Draft Feb. 25, 2019. Forthcoming in *Uniform Law Review* Special Issue on Transnational Commercial Law and Technology/Digital Economy. Comments welcome.

ADVANCED ARTIFICIAL INTELLIGENCE AND CONTRACT

John Linarelli*

Conjectures are of great importance since they suggest useful lines of research. – Alan Turing[1]

Abstract

The aim of this article is to inquire whether contract law can operate in a state of affairs in which artificial general intelligence (AGI) exists and has the cognitive abilities to interact with humans to exchange promises or otherwise engage in the sorts of exchanges typically governed by contract law. AGI is a long way off but its emergence may be sudden and come in the lifetimes of some people alive today. How might contract law adapt to a situation in which at least one of the contract parties could, from the standpoint of capacity to engage in promising and exchange, be an AGI? This is not a situation in which AI operates as an agent of a human or a firm, a frequent occurrence right now. Rather, the question is whether an AGI could constitute a principal – a contract party on its own. Contract law is a good place to start a discussion about adapting the law for an AGI future because it already incorporates a version of what is known as weak AI in its objective standard for contract formation and interpretation. Contract law in some limited sense takes on issues of relevance from philosophy of mind. AGI holds the potential to transform a solution to an epistemological problem of how to prove a contract exists into solution to an ontological problem about the capacity to contract. An objection might be that contract law presupposes the existence of a person the law recognizes as possessing the capacity to contract. Contract law itself may not be able to answer the prior question of legally recognized personhood. The answer will be to focus on how AGI cognitive architecture could be designed for compatibility for human interaction. This article focuses on that question as well.

INTRODUCTION

Artificial intelligence (AI) will bring about the next big change in the law. It will compel us to reconsider who or what will be the subjects upon which the law imposes duties, confers powers, and allocates rights. AI already challenges the allocation of legal duties, powers, and rights to persons already recognized as subjects of the law. But the challenge to-date has been incremental. The next big change will be transformational. It will be more significant than the legal recognition of the limited liability firm, the last

* Professor of Commercial Law, Durham University Law School.

[1] Alan M. Turing, *Computing Machinery and Intelligence*, 59 *Mind* 433, 442 (1950).

great move in legal systems to alter who (or more precisely what) gets legal rights and bears liabilities.[2] An artificial person with limited liability such as a corporation still must operate through human agents,[3] though American law may permit an automated system to operate but not create a member-less limited liability company.[4] A sufficiently advanced AI, a so-called artificial general intelligence (AGI) will not need to operate through agents. In the study of the history of science we are admonished to avoid extending Thomas Kuhn's notion of a paradigm shift beyond its intended purposes.[5] It may be no exaggeration to apply the notion of the paradigm shift here to argue that the next big change in the law to accommodate AGI will not be normal legal change.

We do not know where advances in AI are headed. AGI does not yet exist. AI currently falls short of human capabilities generally though AI already outperforms human intelligence in many specific domains.[6] The Stanford study, *Artificial Intelligence and Life in 2030,* states that to date, "no machines with self-sustaining long-term goals and intent have been developed, nor are they likely to be developed in the near future."[7] But AGI will come to exist in the future. Claims that AGI could never exist are implausible. To base predictions on the current state of AI technology would not be rational. To hold

---

[2] *See* JOHN MICKLETHWAIT & ADRIAN WOOLRIDGE, THE COMPANY: A SHORT HISTORY OF A REVOLUTIONARY IDEA (2005); Henry Hansmann, Reiner Kraakman, & Richard Squire, *Law and the Rise of the Firm* 119 HARV. L. REV. 1335 (2005).

[3] Horst Eidenmüller, *The Rise of Robots and the Law of Humans* (Oxford Legal Studies Research Paper No. 27/2017, 2017), https://ssrn.com/abstract=2941001; Scott Bayern, Thomas Burri, Thomas D. Grant, Daniel M. Häusermann, Florian Möslein & Richard Williams, *Company Law and Autonomous Systems: A Blueprint for Lawyers, Entrepreneurs, and Regulators*, 9 HASTINGS SCI. & TECH. L. J. 135 (2017).

[4] Shawn Bayern, *The Implications of Modern Business-Entity Law for the Regulation of Autonomous Systems,* 7 EUR. J. RISK REG. 297 (2016); Shawn Bayern, *The Implications of Modern Business Entity Law for the Regulation of Autonomous Systems,* 19 STAN. TECH. L. REV. 93, 104 (2015); Scott Bayern, *Of Bitcoins, Independently Wealthy Software, and the Zero Member LLC,* 108 NW. U. L. REV. 1485, 1496-97 (2014). For concern about threats to humanity from entities controlled by AI, *see* Lynn M. Lopucki, *Algorithmic Entities,* 95 WASH. U. L. REV. 1 (2018).

[5] THOMAS KUHN, THE STRUCTURE OF SCIENTIFIC REVOLUTIONS (1962).

[6] NICK BOSTROM, SUPERINTELLIGENCE: PATHS, DANGERS, STRATEGIES 14 (2014). But certainly not in all areas. *See* Harry Surden, *Machine Learning and Law* 89 WASH. L. REV. 87 (2014).

[7] PETER STONE ET AL., STANFORD UNIV., ARTIFICIAL INTELLIGENCE AND LIFE IN 2030: REPORT OF THE 2015 STUDY PANE 4 (2016).

predictions hostage to some presupposed uniqueness of human understanding or consciousness is unsupportable.[8] Nick Bostrom tentatively predicts the existence of human level AI by mid-century and soon thereafter a super-AI vastly exceeding the cognitive abilities of humans.[9]

And we do not know of the progress *AI itself* will make in advancing towards AGI. It is possible that the basic components of AI cognitive architecture (code or otherwise) will become something roughly analogous to "gene" and AI evolution will occur independently of the initial human intervention of AI creation, if humans engineer such evolution to be possible. Machine learning already accomplishes this in limited domains.[10]

The aim of this article is to assess the feasibility of investing an AGI, from a legal point of view, with the power to enter into contracts, either with humans or with other AGIs. The argument made in this article is that an AGI can be a party to a contract. It can be legally obligated for promises it makes in contracts, have and enforce rights as a matter of contract, and bear contractual liabilities. Part of the answer rests on contract law itself and part rests on the need to design AGI cognition to be compatible with the

[8] A longstanding disagreement in philosophy of mind is between those who argue that AI, or at least what is known as strong AI, is impossible. Strong AI is AI that *actually* thinks, is conscious, has a phenomenology of the particular experiences of life, and has the properties of intentionality that humans have. Weak AI is AI that acts *as if* it is thinking, conscious, and acting with intentionality. *See* STUART RUSSELL & PETER NORVIG, ARTIFICIAL INTELLIGENCE: A MODERN APPROACH 1020-1033 (3d ed. 2010)(overview). While staking out a middle ground David Chalmers offers a good summary of the arguments. DAVID J. CHALMERS, THE CONSCIOUS MIND: IN SEARCH 6OF A FUNDAMENTAL THEORY 313-332 (1996). This is an argument about whether AI must have some form of inner life, some phenomenology of conscious experience, or true understanding or whether a simulation of these things will suffice, is beyond our scope here. This article rests on the argument that weak AI is sufficient for contracting and so no need exists to engage in this debate. AI researchers take a similar view and go further, arguing that philosophers are asking the wrong questions that are too open ended for science. The argue, for example, that asking whether AI has consciousness is akin to asking if aircraft can fly because they are not birds or asking whether submarines can swim because they are not fish. RUSSELL & NORVIG, *supra*, at 1021.

[9] BOSTROM, *supra* note xx, at 24-25.

[10] *See supra* note xx and accompanying text.

requirements of contract law and the properties of exchange in social and market contexts.

At the outset, distinguish contracts with and by AGIs from smart contracts. This article does not deal with smart contracting. The term "smart contract" is ambiguous.[11] It is not a legal concept. Very simply, a smart contract is a contract (a legally enforceable agreement) for which some or all contract performance is executed and enforced digitally and without the need for human intervention except at the level of writing code to automate contract performance.[12] Distributed ledger technology has advanced substantially the ability of contract parties to write and use smart contracts. The combination of the distributed ledger, the network, and the consensus mechanisms built into distributed ledger technology facilitate trust between contract parties and replace humans in institutions operating as intermediaries.[13] In short, smart contracts substitute algorithmic for human contract performance and enforcement. The next-generation conceptualization of smart contracts might be as "algorithmic" contracts in which algorithms operate constructively as agents for humans.[14] A more general notion is Scott Bayern's process-agreement equivalence principle: "at least as a matter of conceptual logic, a legally enforceable agreement may give legal significance to arbitrary features

[11] Harry Surden has developed a typology to clarify the evolution and categories of smart contracting. He classifies the evolution of digitized agreements, starting with online contracts, moving to data-oriented contracts that specify obligations in code, and then on to computable contracts that assess contract performance and produce consequences. Harry Surden, *Computable Contracts,* 46 U.C. DAVIS L. REV. 629, 631-42 (2012). He developed this typology before the rise of blockchain. Kevin Werbach & Nicholas Cornell, *Contracts Ex Machina*, 67 DUKE L. J. 101, 108-112 (2017).

[12] Various authors have offered definitions of a smart contract. Nick Szabo is credited with inventing the phrase. Werbach & Cornell, *supra* note 9 at 107. Szabo defines a smart contract as a "set of promises, specified in digital form, including protocols within which the parties perform on these promises." Nick Szabo, Smart Contracts: Building Blocks for Digital Markets (1996), available at http://www.fon.hum.uva.nl/rob/Courses/InformationInSpeech/CDROM/Literature/LOTwinterschool2006/szabo.best.vwh.net/smart_contracts_2.html (accessed Feb. 7, 2019). Max Raskin describes smart contracts as "agreements wherein execution is automated, usually by computers." Max Raskin, *The Law and Legality of Smart Contracts,* 1 GEO. L. TECH. REV. 305, 306 (2017); Werbach & Cornell define a smart contract as an "agreement in digital form that is self-executing and self-enforcing." Werbach & Cornell, *supra* note xx at 108.

[13] Werbach & Cornell, *supra* note xx at 118.

[14] Lauren Henry Scholz, *Algorithmic Contracts,* 20 STAN. TECH. L. REV. 128 (2017).

of the state of any process (such as an algorithm or physical system) by specifying legal conditions satisfied by features of that state."[15]

Smart contracts are in widespread use now and their use is increasing. An example of a smart contract is Fizzy, the automated flight delay insurance system used by the French airline AXA, running on the Ethereum blockchain. Fizzy allows passengers to be indemnified for late fight arrivals as soon as they arrive at their destination. Passengers need do nothing other than buy the insurance on the AXA app. With Fizzy, if a flight is more than two hours late, the passenger will receive an automatic notification with compensation options. The code sends the compensation directly to the bank or credit card the customer has chosen.[16] Compare the Fizzy smart contract with the "dumb" contracts that train operators use in the United Kingdom.[17] In the United Kingdom, the typical terms of the franchise by government to train operators requires the train operator to adhere to a passengers' charter mandating passenger compensation for some train delays, depending on cause and duration. To claim compensation, passengers must enter details online along with a readable scan of their paper tickets or mail the tickets along with a form completed by hand. Train operators promise compensation in 28 days. Train passengers must endure this process even though the train operators have in their computer systems all the information they need about passengers and delays.

The difference between contracting with or by an AGI and smart (or algorithmic) contracts is that humans or legal persons in the form of entities such as corporations are the actual parties to smart contracts, whereas contracts with AGIs involve at least one

[15] Bayern, *supra* note xx, at 300; Bayern, Burri *et al.*, *supra* note 2, at 136. On the limits of smart contracts, *see* Werbach & Cornell, *supra* note xx; Jeremy M. Sklaroff, *Smart Contracts and the Cost of Inflexibility*, 166 U. PA. L. REV. 263 (2017).

[16] AXA Goes Blockchain with Fizzy, https://www.axa.com/en/newsroom/news/axa-goes-blockchain-with-fizzy (accessed Feb. 8, 2019).

[17] *See, e.g.*, Great Western Railway, Passenger's Charter, https://www.gwr.com/about-us/our-business/passengers-charter (accessed Feb. 8, 2019). *See also* Jeffrey Lipshaw. *The Persistence of "Dumb" Contracts*, 2 STAN. J. BLOCKCHAIN L. & POL'Y 1 (2019).

contract party that is neither a human nor a currently recognized legal person such as a corporation, limited liability company, or partnership. This difference is fundamental. This article examines what it might mean for contract law when AI is sufficiently advanced that it could have the capacity, in a cognitive sense, to make and perform contracts.

This article is organized as follows. Part I examines how contract law, already at least partly and with some adaptation, has answered the question of how to treat an AGI as a contract party. The objective theory of contract, prevalent in American and English common law, informs us that in determining whether a contract came into existence and what its terms might be, courts do not inquire whether persons in their minds actually possessed intent to be bound in contract, but look to external evidence in the form of an outward manifestation of assent by words, behavior, and action. The objective theory coincides closely to the Turing test for assessing whether an AI exists. The Turing test requires an evaluator to evaluate text-based communications between a human and a machine and if the evaluator cannot reliably predict who is human and what is machine, then the machine has passed the test. The subjective-objective debate in the common law of contract, long settled in favor of the objective standard, reflects with remarkable consistency debates among philosophers and AI researchers on how to assess intentions of AI and test whether AI could plausibly have consciousness. That the objective theory of contract may be meant to deal with an epistemological problem, that of determining whether a contract has come into existence, and not an ontological problem, that of determining whether the actors who are attempting to contract are eligible as persons to be bound by the law, does not pose an obstacle to using the theory to support recognition of an AGI as a contract party. Part II addresses how to design an AGI to allow it to interact with humans in the domain of contract. This is a question of

how to code the cognitive architecture of AGI. Part II does not argue that AI researchers must replicate or copy humans for the sake of doing so, if that will ever be possible. Rather, the focus should be on interactivity and cognition associated with transacting and exchange. Because contract law evolved as a social institution to enforce group norms for humans and reflects normative concepts found in human cognition, any AGI, to form contracts with humans, will have to possess some cognitive adaptations that support exchange-level interactions of a contractual kind with humans. Developing AGIs as contractors is a demanding task, and while the investigation of these questions is speculative at this point, it might be an eminently practical sooner than we think.

I. THE TURING TEST ALREADY IN CONTRACT LAW

In his 1950 article, *Computing Machinery and Intelligence*, Turing starts by stating his question as "can machines think?"[18] He eventually finds this question "too meaningless to deserve discussion"[19] and replaces it with "are there any imaginable digital computers which would do well in the imitation game?"[20] The imitation game proceeds as follows. The players in the game are a human being and a machine. The game also includes a human interrogator. The interrogator is in a separate room from the human and the machine. The interrogator knows the others are labelled *X* and *Y*, knows one is human and the other machine, but does not know which. The object of the game is to test whether the interrogator can tell the difference between the human and the machine though a series of questions. The interrogator is to ask questions to the machine and the individual through a text channel to avoid revealing which is the human and which the machine. If the interrogator cannot reliably distinguish the human from the machine,

[18] Turing, *supra* note 1.
[19] *Id.*, 442.
[20] *Id.*

then the machine passes the test. The test is meant to assess whether a machine can give answers that consistently resemble those a human would give such that humans cannot tell the machine apart from a human, in terms of the cognitive performance of the machine.[21]

There has been much debate about what the Turing test tells us. Turing offered his test at a time when behaviorism was ascendant in psychology, and cognitive science, a field closely aligned with artificial intelligence and focusing on mental representations and not only on behavior, was only beginning to emerge as a field of study.[22] Gilbert Ryle's, *The Concept of Mind,* a thorough rejection of Cartesian dualism, was published a year before Turing's paper.[23] While Ryle said that his work was limited to rejecting mind-body dualism as it is understood in philosophy, he did point out that a methodological problem with psychology before behaviorism was that "the reputed deliverances of consciousness and introspection are not publicly checkable."[24]

AI definitions vary based on whether to focus on behavior and action or on whether machines actually think.[25] Turing focused on the former. His test is now understood as a way to look for "weak AI," enough for classifying a machine as an AI if the machine can act *as if* it is thinking and has intentionality.[26] Distinguish strong AI: for a machine to qualify as intelligent in a strong AI sense, it actually has to think and have actual intentions associated with its actions and not just simulate thinking and intentions.[27] In

[21] *See* Stevan Harnad, *The Turing Test Is Not A Trick: Turing Indistinguishability Is A Scientific Criterion*, 3 SIGART BULLETIN 9 (1992).
[22] The first academic gathering that got cognitive science as a discipline started was the Hixon symposium, "Cerebral Mechanisms in Behavior," held at the California Institute of Technology in 1948. HOWARD GARDNER, THE MIND'S NEW SCIENCE: A HISTORY OF THE COGNITIVE REVOLUTION 10 (1985).
[23] GILBERT RYLE, THE CONCEPT OF MIND (1949).
[24] *Id.*, 327.
[25] Russell & Norvig, supra note xx, at 1, 1020. *See also* note __ and accompanying text.
[26] Russell & Norvig, *supra* note xx, at 1, 1020; WENDALL WALLACH & COLIN ALLEN, MORAL MACHINES: TEACHING ROBOTS RIGHT FROM WRONG 74-75 (2010).
[27] Russell & Norvig, *supra* note xx, at 1026-27; Wallach & Allen, *supra* note xx, at 74-75; John Searle, *Minds, Brains, and Programs,* 3 THE BEHAVIORAL & BRAIN SCIENCES 417 (1980).

other words, strong AI posits the notion that a machine has a "mind" that can understand and have mental states just like humans. AI researchers do not care about these distinctions and accept weak AI as sufficient to determine whether machine intelligence can be classified as AI.[28]

Weak AI is sufficient for purposes of determining whether an AGI could be a party to a contract, in terms of understanding the question as one internal to contract law. The Turing test has been effectively embedded into Anglo-American contract law, in the objective theory of contract.[29] The focus of this tradition in contract law is on outward appearances – on what can be proven as a matter of evidence independent of the mental states the parties may or may not have. The objective theory of contract tells us that intention to be bound to or form a contract is determined by evidence external to the actual intentions of the parties. Judge Learned Hand has said:

> A contract has, strictly speaking, nothing to do with the personal, or individual, intent of the parties. A contract is an obligation attached by the mere force of law to certain acts of the parties, usually words, which ordinarily accompany and represent a known intent. If, however, it were proved by twenty bishops that either party, when he used the words, intended something else than the usual meaning which the law imposes upon them, he would still be held, unless there were some mutual mistake, or something else of the sort.[30]

Judge Frank Easterbrook has explained that intention to be bound "does not invite a tour through [a contract party's] cranium" but must necessarily be derived from a consideration of the words, written and oral, and actions of the part parties.[31] Often quoted

[28] Russell & Norvig, *supra* note xx, at 1020. Searle, widely known as an avid critic of the notion of the possibility of AI, seems only to object to the possibility of strong AI. Searle, *supra* note xx, at 417.
[29] *See* Joseph M. Perillo, *The Origins of the Objective Theory of Contact Formation and Interpretation*, 69 FORDHAM L. REV. 427 (2000); Timothy A. O. Endicott, *Objectivity, Subjectivity, and Incomplete Agreements*, *in* OXFORD ESSAYS IN JURISPRUDENCE FOURTH SERIES 151 (Jeremey Horder ed., 2000).
[30] Hotchkiss v. National City Bank, 200 F. 287, 293 (S.D.N.Y. 1911), *aff'd,* 201 F. 664 (2d Cir. 1912), *aff'd,* 231 U.S. 50 (1913).
[31] Skycom Corp. v. Telstar Corp., 813 F.2d 810, 814 (7th Cir. 1987).

on the objective theory of contract formation and interpretation is the New Hampshire Supreme Court, itself quoting Oliver Wendell Holmes Jr.:

> A contract involves what is called a meeting of the minds of the parties. But this does not mean that they must have arrived at a common mental state touching the matter at hand. The standard by which their conduct is judged and their rights are limited are not internal but external. In the absence of fraud or incapacity, the question is: What did the party say and do? "The making of a contract does not depend upon the state of the parties' minds; it depends upon their overt acts."[32]

The U.S. Restatement (Second) of Contracts, intended to reflect a consensus about contract law in the United States, does not contain any section explicitly titled on intention to form a contract. It advises us that American contract law has likely abolished the idea of intention to be legally bound. Restatement (Second) section 21 provides that "neither real nor apparent intention that a promise be legally binding is essential to the formation of a contract. . . ."[33] American contract law instead relies on what is known in American law as manifestation of mutual assent, which requires each party either to promise, objectively understood, or perform.[34] While English law does not reflect this Restatement (Second) language of manifestation of mutual assent, it is substantially similar in adhering to an objective theory of contract formation and interpretation.[35] In English law, intention to create legal relations is traditionally only used to distinguish promises the parties want the law to enforce and promises they do not want the law to enforce.

---

[32] Woburn National Bank v Woods, 77 N.H. 172, 89 A 491, 492 (1914)(citation omitted), *quoting* OLIVER WENDELL HOLMES, JR., THE COMMON LAW 307 (1881).
[33] Restatement (Second) of Contracts §21 (1981).
[34] *Id.,* §18. The "manifestation" language is pervasive in American contract law and reflects the notion of contracting making sense only in the form or external representations to other persons. *See, e.g.,* Restatement (Second) of Contracts §2 (1981), which defines core concepts such as a promise as "a manifestation of intention to act or refrain from acting in a specified way, so made as to justify a promisee in understanding that a commitment has been made." Comment b explains that a manifestation of intention is an "external expression" as opposed to "undisclosed intention." Restatement (Second) of Contracts §2 cmt. b (1981).
[35] Endicott, *supra* note xx; EDWIN PEEL, TREITEL ON THE LAW OF CONTRACT 1 (14th ed., 2015)(§1-002, 'The objective principle")..

No "Chinese room" problem confronts the thin version of intentionality in American and English contract law. The philosopher John Searle posed the Chinese room problem, a thought experiment[36] to argue that strong AI is impossible. In summary form the Chinese room problem proceeds as follows. Imagine yourself alone in a room following a set of rules in English, a language you understand, responding to Mandarin Chinese characters slipped under the door. You understand no Mandarin, but you are following the rules provided to you, analogous to a computer program, to produce appropriate responses in the form of Mandarin Chinese characters to questions posed in Mandarin Chinese to you under the door. You can do this to an adequate level of proficiency that you fool the person on the other side of the door into believing that you actually know Mandarin Chinese. The thought experiment is designed to show that programming a machine might make the machine appear to understand a language, but it does not actually understand the language. Searle argues that the thought experiment illustrates that computers only use syntactic rules to manipulate symbol strings but have no understanding of the meaning or semantics of those symbols.

Several conclusions can be drawn from the Chinese room problem that go far beyond our scope here.[37] For our purposes the Chinese room problem is useful in helping us understand why weak AI just might be enough for contract formation and interpretation. Think about how the common law of contract does not recognize a "secret" intent to form a contract or not to form a contract, or for a particular provision to be or not to be in a contract. Contract law does not make mental states or "mind" relevant. Simulating "real" intent does not matter to contract law. Analogizing to the Chinese

[36] On the use of thought experiments in philosophy *see* Kimberly Brownlee & Zofia Stemplowska, *Thought Experiments, in* METHODS IN ANALYTICAL POLITICAL THEORY 21 (Adran Blau ed., 2017)
[37] For an extended discussion, *see* David Cole, *The Chinese Room Problem, in* STANFORD ENCYCLOPEDIA OF PHILOSOPHY (Edward N. Zalta, ed., 2014), https://plato.stanford.edu/entries/chinese-room/#toc (last visited Feb. 16, 2019).

room thought experiment, the symbols that come out of the room matter for contracting, regardless of how they are produced. The semantics of contract terms does not rely on any "true" understanding, in terms of what the parties actually mean, but only what their actions or outward appearances convey what they mean. This is weak AI. The law of contracts only recognizes weak forms of intelligence, natural or artificial. Simulated or real, whatever that may mean in the study of consciousness and the mind, are simply irrelevant to contract law.

The argument for recognition of AGI as a potential contracting party using the objective theory of contract is not an extension of objective theory beyond its purposes. A possible objection to the argument, which would ultimately prove unsuccessful, might proceed along the following lines. The objective theory of contract is meant to be epistemological and not ontological. Objectivity in contract law is meant to solve problems of proof and evidence. It is about providing tractable means by which to prove to a fact-finder (a judge or a jury) whether the parties formed a contract and on what terms.[38] My argument, so the objection goes, attempts to answer a prior question, about the nature of the contract parties themselves, about who can be said to qualify by law to form an intent in the first place. These questions will be taken up in part II below in a more policy-oriented sense, but this part will address it from the internal perspective of the objective theory of contract itself.

The objective theory of contract is more than about what one can prove in court. It is not a set of propositions from the law of evidence but a set of propositions about the concepts or properties of contract. Objectivity goes to a claim that differs from how to prove. It goes to whether a contract exists, or not. Nothing about contract law is first person or phenomenological. Its doctrines are functional. Contract law does not care

[38] *See* Perillo, *supra* note xx.

about the nature of the cognitive systems of a contracting party but how they function. If an intelligence is capable of manifesting to the external world that it can promise, bargain, and discharge contractual rights and obligations, then its outward appearances in the form of objective evidence is not subject to refutation, as a matter of contract law.[39]

We see the focus on function in the law on capacity to contract. Relying on existing law risks anthropocentrism and setting an inappropriate standard, but given law's emphasis on text, tradition, and precedent, let us give it a go.

Preliminarily, avoid tautologies. It is true that capacity to contract can exist because the law makes it so. To rely on such a claim would be an exercise in empty formalism. It is equivalent to saying *p* because *p.* A stipulative approach works for delimiting capacity as it relates to the age of individuals and to artificial persons such as corporations. These persons derive their capacity to contract through statute or well-settled common law principles.[40] Circular reasoning will not solve the problem. It may come to pass that AGI does receive statutory or even case law recognition as a person, but such recognition usually hinges on some deeper need or justification and that is what we explore here.

The law on contractual capacity of natural persons does not depend on the "nature" or qualities of the person who is the contracting party, and so it is at best unclear why it should do so for AGI. There is scant recent common law on capacity and the subject

[39] *See supra* note xx. Very crudely, contract law could be said to support or be consistent with a functionalist philosophy of mind. There are many approaches to functionalism, and it is hazardous to generalize. The gist of functionalism is this: a mental state does not depend on its internal composition but on how it functions or the role it plays for agents. It does not matter whether the actual stuff making these functions happen is a biological brain of a primate or the silicon chips of a computer. For a prominent version of functionalism, *see* DANIEL C. DENNETT, CONSCIOUSNESS EXPLAINED (1991); DANIEL C. DENNETT, THE INTENTIONAL STANCE (1989).

[40] Restatement (Second) of Contracts §12 cmt. e (1981); E. ALLEN FARNSWORTH & ZACHARY WOLFE, FARNSWORTH ON CONTRACTS 220-27 (4th ed., 2018); ANDREW BURROUGHS, A RESTATEMENT OF THE ENGLISH LAW OF CONTRACT 34 (2016).

is covered very briefly here. American Restatement (Second) of Contracts Section 15, entitled "Mental Illness or Defect," is illustrative. It provides that a person with a mental illness or defect lacks the capacity to contract and incurs only voidable contractual duties if she or he cannot "understand in a reasonable manner the nature and consequences of the transaction" or cannot "act in a reasonable manner in relation to the transaction and the other party has reason to know of his condition."[41] Moreover, when the contract is fair in its terms and the other party does not know of the mental illness or defect of the other party, the power of avoidance terminates to the extent the contract is performed or avoidance would be unjust because of changed circumstances.[42] Comment c to Section 15 explains that to prove incapacity it is essential to prove "irrational or unintelligent behavior" and that "almost any conduct of the person may be relevant."[43] These Restatement provisions illustrate that the law on incapacity to contract, consistently with the objective theory of contract, embeds a Turing-like test in its terms. They illustrate how contact law relies only on the external appearance of capacity.[44] The person whose capacity is in issue must be unable to understand or act "in a reasonable manner" to put the other party on notice of the incapacity. It removes the availability of the incapacity defense in situations in which the other party does not know the incapacity and the contract is basically fair and already performed. There is no mental state-type of evaluation because such an approach has no place in contract law. Capacity has to do, at least partly, with the manifestation of assent to contract, which, as explained

[41] Restatement (Second) of Contracts §15(1) (1982).
[42] *Id.*
[43] *Id.* §15(2).
[44] *Id.*, cmt. c. There seems to be some difference in English law, in which an individual lacks mental capacity if at the time of contracting she cannot decide for herself to enter the contract "because of an impairment of, or a disturbance in the functioning of, the mind or brain," but this can be read as agnostic on the question of the relevance of mental states and this principle is qualified by the fact that English law seems to require that the other party actually know of the incapacity. BURROUGHS, *supra* note xx. So, English law requires both (i) incapacity of the contracting party and (ii) actual knowledge of that incapacity of the other contracting party. Knowledge must necessarily be obtained by external means. English contract law rests on objective theory.

above, is determined through observable behavior and action. The appearance of assent to the other party can be determinative in any case in which incapacity is asserted as a defense.

In summary, nothing in contact law itself prevents an AGI from being a party to a contract. To the contrary, contract law supports AGI recognition as a contracting party. Of course, many practical difficulties will arise in cases in in which an AGI will purport to be a party to a contract but from the standpoint of contract law itself, any difficulties can be overcome. Putting doctrine into practice is often more difficult than conceptualizing how the law would address a particular social problem in theory.

II. AGIs as Bearers of Contract Rights and Liabilities

As explained above, a possible objection to the above argument for AGI capacity or eligibility as a contract party is that it is question begging or at the very least sets overly narrow parameters by focusing only on contract law and not on the necessary conditions for contract law to operate in the first place. The prior question, so the argument goes, is not about contract but about status, about the status of AGI as persons. Before we begin to talk about contracting by AGIs, we need to resolve whether an AGI is entitled to status or recognition as a person under the law. The preceding discussion assumes there is a candidate or subject who (or which) can have capacity to contract. This part will take up these broader and more policy-oriented questions. The focus will be on subjects who can possess contract rights and bear contract liability.

At the outset, rule out "artefact" arguments against recognition of AGI as subjects of the law.[45] That an AGI is not a natural person tells us nothing about legal status other

[45] *See* Stephen Wettig & Bernhard Zehendner, *A Legal Analysis of Human and Electronic Agents,* 12 Artificial Intelligence & L. 111, 123-24 (2004); Lawrence B. Solum, *Legal Personhood for Artificial Intelligences,* 70 N. C. L. Rev. 12311276-1279 (1992).

than that the law should not recognize an AGI as a natural person. That an AGI is not a natural person in the form of a human does not lead to the conclusion that it is to have no legal status at all. The law recognizes many kinds of persons that are not natural persons. Humans create other persons with no biology but with legal recognition, such as corporations. Biology is an arbitrary feature that cannot pick out the necessary conditions for personhood. In the future, some AGI may be biological in at least some or even its main features and humans may incorporate elements of machine intelligence into their brains and bodies.[46]

That an AGI is created by humans does not lead to the conclusion that an AGI should have no legal status as a person under the law or subject to the law. Several defects are apparent in this "creation" argument. Humans are created by humans too. The theological argument that humans are created by God cannot work. The proximate creators of humans are other humans at the present time though this may change.[47] AGIs are also created by God because everything is ultimately God's creation, if one accepts the teachings of mainstream religions. AGIs, moreover, will probably be able to create other AGIs in the future, as this is what being a general domain AI may entail.

That AI is "coded" in the sense that their intelligence may derive from a form of programming does not matter either for purposes of legal recognition of the status of AGI as subjects to which the law makes authoritative pronouncements. Theory and evidence about evolution informs us that humans are coded too.[48] The "technology" or

[46] This article focuses on machine intelligence. BOSTROM *supra* note xx at 26-62 examines the various potential paths to super-intelligence, including whole brain emulation, biological cognition, and brain-computer interfaces.
[47] Solum *supra* note xx at 1278.
[48] Evolutionary approaches are influential and perhaps dominant in the study of human cognition. This article work investigates human thought in a legal context but a look at how it has developed in the context of moral cognition may be revealing. For an overview in the context of the study of morality, *see* MICHAEL TOMASELLO, A NATURAL HISTORY OF HUMAN MORALITY 137-142 (2016). Wallach & Allen identify the differences in thinking about morality along the lines of top-down versus bottom-up. The top-down approach to understanding morality is the approach of traditional moral philosophy, focusing on its three branches of deontology, consequentialism, and virtue ethics. WALLACH & ALLEN, *supra* note

source of intelligence or intentionality should not be determinative of legal status or recognition.

But there is a human constraint on AGI recognition as a person for purposes of contract law. The constraint is not special to contract law, but contract law adds special elements to it. It derives from the proposition that law is what anthropologist Michael Tomasello characterizes as a "conventional cultural practice" that has evolved through time to reflect and enforce group social norms of a special kind.[49] For the law to be applied to both humans and AGI on terms of equal respect for each, AGIs will have to possess or be able to simulate a collective intentionality that so far has been special to humans.[50] In broad strokes and with some adaptation the argument proceeds as follows. Collective intentionality is necessary for the existence of conventional cultural practices such as law.[51] Law does not simply operate on human cognition; it is also created by it.[52] Humans have evolved so that their cognition includes substantial elements of group-mindedness and pro-sociality.[53] These features of human cognition gave humans the ability to construct cultures common in their groups through conventions that include social norms and institutions such as law.[54] Contract law, money, social norms on what constitutes fair exchange and the psychology of markets all play a role. It will

xx at 83-97. The bottom-up approach to understanding morality is the work in psychology, anthropology, and socio-biology, using empirical and experimental methods to explain the cognitive structure of human thinking about morality. The jury is still out on the approach to morality programming for AI. The answer will likely be driven by what is technologically plausible. *See* BOSTROM, *supra* note xx, at 30-32.

[49] MICHAEL TOMASELLO, A NATURAL HISTORY OF HUMAN THINKING 90-92 (2014).

[50] For a complementary approach that does not deal with a specific area of human activity such as contracting, *see* David J. Calverly, *Imagining a Non-Biological Machine as a Legal Person*, 22 ARTIFICIAL INTELLIGENCE & SOC. 523 (2008). On collective intentionality and humans, *see Id.* 80-123; PASCAL BOYER, MINDS MAKE SOCIETIES: HOW COGNITION EXPLAINS THE WORLD HUMANS CREATE 163-202 (2018); JOSHUA GREENE, MORAL TRIBES: EMOTION, REASON, AND THE GAP BETWEEN US AND THEM 1-27 (2013); SAMUEL BOWLES & HOWARD GINTIS, A COOPERATIVE SPECIES: HUMAN RECIPROCITY AND ITS EVOLUTION (2011).

[51] TOMASELLO, *supra* note xx at 90-92. A linkage may exist here to Scott Shapiro's planning theory of law. SCOTT SHAPIRO, LEGALITY 2011).

[52] BOYER, *supra* note xx.

[53] *See infra* notes xx and accompanying text.

[54] *Id.;* TOMASELLO, *supra* note xx.

be difficult for AGI to engage in contracting or to be subject to contract law if it cannot interact with humans as humans have evolved to interact and to be regulated by institutions humans have constructed to reflect the intentions humans hold in common. These cultural practices existed before many humans are born and yet they are subject to them and indeed cultural norms are transmitted intergenerationally through social learning.[55] So too something similar will have to occur or be simulated for AGI.

Lon Fuller presaged the contemporary research in psychology and anthropology in a way that is relevant to the question of AGI legal agency. He states the relationship of human agency to the law as follows:

> To embark on the enterprise of subjecting human conduct to the governance of rules involves of necessity a commitment to the view that man is, or can become, a responsible agent, capable of understanding and following rules, and answerable for his defaults.[56]

The question is, to Fuller, whether a person subject to the law "is, or can meaningfully strive to become, a responsible, self-determining center of action."[57] Reading the relevant passages in Chapter 4 of Fuller's *The Morality of Law* informs us that he was responding to the behaviorism of his time, most notably that of psychologist B.F. Skinner, for whom intentionality and related aspects of cognition were black boxes not to be explored.[58] Beyond our scope here is an exploration how Fuller's principle of self-determination does or does not reflect some form of intentionality. The point here is limited: that some "traditional" or non-naturalist legal philosophy may have something to tell us about who can be law's subjects.

Continuing on from Fuller's responsible agent condition with a more traditional philosophical approach, we can examine the question of extending law's rights and

[55] TOMASELLO, *supra* note xx, at 115.
[56] LON L. FULLER, THE MORALITY OF LAW 162 (2d ed., 1969).
[57] *Id.*, 163.
[58] *Id.*, 163-67.

liabilities to AGI through the lens of good old-fashioned moral philosophy about justification of the law. How might we justify the extension of the law of contract to AGIs? Justification of this kind involves investigating two relationships: (1) the relationship of the law, a social institution backed by the coercive power of the state, to bearers of its legal rights, duties, and obligations and (2) the relationship the law establishes between persons.[59] At least from a human perspective, we can rely on a contractualist approach, developed by T.M. Scanlon and others:[60] do humans have reasons to reject an AGI as having the status of a contract party, generally as a matter of principle and not specific to any transaction? AGIs may also deserve the same question to be answered if we are amenable to a level of abstraction that permits us to consider whether they are responsible agents who are capable of desert, blame, and other attributes associated with how the normative structure of the law has authority to regulate the persons it governs.[61]

We will likely reach the same conclusion by the philosophical route that we have reached by the anthropological and psychological route. If we want to justify the application of the law as it has been constructed by and for humans to AGI, then we need to develop AI with human-like values and dispositions – with human-like cognitive architecture or that at least simulates it with the ability to interact with humans, otherwise humans can reasonably reject the move. We need AI to have these qualities to be able to enter into and perform contracts with humans. A relationally focused cognitive architecture for exchange relationships with humans and for AGI to be recognizable –

[59] *See* John Finnis, *The Priority of Persons in* OXFORD ESSAYS IN JURISPRUDENCE FOURTH SERIES 1 (Jeremy Horder ed., 2000)(law as establishing relationship between persons); T.M. Scanlon, *Promises and Contracts, in* T.M. SCANLON, THE DIFFICULTY OF TOLERANCE: ESSAYS IN POLITICAL PHILOSOPHY 234 (2003)(on the relationship of law to persons); *see* KRISTEN RUNDLE, FORMS LIBERATE: RECLAIMING THE JURISPRUDENCE OF LON FULLER 99 (2012)(on law's relationship to human agency as understood by Fuller).

[60] *See* Scanlon, *supra* note __; T.M. SCANLON, WHAT WE OWE TO EACH OTHER (1998).

[61] *See* Luciano Floridi & J.W. Sanders, *On the Morality of Artificial Agents*, 14 MIND & MACHINE 349 (2004)(on a method of abstraction).

subject to or eligible for - recognition by the law, is necessary. Contract law exists to meet human needs in human societies for voluntary exchange. A sufficient level of mutuality in terms of autonomy, interaction, autonomy, and adaptability[62] would seem to be necessary for the extension of contract rights and liabilities to AGI to survive the test of reasonable rejection. Otherwise cooperation with humans will fail. Worse, an alien system of AGI values and cognitive abilities, based on norms humans do not understand and cannot reasonably accept, may be considered by humans to be harmful or pernicious, or may cause harm to humans.

What would be the properties of collective intentionality or responsible agency relevant to contracting by AGI? We can make a start on answering this question here, though more AI research would need to be done. AGI will need to have a cognitive architecture of voluntary exchange, understood within a social context of cooperation within groups. It would not be pure self-interest as understood in rational choice theory, long since discredited by advances in anthropology, psychology, and economics itself.[63] In Pascal Boyer's words: "Humans area immensely cooperative."[64] Tomasello explains, based in his many years of anthropological work, that humans differ from other apes because we are "group-minded."[65] Humans have only recently separated

[62] John Sullins reconceptualizes moral agency for AI around three criteria: (1) autonomy in the engineering sense, which means not under control of another agent or user; (2) intentionality in a weak AI sense, and (3) responsibility, satisfied with appearance and putting aside metaphysics. John Sullins, *When is a Robot a Moral Agent?*, 6 INT'L REV. INFO. ETHICS 23 (2006). Floridi and Sanders develop a moral agency for AI around the following three criteria: (1) interactivity, which means the agent and its environment can act upon each other; (2) autonomy, the agent can change state without direct response to interaction, and (3) adaptability, meaning the agent's interactions can change the transition rules by which it changes state. Floridi and Sanders add that an action qualifies as a moral if and only if it can cause moral good or evil. Floridi & Sanders, *supra* note xx. For a discussion of these and other conceptions of moral agency for AI, *see* DAVID J. GUNKEL, THE MACHINE QUESTION; CRITICAL PERSPECTIVES ON AI, ROBOTS, AND ETHICS 69-73 (2017).

[63] *See* BOWLES & GINTIS, *supra* note xx; Joseph Heinrich. Robert Boyd, Samuel Bowles, Colin Camerer, Ernst Feht, Herbert Gintis, & Richard McElreath, *In Search of Homo Economicus: Behavioral Experiments in 15 Small-Scale Societies*, 91 AM. ECON. REV. 73 (2001); BOYER, *supra* note xx, at 167.

[64] BOYER, *supra* note xx, at 164.

[65] TOMASELLO, *supra* note xx, at 82-90.

economic exchange from other social aspects of interaction.[66] An AGI would have to have the cognitive architecture associated with the psychology of transacting and exchange, set within the broader social context of how transacting and exchange connects to group identification. Self-interest would have to be represented in this architecture, but self-interest alone will be insufficient.

In the context of contracting and market transactions, AGI cognition will have to take on the first, second, and third person standpoints. It would have to recognize and consider in an appropriate way its own interests (first personal) and the interest of others on an individual-by-individual basis (second personal). It will have to be able to represent the world objectively (third personal), to possess cognitive skills to process information about the world abstractly, as what Tomasello characterizes as "transpersonally," in an agent-neutral way. [67] This third person perspective connects to the objective theory of contract because it includes the ability to process information beyond one's own particulars.

An AGI capable of contracting would have to have the ability to compare utilities at least at the level of the average human. AGIs must be able to measure respective utilities to be able to infer that gaining one thing is worth losing another.[68] They would have to have cognitive capacity to represent ownership cognitively, in terms of knowing what it owns and what it does not own and what passage of title or ownership means. An AGI would have to develop free rider and cheater detection capabilities.

Leda Cosmides has done experimental research on how humans have evolved and their specialized cognitive capacities to process information in the format of "benefit

[66] BOYER, *supra* note xx, at
[67] TOMASELLO, *supra* note xx, at 87-88, 116, 122, 139.
[68] BOYER, *supra* note xx, at 180.

received, cost not paid."[69] Cosmides and Tooby have researched how human cognition has an evolved specialization for deontic reasoning about social exchange between humans, which may be understood as a form of actually bargained for exchange between humans in the same social group, which can be broadly defined.[70] The cognitive architecture for social exchange reliably produces cooperation for mutual benefit between two agents. When the mind registers a situation in the form of a conditional rule, such as: "if X accepts a benefit from Y, it triggers an expectation that at some point that X will have to confer a benefit on Y," a social exchange relationship has been activated.[71] In these contexts, the human mind will apply deontic concepts of obligation and entitlement.[72]

Not all social exchanges will qualify as contracts under the law enforceable by the state. Moreover, market transactions themselves may differ from paradigmatic social exchange. But it is worth understanding social exchange because it serves as a basic template for transacting in goods and services.[73] The point here is basic: humans have just this cognitive structure connecting to transacting and agreeing and hence how contract law operates, and if AGIs are to participate in the social practices of contracts enforceable by law, they will have to share a similar cognitive structure with humans.

Cognitive capacities associated with obligation and entitlement not only connect to contract formation but to contract performance. To be able to transact in a contractual context, an entity must have the ability to detect cheaters and free riding.[74] The

[69] *Id.*, 181-82; Leda Cosmides, *The Logic of Social Exchange: Has Natural Selection Shaped How Humans Reason? Studies with the Wason Selection Task*, 31 COGNITION 187 (1989).
[70] Leda Cosmides & John Tooby, *Can a General Deontic Logic Capture the Facts of Human Moral Reasoning? How the Mind Interprets Social Exchange Rules and Detects Cheaters, in* MORAL PSYCHOLOGY VOL. 1: THE EVOLUTION OF MORALITY: ADAPTATIONS AND INNATENESS 53 (Walter Sinnott Armstrong ed., 2008).
[71] *Id.*
[72] *Id.*, 73.
[73] BOYER, *supra* note xx, at 193-196.
[74] *Id.*, 181-182.

underlying cognitive structure associated with opprobrium for cheating is as a conditional rule violation.[75] It is telling that accidental rule violations will not fully engage cheater detection cognition in humans, but international cheating certainly will.[76] While contractual liability is strict liability, fraud or misrepresentation liability is not in the main legal systems of the world and usually requires some form of proof of intent to defraud or deceive.[77]

Another set of cognitive capacities that will have relevance to contracting are those that relate to what is known in the psychological literature as pro-sociality.[78] Prosocial behavior informs us that reciprocity probably has a great deal of relevance to how human contractors understand bargained for exchange and deviations from contract performance requirements.

Finally, a cognitive architecture for engaging in repeat transactions would be necessary.[79] This would require the ability to recognize and act upon trust, reputation, and tit-for-tat.[80] Human cooperation in repeat transactions is a well-understood evolutionary stable strategy for humans.

We cannot provide a full design specification for AGI as a contracting party here. This article only serves to identify some important parameters for determining what might be needed for an AGI to be a competent contracting party. Some questions remain unanswered. An important set of questions have to do with whether to code AGI for human biases such as loss aversion or other biases that may provide us with a more

[75] Cosmides & Tooby, *supra* note xx.
[76] On accidental versus willful breaches, *see* PETER A. ALCES, A THEORY OF CONTACT LAW: EMPIRICAL INSIGHTS FROM MORAL PSYCHOLOGY 287 (2011); Cosmides & Tooby, *supra* note xx, at 102, 104.
[77] *See* FARNSWORTH, *supra* note xx, at 243-45.
[78] *See* BOYD, *supra* note xx, at 167-168; NATALIE HENRICH & JOSEPH HENRICH, WHY HUMANS COOPERATE; A CULTURAL AND EVOLUTIONARY EXPLANATION (2007); FOUNDATIONS OF HUMAN SOCIALITY: ECONOMIC EXPERIMENTS AND ETHNOGRAPHIC EVIDENCE FROM FIFTEEN SMALL-SCALE SOCIETIES (Joseph Henrich, Robert Boyd, Samuel L. Bowles, Colin Camerer, Ernst Fehr & David Gintis eds., 2004).
[79] BOYD, *SUPRA* NOTE XX, AT XX.
[80] *Id.*, 184,194.

negative framework for understanding how humans do not always act in ways that increase their welfare.[81] Should we program or mimic "disability" into an AGI so it will not be able to out-negotiate a human in the contracting sphere or have some strategic advantage as a result of superior cognitive abilities? As Bostrom explains in the context of ethics, "to the extent that ethics is a cognitive pursuit, a superintelligence could do it better than human thinkers."[82] Moreover, what if AGIs could someday produce other AGIs on their own, which could surpass human capabilities in transactional contexts? These are hard questions to answer about whether AGI will pose a "threat" to humanity and beyond our scope here.

Of course, *how* an AGI would come to obtain legal rights and liabilities as a matter of law is a different question. It would need some form of explicit statutory or other legal recognition, with perhaps some built in protections such as insurance or escrow requirements for assets.[83] These questions are about the conventional legal pragmatics once we have established that an AGI *can* qualify as a contracting party.[84] Legal systems have been down this route before, with the recognition of corporate legal personality. They will do it again.

CONCLUSION

The work of understanding the role of AI in contracting is at its very beginning stages. While limited domain forms of AI are currently deployed in smart contracting,

---

[81] This reflects an old debate in the behavioral sciences as to whether biases and heuristics decrease welfare or play a more positive and adaptive role. *See* Peter B.M. Vranas, *Gigerenzer's normative critique of Kahneman and Tversky*, 76 COGNITION 179 (2000).

[82] Nick Bostrom, *Ethical Issues in Advanced Artificial Intelligence*, *https://nickbostrom.com/ethics/ai.html* (accessed Feb. 21, 2019).

[83] *See* Solum, *supra* note xx, at 1245 on requiring insurance, among other legal requirements.

[84] *See* Ngaire Naffine, *Who are the Law's Persons? From Cheshire Cats to Responsible Subjects,* 66 MODERN L. REV. 346 (2003), which takes us to the philosopher John Dewey, *The Historic Background of Corporate Legal Personality*, 35 YALE L. J. 655, 660 (1926)("There is no general agreement regargding the nature *in se* of the jural subject; courts and legislators do their work without such agreement, sometimes without any conception or theory at all regarding their nature."); Note, *What We Talk About When We Talk About Persons; The Language of Legal Fiction*, 114 HARV. L. REV. 1745 (2001)

the potential for expanding AI engagement in contracting and commercial transactions will come with technological advances in AI. AGIs do not yet exist, but we must think now about how they will interact with humans when they do come to exist. Law is one of the most important normative domains for humans. If AGIs should have some form of legal agency in the world of commerce and exchange, law could impose a significant constraint on how their cognition and resulting behavior should operate in the world. In the context of investing AGI with the capacity to contract, contract law itself will face few problems in recognizing AGI as a contracting party, if AGI has the cognitive machinery by which to engage in contract relationships with humans. AGI will need to be able to interact and value in a way that is compatible to contract law's principles. These principles and the values they express are human values.[85] Arguments that AGI can have different motivations and psychologies[86] than humans seem impractical unless we want AGIs to have their own moral and political communities segregated from human ones, an approach prone to produce inter-group conflict. While AGI may be in the far future now, it will be far more practical to align their cognition to contract law and the values it represents than to change contract law and the humans who invented it. Significant work remains to be done.

[85] *See* Harry Surden, *Values Embedded in Legal Artificial Intelligence,* (U. of Colorado Law Legal Studies Research Paper No. 17-17, 2017), https://ssrn.com/abstract=2932333; Virginia Dignum, *Responsible Autonomy*, PROCEEDINGS OF THE 26TH INT'L JOINT CONFERENCE ON ARTIFICIAL INTELLIGENCE (IJCAI-17) 4698 (2017).
[86] *Cf. Bostrom, supra* note xx (on moral motivations).