 In a unilateral contract, an act is exchanged for a promise. When the contract comes into existence, only the offeror has obligations to fulfill: the offeree has already fulfilled its obligations simply by accepting the offer. See MITCHELL MCINNES, IAN KERR & ANTHONY VANDUZER, *MANAGING THE LAW* 46 (2d ed. 2006). Or, as Corbin puts it: “A unilateral contract is one where only one of the parties assumes a contractual duty and only the other acquires any contractual right.” Corbin, *supra* note 38, at 750.

their enforcement has drifted a long way from the interpersonal transactions that were the theoretical foundation of contracts in days gone by. Many, if not most, of today's contracts are better understood on the basis of attribution than actual agreement.⁴³

Like contracts, privacy relationships are also becoming more one-sided. In both instances, there is considerable uncertainty about what exactly is taking place on the other side of the transaction. Although we are quite used to the idea of a data subject transacting with artificial entities such as corporations, in theory, we still require there to be someone who stands in relation to the data subject according to conditions that are said to give rise to a privacy relationship. Privacy, like contract, desperately tries to achieve this through the reductive recognition that organizations are comprised of individuals and then by attributing some individual to the collection, use, or disclosure with which privacy is concerned.⁴⁴

This approach, however, is much more plausible in the case of contract. Although there may be uncertainty about the other contracting party insofar as *who* exactly agreed to it, the subject matter of a contractual transaction — namely the agreement itself — is usually well known and almost always knowable. There is less justification for the use of an attribution principle in the privacy context, where the subject matter of informational transactions so often take place in a “black box.”⁴⁵ Even in a so-called consent-based interaction, the data subject does not actually know who, if anyone, sits on the other side of the relationship, nor whether, or to what extent, personal information has in fact been collected, used, or disclosed.

43 See generally Ian R. Kerr, *Spirits in the Material World: Intelligent Agents as Intermediaries in Electronic Commerce*, 22 DALHOUSIE L.J. 189 (1999).

44 David Matheson, *An Obligation to Forget*, in THE ROUTLEDGE HANDBOOK OF PHILOSOPHY OF MEMORY 369 (Sven Bernecker & Kourken Michaelian eds., 2017). What distinguishes corporations from robots and AIs is the very fact that corporate actions are *always* reducible to human activities, while emergent behavior in robots and AIs is neither reducible to — nor properly understood by — the correlative investigation of human conduct or awareness. One might say that the whole point of machine learning and other techniques in artificial intelligence is to generate decisions and actions different (and in some way better than) those that would have been produced by a human actor or an organization of humans in the form of a corporation. Consequently, this Article very intentionally does not discuss corporations or other forms of human enterprise. The concerns being addressed here engage the epistemological implications of actions or decisions totally unknown to *any* human actors.

45 FRANK PASQUALE, *THE BLACK BOX SOCIETY: THE SECRET ALGORITHMS THAT CONTROL MONEY AND INFORMATION* (2015).

A. In Automated Transactions, the Conditions for Privacy Reside in Uncertain States.

Privacy's corollary to the automation of contract through standard forms and unilateral offers increasingly takes place though the instrumentality of various forms of automation, data analytics and, more recently, artificial intelligence (AI).⁴⁶ Sophisticated AI and robotic technologies can sense, process, and act upon the world with increasing agency.⁴⁷ This blurring of instrument and actor has caused some to see AI as revolutionary — a field in which the promises of science fiction and current capabilities of science are likewise blurred.⁴⁸ Despite what some would like us to consider, AIs are not capable of consciousness and have not advanced to the realm of “superintelligence” that is so often heralded as the future. Despite this reality, historical debates about “strong” versus “weak” AI⁴⁹ continue to give way to the quest for AGI, which assumes a “human-like” intelligence with broad, multi-hyphenate capabilities.⁵⁰

46 It is of course also true that standard forms and unilateral offers are also used as an engine in obfuscation and circumvention of privacy's other.

47 Technology theorist and legal scholar Mireille Hildebrandt uses the phrase “data-driven agency” to describe the apparent agency of modern smart technologies, which are empowered through knowledge discovery in databases to profile, predict, and possibly even preempt human behaviors. Through the additive capacity of algorithms, compiling consumer data, smart technologies can develop the “autonomy” necessary to transcend mere programming to actually respond to human needs, or, in other words, exhibit agency. See MIREILLE HILDEBRANDT, SMART TECHNOLOGIES AND THE END(S) OF LAW: NOVEL ENTANGLEMENTS OF LAW AND TECHNOLOGY 19-31 (2015). See also Kerr, *supra* note 43.

48 Bill Smart and Neil Richards caution against commission of “The Android Fallacy”: wrongly anthropomorphizing robots in a manner consistent with science fiction. They argue that an important part of proposing legal regulation for AIs is sketching what robots can do, what they cannot do yet, and what they may never be able to do. Recognizing that “never” is a long time, they suggest that “long after we are dead” is a sufficient standard. See Neil M. Richards & William D. Smart, *How Should the Law Think About Robots*, in ROBOT LAW 3 (A. Michael Froomkin, Ryan Calo & Ian Kerr eds., 2015).

49 Calo, *supra* note 2.

50 Daniel Schönberger, *Deep Copyright: Up - And Downstream Questions Related to Artificial Intelligence (AI) and Machine Learning (ML)*, 10 ZEITSCHRIFT FÜR GEISTIGES EIGENTUM [INTELLECTUAL PROPERTY JOURNAL] 35 (2018), <https://ssrn.com/abstract=3098315>.

This Article focuses exclusively on the current, narrow AIs — not AGIs.⁵¹ Current incarnations of robots and AIs can be used as delegates in informational transactions that would have previously required a personal relationship. The difference between AIs and other artificial entities that stand in for a human observer is that some AIs can function completely independently of human oversight, intervention, or knowledge. Consequently, a robot's interaction with a human cannot easily be understood as part of a traditional privacy relationship in situations where no other human being ever comes into play. Even the most outlandish legal fiction of attribution would have nobody to attach itself to when the data subject stands completely alone. In these twenty-first-century informational transactions, there is no human observer as historically presumed in the privacy relationship, and yet the independent operations of an AI or robot can generate devastating consequences that radically alter and diminish data subjects' life chances and opportunities. To take one example, many AI and privacy experts are concerned about the desire in some countries to develop and deploy autonomous robot police or private security guards capable of conducting surveillance and making decisions about people that could result in the projection of lethal force.⁵² These same algorithms could be used as access control providers (gatekeepers), as a means of detaining a subject without a warrant, or to conduct uber-surveillance, *etc.*

The latter possibility was brought to life in June 2013, when the National Security Agency's PRISM (upstream data collection program)⁵³ began reading,

51 As Mireille Hildebrandt brilliantly points out, there is a risk in conceiving robots or AIs as standalone entities. In doing so, we divert attention from the fact that they are only as smart as the institutional clouds of knowledge and information which not only permit but largely determine their operations. We must always keep in mind that worries about robots and privacy are directly related to our existing worries about public- and private-sector big data institutions.

52 Peter Asaro, *Will #BlackLivesMatter to Robocop?*, WE ROBOT (Mar. 1, 2016), <http://robots.law.miami.edu/2016/peter-asaro-on-will-blacklivesmatter-to-robocop/>.

53 PRISM is the U.S. government's surveillance program, which enables the National Security Agency (NSA) and the Federal Bureau of Investigation (FBI) to tap directly into the servers of nine American Internet companies to surveil chat logs, photographs, emails, documents and connection logs. It uses extracted data from Microsoft, Yahoo, Google, Facebook, PalTalk, AOL, Skype, YouTube, and Apple. Originating in 2007 through homeland security legislation under President G.W. Bush, PRISM was top secret until NSA contractor Edward Snowden leaked PowerPoints confirming its operation in June 2013. British intelligence agencies were also accessing the data. See Barton Gellman & Laura Poitras, *U.S., British Intelligence Mining Data from Nine U.S. Internet Companies in Broad Secret Program*, THE WASH. POST (June 7, 2013), <https://www.washingtonpost.com/>

processing, and responding to exabytes⁵⁴ of private communications — mostly without human awareness. Political rhetoric described this merely as “bulk collection.” Because robots and AIs could be delegated the observation task, there was generally no need to keep humans involved in the information collection and processing loop, and therefore there was no “mass surveillance.” Robots and AIs could figure out who the bad guys are so that only criminal suspects or persons of interest would become subject to human scrutiny.⁵⁵ In this way, the use of robots and AI, it was suggested, was not privacy diminishing but, in fact, privacy enhancing. With increasing momentum in so-called bulk collection in both public and private sectors during the five years since — not to mention significant increases in storage capacity, processing power and the sophistication of machine learning techniques — the issue is by now pressing and substantial. Is privacy implicated when only robot eyes are watching,⁵⁶ or does the privacy relationship ultimately require a human observer?

B. Can a Robot or AI Diminish Privacy without Human Intervention, Oversight, Knowledge, or Awareness?

The question is admittedly strange. As stated previously, this is because privacy theory has always had the luxury of focusing on the nature of the data and its subject rather than privacy's other. It has always simply been presumed that privacy's other is a human observer.

A stark example of this presumption is Judge Posner's laser bright line: “[c]omputer searches do not invade privacy because search programs are not

investigations/us-intelligence-mining-data-from-nine-us-internet-companies-in-broad-secret-program/2013/06/06/3a0c0da8-cebf-11e2-8845-d970ccb04497_story.html?utm_term=.46198d242534.

54 One exabyte is the equivalent of 1.5 billion CD-ROM disks.

55 Glenn Greenwald, *The Orwellian Re-Branding of “Mass Surveillance” as Merely “Bulk Collection”*, THE INTERCEPT (Mar. 13, 2015, 2:23 PM), <https://theintercept.com/2015/03/13/orwellian-re-branding-mass-surveillance-merely-bulk-collection/>. See also Glyn Moody, *What's The Difference Between 'Mass Surveillance' And 'Bulk Collection'? Does It Matter?*, TECHDIRT (Jan. 20, 2016, 3:23 AM), <https://www.techdirt.com/articles/20160115/09582933351/whats-difference-between-mass-surveillance-bulk-collection-does-it-matter.shtml>.

56 For an outstanding articulation of the doctrinal answer to this question, see Kevin S. Bankstown & Amie Stepanovich, *When Robot Eyes Are Watching You: The Law & Policy of Automated Communications Surveillance*, Paper presented at the We Robot Conference hosted at the University of Miami (Apr. 5, 2014), http://robots.law.miami.edu/2014/wp-content/uploads/2014/07/Bankston_Stepanovich_We_Robot.pdf.

sentient beings. Only the human search should raise constitutional or other legal issues.”⁵⁷

Judge Posner’s view is supported by the majority of jurists and scholars who have considered the matter. For example, in the context of federal wiretap law, Bruce Boyden argues that “interception requires at least the potential for human awareness of the contents. . . . Under most theories of privacy, automated processing does not harm privacy.”⁵⁸ Matthew Tokson describes such scenarios as “disclosure to . . . automated third parties” and concludes that “individuals whose information is exposed only to automated systems incur no cognizable loss of privacy.”⁵⁹

But is it true that when privacy’s other is an *AI* or *robot*, there is *no* cognizable loss of privacy?

III. EPISTEMIC PRIVACY

Whether there has been a “cognizable loss of privacy” can be understood in more than one sense. When Tokson, Boyden, and Posner say that there is no cognizable loss of privacy for individuals whose information is exposed only to automated systems, they mean it in a legal sense and say so on the basis of legal doctrine. Kevin Bankston and Amie Stepanovich argue the very contrary position, but on precisely the same grounds. On either side, a “cognizable loss of privacy” is interpreted as something that is determined by a court. But, there are other ways of addressing the question.

Whether a robot or AI can generate a cognizable loss of privacy can also be approached as an epistemological inquiry. Understood in this sense, the question to be asked is whether a robot or AI is capable of meeting the epistemic conditions giving rise to diminished privacy. Responding to this question, I turn to a relatively new body of theory that has not yet acquired mainstream recognition: epistemic privacy. This nascent area of privacy scholarship sees an intuitive connection between privacy and knowledge: “one main reason why one might want to protect one’s privacy is that one doesn’t want certain others to acquire knowledge of certain private facts about oneself.”⁶⁰

57 Richard A. Posner, *Privacy, Surveillance, and The Law*, 75 CHI. L. REV. 245, 254 (2008).

58 Bruce E. Boyden, *Can A Computer Intercept Your Email?*, 34 CARODOZO L. REV. 669, 716 (2012).

59 Tokson, *supra* note 39, at 581.

60 Martijn Blaauw, *Introduction: Privacy, Secrecy and Epistemology*, 10 EPISTEME 99 (2013).

Ironically, the acorns of epistemic privacy were first germinated when the much respected knowledge theorist, Alvin Goldman, tried to separate these two modes of inquiry:

Important as [privacy] is, it does not squarely fall into the domain of epistemology as I have delineated it, because epistemology focuses on the means to knowledge enhancement, whereas privacy studies focus on the means to knowledge curtailment (at least decreasing knowledge in the hands of the wrong people). For this reason, I shall not explore this topic. I do not belittle the importance of privacy as a moral issue; it simply falls, for the most part, outside the scope of epistemology.⁶¹

The idea of “decreasing knowledge in the hands of the wrong people” certainly lends itself well to some of privacy’s most important theories. To take one example, it seems to go hand in hand with Anita Allen’s conception of privacy, “denoting limited access or degrees of inaccessibility of persons, their mental states, and information about them to others.”⁶² Allen’s restricted-access definition offers a very clear linkage between the condition of privacy and the epistemic state of the other:

To say that a person possesses or enjoys privacy is to say that, in some respect and to some extent, the person (or the person’s mental state, or information about the person) is beyond the range of others’ five senses and any devices that can enhance, reveal, trace, or record human conduct, thought, belief, or emotion.⁶³

Although, in 1988, Allen was thinking primarily about the other as a human, she was light years ahead in recognizing the extent to which human cognition is augmented by technology, and she was crystal clear in alerting readers to the fact that our devices play a crucial role in the means by which we gain epistemic access to others.

Perhaps the central contribution of epistemic privacy — something that had not previously been emphasized in the privacy literature — is the recognition that most privacy theories “render an individual’s informational privacy at least partly a function of a kind of inability of others to know personal facts about her.”⁶⁴ In other words, someone not knowing something is a requirement for

61 ALVIN L. GOLDMAN, *KNOWLEDGE IN A SOCIAL WORLD* 173 (1999).

62 ALLEN, *supra* note 21, at 14.

63 *Id.* at 15.

64 Matheson, *supra* note 15. Matheson’s aim is to address deficits in this common formulation by offering a new account that would render an individual’s informational privacy exclusively a function of others’ ignorance of personal facts about her.

someone else having privacy about that thing and, equally, someone knowing something is a requirement for someone else losing privacy about that thing.⁶⁵

Matheson has further considered what epistemic privacy might entail alongside someone's claim to privacy: "Conjoined with the plausible claim that there is a moral right to privacy, each of the major contemporary accounts of the nature of privacy . . . implies a duty of ignorance."⁶⁶

Understanding one's state of privacy as a function of another's ignorance and one's right to privacy as giving rise to a duty of that other to remain ignorant

65 As discussed further below, I agree with those epistemologists who argue that this account sets the bar too high because "you can lose privacy even if someone has epistemic access to facts about you that falls short of knowledge." See Don Fallis, *Privacy and Lack of Knowledge*, 10 *EPISTEME* 153 (2013). Regardless of whether the appropriate epistemic threshold is knowledge or some lower epistemic standard (e.g., reliable beliefs), it is important to draw a careful distinction between the level of epistemic awareness and what someone (or something) might do with it. It is the doing (or not doing) that usually results in a privacy violation and possible privacy harms. My current epistemological investigation focuses on how the epistemological status (of knowing, believing, *etc.*) impacts the *state* of privacy, only briefly mentioning the additional consideration of what happens when that knowledge (belief, *etc.*) is (or is not) acted upon. For more on the latter, consider Lisa Austin's insightful discussion in this volume of Alan Westin's account of the norms of tact and discretion. See Lisa M. Austin, *Re-reading Westin*, 20 *THEORETICAL INQUIRIES L.* 53 (2019). These norms sometimes require knowers to act as though they do not know something, or to act as though they do not know that the data subject knows they know. (It's very Rumsfeldian!) If these others act in the appropriate manner, they will not violate privacy and, arguably, may not reduce the data subject's privacy by much — depending on who else comes to know the relevant facts about the data subject and, for example, whether they interfere with the data subject's self-presentation. As Austin points out, the kind of relationships between the various knowers and even the particular relational approach to privacy taken may make an important difference. It is interesting to think about whether Westin's norms of tact and self-discretion would play out the same way if the 'knower' was a machine. That would require a much more fulsome relational account of robots than I have set out to achieve here. Thanks to Lisa Austin for the useful point of connection and for highlighting the need for greater clarification.

66 David Matheson, *A Duty of Ignorance*, 10 *EPISTEME* 193, 194 (2013). It is worth noting that epistemic privacy is "an account of what it means to simply lose privacy rather than an account of what it means for privacy to be violated. In other words, it does not have anything to say about whether a person who has lost privacy has been harmed or wronged in some way." See also Fallis, *supra* note 65, at 156.

is useful in determining whether individuals whose information is exposed only to automated systems incur a cognizable loss of privacy. But before moving directly to that question, it is worth gaining a further appreciation of what epistemologists of privacy are investigating.

Like all of the privacy theorists discussed above, privacy epistemologists begin with the premise that privacy is relational. In the tradition of analytic philosophy, the structure of that relationship has been meticulously unpacked. As Martijn Blaauw suggests,⁶⁷ being in a state of privacy can be expressed as a ternary relation between **S** (a subject), **P** (a set of true propositions; sometimes called “personal facts”⁶⁸) and **I** (some other; often referred to in epistemology as a “cognizer”⁶⁹).

As such, the relationship can be represented by the following locution:

S has privacy about P with respect to I

Rearranged slightly, one can express the conditions for being in a state of privacy as:

67 Martijn Blaauw, *The Epistemic Account of Privacy*, 10 *EPISTEME* 167, 168 (2013).

68 These are what privacy legislation and privacy professionals usually refer to as personal information or personally identifiable information.

69 The term “cognizer” has a long history. It was used by both Aristotle and St. Thomas Aquinas in their respective attempts to understand the natures of knowledge and thought. Both discuss the identity between the knower and what is known — and whether knowledge in the external world can truly exist. This Aristotelian conception of “what understands and what is understood are the same” marks the emergence of the term. Stemming from the Latin *cognoscens*, it connotes both that which has sensation and that which has intellect. Aristotle posited that cognitive powers are mere potencies for the development of cognition: for actual cognition to occur, the things must impress their potency upon the cognizer, according to its potency to receive them. Relatedly, Aquinas suggested that the act of cognizing, even about an external thing, remains within the cognizer. In this way, what the cognizer knows about a thing and how that thing has been cognized are inseparable. See ARISTOTLE, *DE ANIMA* III.2 (425b26-28) (R.D. Hicks trans., 2008). See also Robert Pasnau, *ABSTRACT: The Identity of Knower and Known* (Apr. 25, 1996) (unpublished manuscript), <http://faculty.fordham.edu/klima/APAPasnau.htm>; STEPHEN L. BROCK, *THE PHILOSOPHY OF SAINT THOMAS AQUINAS: A SKETCH* (2015). Modern use of the term intentionally refers to a cognizer (instead of “knower”) so as to be neutral with respect to **I**’s epistemic state. Although some philosophers, including Anita Allen, suggest that the term “cognizer” refers exclusively to humans, I use the term in an intentionally neutral manner and without specific reference to **I**’s species (*i.e.*, cognizers could be nonhuman artifacts).

*S possesses informational privacy in relation to I about P if and only if I does not know P.*⁷⁰

Although most epistemologists working on privacy agree about the general structure of the relationship, there is disagreement about the second half of the bi-conditional proposition, namely, whether “knowledge” (on the part of I) is the appropriate epistemic quality for determining whether S has suffered a loss of privacy. First, as Blaauw points out, the state of one’s privacy is not an on-off switch. Rather, privacy is a *degree concept* that depends on several variables — including how many personal facts are known and not known, as well as the number of individuals who know and don’t know them.⁷¹ Expressing these refinements by way of our locution,

The more Ps that are known about S by I, the less privacy S has with respect to I.

Likewise,

The more Is that know P about S, the less privacy S has regarding P.

The converse of both variations on our locution is also true.

Perhaps even more important for the ultimate question about whether a robot or AI can generate a cognizable loss of privacy, Blaauw suggests:

[I]t would be wrong to maintain, as Matheson does, that *only knowledge* of personal propositions is capable of diminishing one’s privacy. Weaker epistemic states, as well, can diminish one’s privacy. Weaker epistemic states do represent some sort of connection to the true proposition. And however weak the connection may be, it can still impact on one’s degree of privacy.⁷²

Over a long history of engagement with theories of knowledge — stretching back at least as far as Plato’s *Theaetetus* and *Meno* — epistemologists have enumerated several types of epistemic relations in which a cognizer might stand *vis-à-vis* a particular proposition. Much thinking and thousands if not millions of written pages have been generated about this range of epistemic relations in gloriously painstaking, intricate, and practically unimaginable detail by way of millions of examples, counterexamples, counterfactuals conditionals, and so forth. A typical list of the range of epistemic states includes:

- (1) A mere belief that P

70 This basic formulation is known as the “Broad Ignorance Theory” and was first articulated by David Matheson. See Matheson, *supra* note 15, at 259.

71 Blaauw, *supra* note 67, at 170 (emphasis in original).

72 *Id.* at 173.

- (2) A true belief that P
- (3) A justified true belief that P
- (4) A deGettierized⁷³ true belief that P
- (5) A rational true belief that P
- (6) A warranted true belief that P
- (7) Knowledge that P
- (8) Certainty that P⁷⁴

These various epistemic relations can be understood on a continuum ranging from zero knowledge (levels 1-2) to complete certainty (level 8). The stronger an epistemic intrusion becomes the more diminished one's privacy is. But, on this view, complete certainty is unnecessary — even weaker epistemic states (level 3-5) have significant privacy-impacting potential. In other words, as explained in greater detail in Part IV, you can lose privacy even if a cognizer has epistemic access to facts about you that fall well short of knowledge.

The point here in Part III is not to enter the epistemological fray about the exact privacy consequences for each level of epistemic intrusion, nor to determine which level of epistemic intrusion constitutes the appropriate threshold for diminished privacy. Rather, my main takeaway is that epistemic privacy offers a useful analytic framework for understanding the kind of *cognizance*⁷⁵ that implicates privacy. Appreciating the importance of determining the exact

73 The Gettier problem refers to a philosophical query posed by Edmund Gettier regarding the justified true belief (JTB) account of knowledge. JTB proposes that knowledge consists of three components, justification, truth, and belief, all three of which must be met to satisfy a claim of knowledge. A Gettier-type problem, or a Gettierized problem, proposes a counterfactual to challenge this proposition: it attempts to illustrate that JTB accounts of knowledge can fail if an individual holds a belief which, while justified and believed to be true, turns out to actually be false. Because of deception, it is possible that truth is hard to come by. A deGettierized belief, therefore, is one that has considered the Gettier problem and ensured that the environment in which a justified true belief is held is one that allows for truth. For further discussion on this point, see E.J. Coffman, *Lenient Accounts of Warranted Assertability*, in *EPISTEMIC NORMS: NEW ESSAYS ON ACTION, BELIEF, AND ASSERTION* 32, 38-41 (Clayton Littlejohn & John Turri eds., 2014).

74 *Id.* at 171.

75 Don Fallis defines cognizance with a connection between the personal fact and the cognizer's belief that is sufficient for a loss of privacy. Channeling Alvin Goldman's *causal theory of knowledge*, Fallis holds that A loses privacy about a personal fact *p* with respect to S if "the fact that *p* is causally connected in an 'appropriate' way with S's believing *p*." Fallis, *supra* note 65, at 160-61. Others, including Martjin Blaauw, cast an even wider net, allowing that *any*

nature of the connection between personal facts and a cognizer's belief, the epistemologists of privacy have provided strong and convincing theoretical grounding for an argument that Posner's bright-line threshold of sentience or Boyden's and Tokson's single brushstroke requirement of human level cognition are misplaced.

IV. UNCERTAIN STATES

What, then, can theory tell us about whether artificial cognizers can meet the epistemic conditions for diminished privacy? It would seem that the epistemic conditions — like Schrödinger's cat — reside either dead or alive inside the robot's black box. And, we don't know which. How / Can we posit privacy in such uncertain states?⁷⁶

In the spirit of Schrödinger, one could “set up ridiculous cases”⁷⁷ that try to imagine what is going on inside the box. But aren't human cognizers in a similar predicament? As Wittgenstein keenly observed, “[i]f God had looked into our minds he would not have been able to see there whom we were speaking of.”⁷⁸ One might even say that AIs and robots are more easily known;⁷⁹ after all, their hardware and software is built, programmed and, for the most part, understood.⁸⁰ Further, one can apply the epistemological

type of epistemic relation (levels 3-8) would have privacy-impacting potential. Blaauw, *supra* note 67.

76 Eldar Haber suggests that the uncertain state of privacy is expressed in the lack of trust people have once their personal facts are no longer within their exclusive control (*i.e.*, stored within a robot or the cloud). Because neither the robot nor a third-party provider can guarantee perfect security, our privacy is simultaneously violated and not violated once others have the potential to access our personal facts. While Haber believes that privacy is in a state of superimposition as soon as one's personal facts are in the control of some data recipient, such a view would require that a perceived violation of trust is equivalent to a diminished state of privacy, which it is not.

77 Schrödinger, *supra* note 1.

78 WITTGENSTEIN, *supra* note 14, at §217.

79 Thanks to Joelle Pineau for constantly reminding me of this point.

80 Machine learning, however, is designed to go beyond mere execution of preprogrammed lines of code. Instead, the process relies on pattern recognition algorithms and large datasets, enabling the machine to learn by example. This can result in far more complicated algorithms to respond to different situations than human programmers would assign, amplifying the existing problems of opacity to inscrutability: where it becomes impossible for a human to understand the methodology the machine has developed. For more on the problem of

framework set out in Part III to determine whether AIs and robots can meet the privacy-diminishing threshold of level 3 (and above) epistemic quality.

Following Matheson, anyone claiming a right to privacy has a justified expectation of ignorance with respect to any personal facts that fall within the scope of the right.⁸¹ One might therefore say that privacy rights holders are entitled to a *presumption of ignorance* with respect to those personal facts.⁸² On this view, the question is whether one can rebut this presumption in the case of an artificial cognizer.

A. Can an Artificial Cognizer Form Reliable Beliefs about People?

As Fallis and others have suggested, Goldman's causal theory of knowledge⁸³ "may be just what is needed in a theory of privacy."⁸⁴ Briefly alluded to above via the concept of *cognizance*, Goldman seeks truth-indicating properties in a belief by ensuring a proper connection between the cognizer's belief and the personal fact that makes the belief true. According to Goldman's more recent formulations, this requires the cognizer to use a reliable belief-forming process.⁸⁵ One important question therefore is whether an artificial

inscrutability in machine learning, see Andrew D. Selbst & Solon Barocas, *The Intuitive Appeal of Explainable Machines*, 87 *FORDHAM L. REV.* (forthcoming 2018).

81 Matheson not only asserts the duty of ignorance but extends it to a broader range of morally obligated epistemic activities including what he calls "an obligation to forget." He ascribes this duty to humans but also to artificial entities such as online search engines and corporate agents. Matheson, *supra* note 44.

82 I intentionally shift the discourse here towards a generalized presumption of ignorance to avoid the seductive (but almost always unhelpful) temptation to frame the issue in terms of a robot's rights and duties. While I am extremely sympathetic to Peter Kahn *et al.*'s "new ontological category hypothesis," I do not think it is productive to talk in general about robots as having duties or in particular about a robot's duty of ignorance. The more elegant and important point is that robots and AIs exist in states of non-ignorance and their epistemic states have serious implications for privacy.

83 Alvin Goldman, *A Causal Theory of Knowing*, 64 *J. PHIL.* 357 (1967).

84 Fallis, *supra* note 65, at 160.

85 Goldman's theory is paradigmatic of a particular bent of epistemology: reliabilism. According to Comesaña, "[t]he reliabilist theory of epistemic justification has become known as one of the most influential theories in recent epistemology." Comesaña characterizes reliabilism as "an *explanation* thesis ... reliability is what justification *consists in*, as opposed to a condition that is independent of justification." Juan Comesaña, *Reliabilism*, in *THE ROUTLEDGE COMPANION TO EPISTEMOLOGY* 176-77 (Sven Bernecker & Duncan Pritchard eds., 2010).

cognizer can form beliefs and, if so, whether it can do so through reliable belief-forming processes.

Belief-forming processes are simply cognitive processes that end in the formation of a belief. Beliefs themselves are merely affirmations of the truth of propositions or dispositions to affirm the truth of propositions. For example, a belief that Kashmir is pregnant is simply one's affirmation of the truth of that proposition. But it is how the belief was formed that is crucial here. The belief may be formed through a *perceptual* process: looking at Kashmir (who is, let's say, in her fourteenth week of pregnancy), one sees her very noticeable baby bump and one ascertains that she is pregnant. Or the belief may be formed through an *inferential* process: without even seeing her, one notices that Kashmir's shopping cart contains cocoa-butter lotion, a large duffle bag, zinc and magnesium supplements and a bright blue rug, and comes to believe (correctly it turns out) that she is pregnant. Reliable belief-forming can, of course, also be a combination of perceptual and inferential processes.

In either of the above examples, the cognizer could be human. But the cognizer could just as easily be an AI. In the first example, an AI using machine learning to recognize objects could easily identify a baby bump as an expectant belly.⁸⁶ The second example is in fact drawn from Target's actual use of an algorithm to market consumer goods to women the algorithm believed to be pregnant.⁸⁷ Both examples demonstrate reliable, truth-promoting belief-forming

86 Machine learning can analyze enormous datasets for all sorts of nonobvious connections. While identifying pregnant bodies has not yet been achieved, it is easily imaginable, in light of other emerging algorithms. A recent Stanford study analyzing facial images was able to predict homosexuality of test subjects with 81% accuracy in males and 74% accuracy in females — essentially, machine learning “gaydar.” See Sam Levin, *New AI Can Guess Whether You're Gay or Straight from a Photograph*, THE GUARDIAN (Sept. 8, 2017), <https://www.theguardian.com/technology/2017/sep/07/new-artificial-intelligence-can-tell-whether-youre-gay-or-straight-from-a-photograph>. Similarly, in a 2016 study, machine learning-boosted regression algorithms and logistic regression were used to identify 67 biomarkers for depression in large epidemiological datasets. See Joanna F. Dipnall et al., *Fusing Data Mining, Machine Learning, and Traditional Statistics to Detect Biomarkers Associated with Depression*, 11 PLoS ONE (2016), <http://journals.plos.org/plosone/article?id=10.1371/journal.pone.0148195>. It is of course crucial to recognize that a system's overall predictive success does not mitigate its serious potential for bias. See generally ALEX CAMPELO et al., AI NOW 2017 REPORT 14-20 (Andrew Selbst & Solon Barocas eds., 2017), https://ainowinstitute.org/AI_Now_2017_Report.pdf.

87 Kashmir Hill, *How Target Figured Out a Teen Girl was Pregnant Before Her Father did*, FORBES (Feb. 16, 2012), <https://www.forbes.com/sites/>

processes of the sort required by a reliabilist theory of knowledge. They also demonstrate that such beliefs can be formulated by human or nonhuman cognizers. It makes perfect sense in both cases to say that the artificial cognizer has formed truth-promoting beliefs that are justified. And, in both cases, it is plausible to say that Kashmir's privacy has been diminished as a result.⁸⁸

On the standard reliabilist account, it is important to recognize that what matters for knowledge is reliability in fact. Thus, for pure reliabilists, it is not necessarily important to have reasons to believe that one's belief-forming process is reliable. As Peter Lipton illustrates, "[t]he cat knows birds when it sees them, but of course can give no reason to believe that its visual system is a reliable bird-detector."⁸⁹ Others, such as Wilfrid Sellars, reject reliabilism in its pure form on the view that dumb machines — like thermometers — would be said to have observational knowledge when, in fact, they don't. Sellars believes that a cognizer must not only be a reliable reporter but must also recognize that reliability by way of a capacity to connect the report with what it is a report of, as well as a capacity to discern the veracity or correctness of the report.⁹⁰ This has become known as the Reflexivity Requirement (RR).⁹¹

The stringency of RR rightly excludes dumb machines as epistemic cognizers, but it also has the potential to exclude children and nonhuman animals. This is the case even though it is evident that some young children know exactly when it is time for a diaper change, or that most puppies know exactly where the treats are kept. Consequently, other epistemologists have softened RR. For Robert Brandom, "observational knowledge has two plies: (1) it must result from the exercise of a reliable ability to respond to the relevant facts and (2') the observer must know what would rationally support

kashmirhill/2012/02/16/how-target-figured-out-a-teen-girl-was-pregnant-before-her-father-did/#247d2b896668.

88 On the assumption that the personal fact of her being pregnant was previously unknown.

89 Peter Lipton, *Popper and Reliabilism*, in KARL POPPER: PHILOSOPHY AND PROBLEMS 36 (Anthony O'Hear ed., 1995).

90 WILFRID SELLARS, EMPIRICISM AND THE PHILOSOPHY OF MIND 66-75 (Richard Rorty & Robert Brandom eds., 1997).

91 The Reflexivity Requirement is the requirement that the person holding a belief know that her belief is reliable. Knowledge therefore acquires a reflexive component. The seminal explanation is the thermometer in the main text: it may well produce reliable information, but it is unable to know that its information is reliable. The thermometer fails the Reflexivity Requirement; therefore, it cannot possess knowledge. See Willem deVries, *Wilfrid Sellars*, STAN. ENCYCLOPEDIA PHIL. (July 11, 2016), <https://plato.stanford.edu/entries/sellars/>.

it and what it would be supported by.”⁹² Chauncey Maher has suggested that 2’ is Brandom’s way of saying that the observer must only “have the relevant concepts and be in the space of reasons.”⁹³ In other words, Brandom leaves open the possibility that a cognizer can know something without being able to articulate the reasons for it.

Still, it is true that cats, puppies, and little kids who are not yet able to supply beliefs about why they hold beliefs have limited abilities as cognizers. For this reason, there has not been much written about them in the privacy context, where diminishing another’s privacy usually requires more sophisticated epistemic activity. What about *AI*s and *robots* — do, or could, any nonhuman cognizers approximate RR or Brandom’s condition 2’?

V. THE ROBOTS OF TODAY AND TOMORROW

For a couple of decades now, we have been building robots with sophisticated sensors, powerful processors, and effective actuators. With an increasing facility in pattern recognition and powered by probabilistic reasoning and techniques in machine learning, many of these artificial cognizers “have the relevant concepts” and are “in the space of reasons” required by the more stringent reliabilist theories of knowledge. These robots use probabilistic values to make decisions and use reliable belief-forming processes that meet Brandom’s conditions for observational knowledge discussed above.

Consider Xavier, a robot built and programmed at Carnegie Mellon University, first pressed into service over 20 years ago. Xavier uses a navigation architecture that harnesses the partially observable Markov decision process (POMDP) model to develop its own beliefs about its movement, in terms of position and orientation. Xavier is able to “explicitly account for various forms of uncertainty: uncertainty in actuation, sensing and sensor data interpretation, uncertainty in the initial pose (position and orientation) of the robot, and uncertainty about the environment, such as corridor distances and blockages (including closed doors).”⁹⁴ For comparison’s sake, imagine a human forming

92 CHAUNCEY MAHER, THE PITTSBURGH SCHOOL OF PHILOSOPHY: SELLARS, McDOWELL, BRANDOM 93 (2012).

93 *Id.*

94 See Sven Koenig & Reid G. Simmons, *Xavier: A Robot Navigation Architecture Based on Partially Observable Markov Decision Process Models*, in ARTIFICIAL INTELLIGENCE AND MOBILE ROBOTS 91 (1998), <https://pdfs.semanticscholar.org/66e4/162ddee5bfd10ad6273f34f05dc03ae19101.pdf>. For similar applications, see Aastha Nigam & Laurel D. Riek, Social Context Perception for Mobile Robots, Paper presented at IEEE/RSJ International Conference on Intelligent Robots and

a belief about whether it is safe to walk to their car alone based on a number of data points: the neighborhood in which it is parked, the distance from their current location to the car door, the level of illumination of the potential route, the number of other people in the vicinity, *etc.* In order to carry out a plan of getting to the car safely, a quick cognitive assemblage of these data points allows for an estimate of the probability of safely walking alone to the car.⁹⁵ Similarly, robots like Xavier that employ a Markov decision process do not merely execute preprogrammed commands, but instead harness all of the available data points from their sensors and then process them through algorithms designed to form a belief about where they are located in space.⁹⁶ They can then actuate these estimates accordingly by carrying out a mobility plan.⁹⁷ While Xavier does not know precisely where it is in the same way that a human cognizer does (*i.e.*, it may not know or appreciate that it is in the kitchen), it does know its true pose and is able to continuously update its belief about its current pose relative to its true pose. With an ability to reckon and reconsider its beliefs, Xavier is never truly lost.

Robots like Xavier are already sophisticated enough to be unchained from the factory floor. If we continue letting robots loose into the world, we will need to think more carefully about the broader social implications of their coexistence with humans in private and public spaces. Mobilizing robots with increasing cognizance will not only amplify data collection problems but will also implicate privacy and safety, depending on how the beliefs formed by robots are actuated in the real world. The safety risk is well illustrated in the extreme worries of those who campaign against autonomous weapons. Campaigners recognize the danger inherent in letting robot soldiers actuate on beliefs about who is a suspect or combatant without meaningful human control.⁹⁸ Comparable concerns arise if robots and AIs are permitted to actuate

Systems (IROS) (Sept. 28 – Oct. 2, 2015); Darren Chan, Angelique Taylor & Laurel Reik, Faster Robot Perception Using Salient Depth Partitioning, Paper presented at the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (Sep. 24, 2017).

95 This excellent example was suggested to me by Katie Szilagyi. While it may be the case that Xavier's POMDP does not currently seek to formulate beliefs about norm-infused choices such as the "safest" route, it is not difficult to imagine that it could. In fact, one could well imagine that, very soon, artificial cognizers will outperform average human cognizers at such belief-formation processes.

96 Koenig & Simmons, *supra* note 94, at 1.

97 *Id.*

98 The debates regarding use of lethal autonomous weapons in battlespace are preoccupied with the level of human control in targeting and kill decisions. This is referred to as meaningful human control, or MHC. See Rebecca Crotoft, *A*

on beliefs that they form in the law enforcement context, such as NSA's PRISM, discussed above in Part II. Related concerns arise in the consumer context. It is uncontroversial that automated decision-making risks bias, social sorting, and discrimination.⁹⁹ Those problems change by order of magnitude when robots automatically act upon their own beliefs without human oversight. Calls for fairness, accountability, and transparency will be rendered ineffectual if the response is that those decisions and their consequences don't implicate privacy because there was no particular human observation or awareness. A single individual coexisting with an island full of law enforcement robots is not alone. Privacy's presumption of ignorance is easily rebutted. Let us hope the presumption of innocence is not.

Of course, artificial cognizers do not have to be mobile, embodied, or armed to pose privacy threats on the basis of their ability to actuate beliefs that they have formed. Current semantic AI systems combine common sense ontological methods with powerful reasoning engines and natural language interfaces to enable the development of novel knowledge-intensive applications that are sure to implicate privacy, even without human participation. Systems like Cyc use knowledge modelling language, a vast and broad-reaching library of formally represented knowledge, and a powerful, efficient, and extensible inference engine to answer questions, draw conclusions, and solve problems.¹⁰⁰ These applications cover a broad range of functionality — from answering questions for medical researchers, to identifying dangerous scenarios, to optimizing workflows, to efficient resource utilization, to intelligent decision support.¹⁰¹

Because the beliefs generated by such systems can also be programmed to actuate on their own — *i.e.*, to cause people to act or be treated in particular ways, or machines to carry out certain operations independently of human action, oversight, intervention, awareness or knowledge — these AIs can not only diminish our privacy but are in a position to violate it and cause harm. They can, quite literally, affect our life chances and opportunities based on the beliefs they have formed.

If the history of computing continues to maintain its current exponential trajectory, one can only imagine that the potential risks described above will be greatly amplified. We can envisage a world in which there is more and more interaction with artificial cognizers. Not the imaginary robots of Nick Bostrom,

Meaningful Floor for "Meaningful Human Control", 30 TEMPLE INT'L & COMP. L.J. 53, 54 (2016).

99 CAMPELO et al., *supra* note 86. See also Julie Cohen, *Turning Privacy Inside Out*, 20 THEORETICAL INQUIRIES L. 1 (2019).

100 *About Cyc*, CYCORP, <http://www.cyc.com/about/> (last visited Dec. 10, 2017).

101 *Id.*

Max Tegmark, or Elon Musk,¹⁰² but real robots. Machines with exquisite sensor capability, unfathomable facility at pattern matching, powered by AI truly capable of deep learning. Machines with high precision, fantastically sophisticated actuators. Not sentient creatures, but entities capable of acting upon the world with substantial autonomy and the ability to significantly affect the life chances and opportunities of people. As Ryan Calo has described them,¹⁰³ they will be embodied, charged with social valence, and will behave in ways that are emergent — unpredictable and often inscrutable not only to the people interacting with them, but also to those who designed and programmed them. In spite of all this, they will not be sentient or conscious, and they will not be people.

As I have argued in the preceding sections, even though tomorrow's robots will lack these qualities, and even if they are *never* capable of achieving human-level knowledge, many artificial cognizers already meet the epistemic conditions of reliable belief formation and are capable of acting on those beliefs (without human intervention or oversight) in ways that implicate privacy. Consequently, it is my contention that legal theory and doctrine will have to expand their understanding of privacy relationships to include robots and AIs that meet these epistemic conditions. Such machines are no longer ignorant. They possess epistemic qualities that force us to rethink our understanding of privacy relationships.

VI. RELATIONAL ROBOTS

Having considered robots and AIs from the perspective of epistemic privacy, we can now begin to see more clearly how and why artificial cognizers are becoming inextricably intertwined in the dense network of relationships in which many of us live. This will take on even greater significance if social and collaborative robots flourish.

Social and collaborative robots are an important frontier for human-robot interactions. Collaborative robots are used primarily in industrial settings, employed alongside human coworkers and registering their presence while collaborating and completing tasks. The rationale for introducing collaborative

102 See generally BOSTROM, *supra* note 5; MAX TEGMARK, *LIFE 3.0: BEING HUMAN IN THE AGE OF ARTIFICIAL INTELLIGENCE* (2017); Maureen Dowd, *Elon Musk's Billion Dollar Crusade to Stop the A.I. Apocalypse*, VANITY FAIR (Apr. 2017), <https://www.vanityfair.com/news/2017/03/elon-musk-billion-dollar-crusade-to-stop-ai-space-x>.

103 Ryan Calo, *Robotics and the Lessons of Cyberlaw*, 103 CALIF. L. REV. 513 (2015).

robots to the workplace is that, in some situations, humans and bots, working together, will be able to accomplish more than either would working alone. Social robots (which can also be collaborative) are specifically designed and programmed to project social behavior. This is in part accomplished through “anthropomorphic design” — the art and science of building robots that project human features and reactions. Imbuing human qualities in machines adds the complexity of empathy and social interaction to preexisting issues of privacy, as robotic sensors collect, store and process data in ways that are rendered invisible. Although designing bots to appear human can augment trust and increase uptake in their use, these added complexities challenge current models of privacy protection and a range of related privacy values.¹⁰⁴ For example, anthropomorphic design may be used to impact a robot’s perceived epistemic status. We can imagine dumbing-down a robot in order to obscure its actual epistemic capacities, so that those interacting with it are unaware of the actual privacy risks (*e.g.*, some next generation Hello Barbie). One can likewise imagine using anthropomorphic design to amplify a robot’s perceived epistemic capabilities to make people feel as though they are under surveillance (*e.g.*, Elf on a Shelf).¹⁰⁵ Of course, the very same techniques can be used to invoke a sense of trust,¹⁰⁶ or a desire to help the robot.

Consider hitchBOT, the Canadian hitchhiking robot designed to travel the world dependent on the kindness of strangers. Fueled by an entirely manipulated altruism, hitchBOT’s extensive travels were intriguing in many

104 Ian R. Kerr, *Bots, Babes, and the Californication of Commerce*, 1 U. OTTAWA L. & TECH. J. 285 (2004).

105 I owe this important point to Alon Jasper and a related one to Eldar Haber. Jasper and I agree that we recreate the gaze of surveillance in private spheres when we invite such robots into our homes and (un)willingly anthropomorphize them. With Haber and Jasper, I also agree that there are potential privacy implications with or without disclosure by the robot to third parties (including disclosure to other robots). These scenarios illustrate and partially explain our perceived privacy violations. However, in my view, current robots such as Hello Barbie, Elf on the Shelf, and Alexa do not diminish our privacy without disclosure to more sophisticated cognizers since these machines are incapable of forming beliefs or higher-level epistemic states.

106 Consequently, as Mireille Hildebrandt and Eldar Haber have suggested, epistemological considerations of privacy merit discussions about trust and trustworthiness. A trustworthy AI could conceivably collect, use, or disclose our personal information in ways that ensure no violations of our privacy expectations. This invites the possibility that AIs need not (or would not) infringe privacy or generate privacy harms. While these are important considerations, as mentioned previously, this article does not delve into such normative questions.

respects. Its social programming, combined with its sleek technological design, mirrored human behavior in important ways. At the same time, hitchBOT also possessed advanced data-gathering capabilities. When hitchBOT was brutally vandalized in its travels across the continental United States — all of which was captured by the robot's sensors — its demise raised questions about crimes against robots, privacy and surveillance, human empathy, and whether robots deserve the same empathy given their designed proximity to humankind. The continued enhancement and complexity of human-robot relations will surely contribute further to the density and intensity of our network of relationships, which will doubtlessly contain people, smart technologies, and even smarter robots.

Such developments will demand a careful elaboration of relational privacy as theorized by Nedelsky, Cohen, and Hargreaves. As I have suggested, we will need to (re)configure not only the networked self, but also the networked other. More likely, we will need an entire relational account of robots¹⁰⁷ — of the sort that other scholars have diligently carried out with animals.¹⁰⁸ Although a comprehensive relational account of robots is not my present aim, I do hope this article provides at least the first footnote in the recognition that a theory of relational privacy must count robots and AIs as integral to the configuration of the networked other.

VII. SCHRÖDINGER'S ROBOT

The steel chamber sits off in the distance, the robot still penned inside.

Despite my desired fidelity to Schrödinger's fantastic thought experiment, I am unwilling to go to quite the same lengths to stack the deck. As he brilliantly concocts the example, Schrödinger's cat, a strange biological sensor, is both dead and alive until the moment of the grand reveal. But once the box is cut away, a very awkward binary animates its biology (or not, as the case may

107 Thank you to Steven Bero for reinforcing this point.

108 See generally Maneesha Deckha, *Critical Animal Studies and Animal Law*, 18 ANIMAL L. 207 (2012); SUE DONALDSON & WILL KYMLICKA, *ZOOPOLIS: A POLITICAL THEORY OF ANIMAL RIGHTS* (2011); DONNA J. HARAWAY, *WHEN SPECIES MEET* (2007). Although my epistemic analysis crudely lumps robots and animals together as nonhuman cognizers, there is a noteworthy difference between them as it concerns relational theory. The purpose of investigating the relationality of robots is not to adduce how to treat nonhuman cognizers (as is the case when relational theory is applied to animal ethics). The goal of developing a relational theory of robots is in fact to better understand human ethics, including the moral right to privacy.

be): the cat can be only one of dead or alive. Conversely, as I have tried to demonstrate throughout this Article, the epistemic qualities of an AI-driven robot are non-binary. As much as Judge Posner and others would prefer a simple privacy solution triggered by the threshold of human sentience, privacy, much like the epistemic relations between facts and cognizers (whether human or artificial), is a degree concept.

What then of the theoretical basis for legal doctrine which assumes human sentience as its threshold?

Our accepted theories of knowledge make it quite clear that much weaker epistemic states can indeed impact privacy. If privacy requires some degree of ignorance, it is not clear that the conditions for privacy are possible in an unregulated world populated by sophisticated artificial cognizers. This is true even if the robot never escapes Schrödinger's steel chamber to disclose its epistemic findings to humans. So long as an artificial cognizer can actuate from inside the box in ways that affect the outside world, privacy and other risks are imminent.

In this Article, I have examined epistemic privacy's notion of a duty of ignorance through a relational lens, plucking from its wisdom a presumption of ignorance as a heuristic for our increasingly complex web of relationships. I have argued that robots and AIs disturb that presumption in epistemologically significant ways. Legal presumptions — for example, the presumption of innocence — form the cornerstone of our legal system. Although we often think about the presumption of innocence as a narrow set of distinct procedural safeguards enshrined in law, social theorists including Ericson and Gandy have rightly taught us that we can also understand this presumption as a broader moral claim.¹⁰⁹ Its broad aim is to provide fair and equal treatment to all by setting boundaries around the kinds of assumptions and beliefs that can and cannot be made about people.

I have suggested that a presumption of ignorance is a useful heuristic for thinking about robots and AIs. Rather than ask whether the robot is sentient, conscious, or has rights, I contend that it is far more useful to recognize and refine the precise epistemic conditions that implicate privacy. Having done so, it seems clear that privacy theory and doctrine must be expanded to incorporate a dense network of relationships that includes robots and AIs.

109 Ian Kerr, *Prediction, Preemption, Presumption: The Path of Law After the Computational Turn*, in *PRIVACY, DUE PROCESS AND THE COMPUTATIONAL TURN: THE PHILOSOPHY OF LAW* 91 (Mireille Hildebrandt & Katja de Vries eds., 2013); Ian Kerr & Jessica Earle, *Prediction, Preemption, Presumption: How Big Data Threatens Big Picture Privacy*, 66 *STAN. L. REV. ONLINE* 65 (2013).

