CONTRACTING BY ARTIFICIAL INTELLIGENCE: OPEN OFFERS, UNILATERAL MISTAKES, AND WHY ALGORITHMS ARE NOT AGENTS

MATTHEW OLIVER*

On what basis are contracts negotiated and agreed by artificial intelligence programs, without direct human supervision, legally enforceable? The answer is not obvious, because a person cannot consent, or even appear to consent, to terms of which they are evidently unaware. This problem was vividly illustrated by a recent case, *Quoine v B2C2*, in which the result turned on whether contracting algorithms were deemed legal agents or mere tools of the people who use them. Several scholars have argued that AI-negotiated contracts are unenforceable within existing contract law doctrines and that we must turn to principles of agency law in order to enforce them. I argue that this is a mistake. Agency is not the right way to understand the legal relationship between an AI program and the person who uses it. This agency solution is also unnecessary, because AI-negotiated contracts are enforceable within existing contract law doctrines. We can explain why AI-negotiated contracts are enforceable by recognising that a person operating an AI contracting program make an open offer to contract on whatever terms the AI program agrees.

* Matthew Oliver (Hon.BA, BPhil, DPhil) is a student at Yale Law School. His research interests include ethics and the philosophy of contract and criminal law.

I INTRODUCTION

Artificial intelligence ('*AI*') has already transformed the way many industries negotiate and enter into contracts, but it has not yet transformed contract law. That transformation may be about to begin. Courts will soon have to decide whether AI programs that autonomously negotiate contracts should be treated as mere tools or as the legal agents of the people who use them. More fundamentally, courts will have to decide whether AI-negotiated contracts can be enforced within existing contract law doctrines. These are not idle questions, because they have serious practical implications for contractual interpretation and liability. Now is the time to develop answers to these questions because cases involving AI-negotiated contracts are starting to make their way into court.

On 24 February 2020, Singapore's Court of Appeals entered a final judgment in *Quoine Pte Ltd v B2C2 Ltd*, ending a multi-year dispute over the validity of automated trades of two popular crypto-currencies: Bitcoin (BTC) and Etherum (ETH).[1] Quoine operated a market on which traders could exchange BTC and ETH.[2] Due to a technical problem, Quoine's market making program was temporarily unable to access external data about the prices of BTC and ETH.[3] This had two consequences. First, it caused one trader, B2C2, to begin offering ETH at its backup price, which had been programmed into B2C2's algorithm as a safeguard.[4] Second, the price of BTC on the exchange fell, triggering margin calls on several trading accounts.[5] These accounts were automatically forced to sell BTC on the exchange at the 'best available price on the [p]latform,' which happened to be B2C2's backup price.[6] As a result, B2C2 bought millions of dollars' worth of BTC at a rate 250 times more preferable than the market rate.[7] All this trading was done by algorithms, and no one noticed what was happening right away. When they eventually noticed, the people running Quoine reversed the transactions on the ground that the

---

[1] *Quoine Pte Ltd v B2C2 Ltd* [2020] SGHC(I) 02 ('*Quoine*').
[2] Ibid 5.
[3] Ibid 13. The technical problem arose due to an attempt to update some of the passwords on the system.
[4] Ibid 14. This backup price is known as a 'deep price.' The court found that it was programmed into the algorithm to protect B2C2 in case of unexpected events on the exchange.
[5] Ibid 13–14.
[6] Ibid.
[7] Ibid 2. 3092.517116 BTC was debited from the Counterparties' accounts and credited into B2C2's account, while 309.2518 ETH was debited from B2C2's account and credited into the Counterparties' accounts. The historical conversion rates can be found at 'Bitcoin USD', *Yahoo Finance* (Web page, 31 January 2021) <https://finance.yahoo.com/quote/BTC-USD/history>; 'Ethereum USD', *Yahoo Finance* (Web page, 31 January 2021) <https://finance.yahoo.com/quote/ETH-USD/history/>.

transactions were clearly mistakes. B2C2 sued, claiming that the transactions were binding contracts.[8]

Should these contracts be enforced? Legally, that depends on whether these trading programs should be treated like tools used to communicate and enter into contracts, or like legal agents negotiating contracts on behalf of their employers. Under Singaporean law, a contract can be void for unilateral mistake only if the non-mistaken party had 'actual knowledge' of the mistake.[9] Yet only algorithms were aware of the disputed transactions at the time they occurred. Quoine argued the algorithm's knowledge of the trades could be enough to void the contract, because the algorithm could be treated as a 'legal agent'.[10] After all, they argued, if human traders 'had met on the floor of the exchange' and completed the contracts at issue, these contacts would be void for unilateral mistake.[11]

The court disagreed, following the reasoning of the lower court, which held that since the programs in question were 'deterministic' and 'do not have a mind of their own,' the court should treat them as 'mere machines.'[12] The relevant knowledge, both courts held, is the knowledge of the programmer.[13] The programmer was not aware of the mistake, because they were not aware of the transaction, so there could be no reversal for unilateral mistake. In the end, the court enforced the contracts because it concluded that these particular algorithms were mere machines, not legal agents. The court sidestepped the question of how to deal with more sophisticated AI

---

[8] *Quoine* (n 1) 15.

[9] In *Quoine* the court explains the Singapore law of unilateral mistake, which is slightly different from the Australian or US law. In Singapore, a transaction is legally void for unilateral mistake only if the non-mistaken party had 'actual knowledge' of the mistake and is voidable in equity if the non-mistaken party had constructive knowledge of the mistake and behaved unconscionably in relation to the mistake: at 38. For the purposes of the opinion only the legal doctrine was relevant since the court found no basis for concluding that B2C2 acted unconscionably: at 54. For a discussion to the US law of unilateral mistake, and its application in a case like *Quoine* see below Part V.

[10] These arguments are explained in the lower court's opinion, which was affirmed on appeal. See *B2C2 Ltd v Quoine Pte Ltd* [2019] SGHC(I) 03, 86 ('*B2C2*').

[11] Ibid 86. The same issue is discussed by the appellate court: *Quoine* (n 1) 46.

[12] See *B2C2* (n 10) 89. The court stated 'They are, in effect, mere machines carrying out actions which in another age would have been carried out by a suitably trained human. They are no different to a robot assembling a car rather than a worker on the factory floor or a kitchen blender relieving a cook of the manual act of mixing ingredients. All of these are machines operating as they have been programmed to operate once activated'. The appellate court explicitly endorsed and ultimately upheld this reasoning stating 'We agree [that] it would be wholly artificial to recast the relevant matrix of fact, which was one where the contracting parties did not in fact know beforehand that they were going to enter into the Trading Contracts or their terms[.] ... We also agree with the Judge that it follows from this that in cases where contracts are made by way of *deterministic* algorithms, any analysis concerning knowledge of a mistake or unconscionably taking advantage of one must be done by reference to the state of mind of the programmers of the algorithms at the time of the programming': *Quoine* (n 1) 46.

[13] *Quoine* (n 1) 46; *B2C2* (n 10) 88.

algorithms.[14] In Singapore and around the world the legal status of AI-negotiated contracts remains an unanswered question that is not going away.

The lower court, and the lawyers for Quoine, cited the work of several contract scholars who have recently argued that AI contracting programs should be treated as the legal agents of the people who use them.[15] These scholars argue that AI-negotiated contracts may be unenforceable under current contract law doctrines. It is easy to see why. To be bound by a contract a person must consent, or appear to consent, to the specific terms of that contract. If an AI program negotiates and completes a contract without human supervision, then there is no consent, and no appearance of consent, to the specific terms of that contract on the part of the people the contract proports to bind. Without consent there is no legally enforceable contract. These scholars propose to solve this problem by treating AI contracting programs as the legal agents of the people who operate them. A legal agent may negotiate and agree to a contract that binds their principal, even though the principal does not know or consent to the specific terms of that contract. By treating AI programs as legal agents, or quasi-agents, agency theorists think they can solve the problem of consent for AI-negotiated contracts.

In this paper I argue that agency is not the right way to understand the legal relationship between an AI program and the person who uses it. The moral and legal background to agency law is the presumption that each person is responsible for themselves and is not responsible of others. Agency law makes a limited exception to this background principle. Agency law is a poor foundation on which to build the law governing artificial intelligence, because AI programs are not legal persons that can be responsible for themselves. Applying agency law to AI programs would create perverse incentives for the oversight of AI programs and would allow the people who benefit from AI programs to avoid liability for some of the risks created by those programs. This matters because the way judges in the coming decades choose to understand AI-negotiated contracts will set important precedents for the way the law of artificial intelligence develops.

[14] The lower court noted that the law here is 'in a formative state': *B2C2* (n 10) 88.

[15] *B2C2* (n 10) 86 citing Lauren Scholz, 'Algorithmic Contracts' (2017) 20(2) *Stanford Technology Law Review* 128 ('*Algorithmic Contracts*'); Samir Chopra and Laurence White, *A Legal Theory for Autonomous Artificial Agents* (University of Michigan Press, 2011).

The agency solution is unnecessary because, properly understood, AI-negotiated contracts are perfectly enforceable within existing contract doctrine. AI programs are not the agents of contracting parties. Rather, AI programs are tools that fill in the specific content of offers made by contracting parties. Contract law already accepts that a person may validly offer to contract on terms that will be filled in by some future event. For example, I could consent to contract 'at market prices on 1 June' or on 'whatever terms John offers to Sally.' Just as I can legally consent to contract on 'whatever terms John offers to Sally,' I can legally consent to contract on whatever terms my AI program offers to you. This kind of underdetermined open offer is risky, but it can also be profitable.[16] The person who makes such an offer must accept the associated risk. I argue that this is a better and more practicable way of understanding the role of AI programs in contract negotiation. Understood this way, AI-negotiated contracts are perfectly enforceable within existing contract law doctrines.

There is a further advantage to this open offer understanding of AI-negotiated contracts. It allows courts to avoid looking to the inner workings of AI programs. The agency law interpretation requires courts to look into the way an AI program works to decide if it is sophisticated enough to be treated as an agent, and to decide whether it can be said to have had a particular belief or intention. A court could make this investigation if it had to, but these are difficult technical and philosophical issues. The open offer understanding of AI contracting avoids these issues. In theory anything could provide the content of an underdetermined open offer; it could be the market price on 1 January, or 'whatever terms John offers to Sally.' If a court applies the open offer interpretation of AI contracting, it can avoid looking into the inner workings of AI algorithms, because on this view it does not matter what kind of thing an AI program is or whether it has certain beliefs and intentions.

The approach we choose—the agency approach or the open offer approach—will have practical implications for contractual liability. As cases like *Quoine* show, the agency solution would allow companies that use AI contract programs to avoid liability for the mistakes those programs make. The open offer solution would require a party to accept liability for unprofitable contracts negotiated by its AI program. Requiring parties to accept liability is the better result,

[16] By an 'open offer' I mean an offer made to all members of the public, or to all the members of a group, rather than to a specific counterparty. See, eg, *Carhill v Carbolic Smoke Ball Co* [1892] EWCA Civ 1, 256 where the court interpreted an advertisement as an open offer conveying the power of acceptance to members of the public.

because parties using unpredictable new technologies like AI should be required to bear the risks created by those technologies. More broadly, the open offer solution sets a better precedent for the future. It avoids the precedent that AI programs can be treated as legal agents and allows the contract law to remain agnostic as to the legal status of AI programs.

The arguments in this paper address foundational principals of the common law of contracts. The paper draws primarily on the statutory and common law of Australia and the United States. However, the paper's conclusions should be applicable across common law jurisdictions. The paper proceeds as follows. In Part II, I begin by describing the current uses of autonomous AI programs to negotiate contracts and the laws specifically addressing electronic contracts. In Part III, I introduce the problem of consent identified by agency theorists and their proposed solution of treating AI programs as legal agents. In Part IV, I show why the agency solution is a mistake. I first show that AI-negotiated contracts are perfectly enforceable under current contract law doctrines by defending an open offer interpretation of AI contracting. I then argue that, despite some obvious similarities, AI programs are not akin to legal agents, and that it would be a mistake for courts to begin treating them as such. Ultimately the relationship between a person and their human agent is legally and morally unlike the relationship between a person and an AI program they use to negotiate contracts.

Finally, in Part V, I return to *Quoine* and the application of the doctrine of unilateral mistake to AI-negotiated contracts. I show how the open offer view of AI contracting provides a simple default rule for cases like *Quoine*: a contract cannot be voidable due to unilateral mistake where the mistake was made by an AI contracting program. This is not, as the *Quoine* court reasoned, because a counterparty would not be aware of the mistake, but because a mistake by an AI contracting program is not the kind of mistake that can ground a defence of unilateral mistake.[17]

[17] This is not to say that the court in *Quoine* was mistaken in its application of Singapore law. My point is about how an Australian or US court should handle a similar case should it arise in their jurisdiction.

II FACTUAL AND LEGAL BACKGROUND

The term 'artificial intelligence' connotes something mysterious, new, and exciting. Legal scholarship surrounding AI echoes this sense of mystery and excitement. To understand the legal significance of artificial intelligence, it is important to demystify it. The more mysterious and exciting we find AI, the more tempting it is to reach for radical and novel legal solutions. We should resist this temptation. Part of my purpose in this paper is to show that AI is already fairly commonplace in some industries and can be accommodated into our existing legal categories.

A *What is Artificial Intelligence?*

Much of the mystery surrounding artificial intelligence stems from a basic misunderstanding of what the term 'artificial intelligence' means. One might think that 'artificial intelligence' defines a distinctive new technology: a new kind of thing that has been brought into the world. This understanding of AI is misleading because the term 'artificial intelligence' does not describe the way a program works, but rather what that program is used to do. There is no formally accepted definition of artificial intelligence, but in general use the term refers to any of a wide range of technologies that are used to do the kind of work that otherwise requires human creativity and intelligence.[18] This functional understanding of artificial intelligence goes some way towards demystifying it.

There are many ways of designing computer programs to do work that otherwise requires human creativity and intelligence. Even a comparatively simple deterministic algorithm can do some kinds of work that otherwise requires human creativity and intelligence. AI programs range from these deterministic algorithms to highly autonomous 'black box' algorithms which behave in ways that cannot be explained by their designers even with the benefit of hindsight.[19] The more

[18] John Daintith and Edmund Wright, *A Dictionary of Computing* (Oxford University Press, 6th ed, 2008). This popular definition is frequently attributed to the AI pioneer John McCarthy, although the source of that attribution is unclear. McCarthy himself defines AI as 'the science and engineering of making intelligent machines,' while conceding that there is no 'solid definition' of intelligence that does not require an analogy to human intelligence: John McCarthy, 'What is Artificial Intelligence?' (Stanford University, 2007) 2.

[19] See Financial Industry Regulatory Authority, *Artificial Intelligence (AI) in the Securities Industry* (Report, June 2020) 3, 12 which explains '[Deep learning models] involve multiple layers and a dynamic, iterative learning process, where the internal learnings are opaque, making it difficult to identify the specific factors and their interrelationships that lead to the final outcome. Despite the challenges, these more complex AI models continue to garner interest from the industry because they offer the potential to be more powerful in identifying patterns and making more precise predictions relative to simpler models'.

autonomous and unpredictable programs, particularly black box algorithms, pose the most serious problems for conventional legal doctrines.[20]

What makes the more autonomous kinds of AI both very useful and potentially problematic is their capacity to solve problems without being told how to solve them.[21] For example, machine learning programs can identify patterns in huge quantities of data, and use those patterns to make predictions, without being programmed with the patterns in advance.[22] How this works more precisely is both complex and constantly changing. This complexity is a crucial point for lawyers to bear in mind. As far as is fair and possible, laws that apply to the use of AI programs should not require courts to understand the precise inner workings of these programs.

The potential problems which AI programs raise for the tort law have received considerable scholarly attention.[23] Tort liability for damages caused by self-driving cars has been much discussed.[24] Of greater immediate practical significance is malpractice liability for damage

---

[20] '[U]npredictability, and partially autonomous behavior may make it hard to verify compliance with, and may hamper the effective enforcement of, [legal rules]': European Commission, *Artificial Intelligence—A European Approach to Excellence and Trust* (White Paper, February 2020) 12.

[21] Ryan Calo, 'Robotics and the Lessons of Cyberlaw' (2015) 103 *California Law Review* 513, 532 identifies 'unpredictably useful behavior' as the ultimate goal of robotics.

[22] 'Machine learning is a field of computer science that uses algorithms to process large amounts of data and learn from it. Unlike traditional rules-based programming, [machine learning] models learn from input data to make predictions or identify meaningful patterns without being explicitly programmed to do so': Financial Industry Regulatory Authority, *Artificial Intelligence (AI) in the Securities Industry* (Report, June 2020) 2.

[23] See, eg, Karni Chagal-Feferkorn, 'The Reasonable Algorithm' (2018) 1 *Journal of Law, Technology and Policy* 111 which argues that the tort law should treat AI programs more like humans, specifically by adopting 'a "reasonable algorithm" standard to non-human decision makers—similar to the "reasonable person" or "reasonable professional" standard that applies to human tortfeasors'; Ryan Abbott, 'The Reasonable Computer: Disputing the Paradigm of Tort Liability' (2018) 86(1) *The George Washington Law Review* 1 which argues that tort law should treat AI programs more like humans particularly that 'where a supplier can show that an autonomous computer, robot, or machine is safer than a reasonable person, the supplier should be liable in negligence rather than strict liability. The negligence test would focus on the computer's act instead of its design, and in a sense, it would treat a computer tortfeasor as a person rather than a product'. See also Iria Giuffrida, 'Liability for AI Decision-Making: Some Legal and Ethical Considerations' (2019) 88(2) *Fordham Law Review* 439 which discusses the challenges AI programs raise for the tort law and proposing a framework for determining the basis of liability; Karni A Chagal-Feferkorn, 'Am I An Algorithm or A Product? When Products Liability Should Apply To Algorithmic Decision-Makers' (2019) 30 *Stanford Law & Policy Review* 61, 69 which proposes a method for determining when the tort law should treat an AI program like a human tortfeasor and when it should apply traditional product liability analysis according to which one 'analyzes the system's characteristics vis-à-vis the rationales behind the products liability legal framework, and identifies those associated with promoting said rationales versus ones adversely affecting the ability to accomplish them'. See also Yesha Yadav, 'The Failure of Liability in Modern Markets' (2016) 102 *Virginia Law Review* 1031 which argues that traditional legal standards for liability derived from the common law of torts are ineffective means of regulating markets dominated by autonomous trading programs and suggests institutional reforms to trading exchanges instead.

[24] See Bryant Walker Smith, 'Automated Driving and Product Liability' (2017) *Michigan State Law Review* 1, 2. Smith argues that 'the current product liability regime, while imperfect, is probably compatible with the adoption

caused by the use of AI programs in hospitals.[25] More generally, scholars have debated the best way to accommodate AI programs into existing legal categories.[26]

There has been comparatively little study of the legal consequences of the widespread use of AI programs for negotiating and consenting to contracts.[27] Yet, the use of AI programs to negotiate and consent to contracts is already widespread in large segments of the economy. While the tort law attempts to anticipate the coming of AI programs to roads and hospitals, AI contracting is already here. Lawyers and judges will have to catch up and adapt legal doctrines to this new contracting reality as cases of AI contracting go wrong and begin to make it into the courts.

B *Contracting by Artificial Intelligence*

Contract law has well-established requirements for creating a legally binding contract. The basic requirement is that two parties manifest their intent to be legally bound by the terms of an agreement in which each party receives something of value from the other.[28] Courts will look to

---

of automated driving systems'. See also Harry Surden and Mary-Anne Williams, 'Technological Opacity, Predictability, and Self-Driving Cars' (2016) 38 *Cardoza Law Review* 121. Surden and Williams explain why self-driving cars can be unpredictable and discussing the implications of self-driving cars for tort law.

[25] See Jason Chung and Amanda Zink, 'Hey Watson, Can I Sue You for Malpractice? Examining the Liability of Artificial Intelligence in Medicine' (2018) 11 *Asia-Pacific Journal of Health Law, Policy and Ethics* 51, which discusses the tort liability of operators of AI powered medical tools, particularly IBM's Watson for Oncology; Jane R Bambauer, 'Dr Robot' (2017) 51 *University of California, Davis Law Review* 383, which discusses 'whether health and medical AI should be regulated more like doctors than like devices, and what difference it would make'; W Nicholson Price II, 'Regulating Black-Box Medicine' (2015) 28(2) *Harvard Journal of Law & Technology* 419.

[26] See Sam N Lehman-Wilzig, 'Frankenstein Unbound: Towards a Legal Definition of Artificial Intelligence' (1981) 13 *Futures* 442, 451, which describes seven possible legal analogues for increasingly sophisticated AI programs, (1) products, (2) dangerous animals, (3) slaves, (4) adults with diminished capacity, (5) children, (6) agents, (7) persons, and noting that agency 'is the most comprehensive and germane with regard to both the essence and function of such a creature'; Matthew U Scherer, 'Of Wild Beasts and Digital Analogues: The Legal Status of Autonomous Systems' (2019) 19(1) *Nevada Law Journal* 259, 260, which discusses whether AI programs could be considered legal persons and concluding that 'agency law provides the most effective and flexible legal analogue for artificial systems'.

[27] For example, recent whitepaper released by the European Union on the legal implications of AI runs to 26 pages without mentioning contracts once. See European Commission, *Artificial Intelligence—A European Approach to Excellence and Trust* (White Paper, February 2020) 12.

[28] See *Uniform Commercial Code* (2012) § 2-204 ('*Uniform Commercial Code*'), which states: 'A contract for sale of goods may be made in any manner sufficient to show agreement, including conduct by both parties which recognize the existence of such a contract'; American Law Institute, *Restatement (Second) of Contracts* § 17 (1981). Of course, 'not mutual assent but a manifestation indicating such assent is what the law requires. Nor is it essential that the parties are conscious of the legal relations which their word or acts might rise to. It is essential, however, that the acts manifesting assent shall be done intentionally. That is, there must be *conscious will* to do those acts; but it is not material what induces that will. Even insane persons may so act; but a somnambulist could not': at § 20, Comment A (emphasis added). It may be best not to read too much into the choice of the phrase

see whether an offer has been made and whether that offer has been accepted.[29] The offer both proposes the terms of the agreement and gives another party the power to conclude the contract by accepting.[30]

Simple machines, such as vending machines, can be used to communicate an offer or acceptance. In one well-known case, *Thornton v Shoe Lane Parking Ltd* ('*Thornton*'), the court accepted that a parking garage operator could enter into a contract with a customer by using an automatic ticket machine to issue tickets to drivers entering the garage.[31] As Lord Denning reasoned,

> The customer pays his money and gets a ticket. ... It can be translated into offer and acceptance in this way. The offer is made when the proprietor of the machine holds it out as being ready to receive money. The acceptance takes place when the customer puts his money into the slot. The terms of the offer are contained in the notice placed on or near the machine stating what is offered for the money.[32]

The ticket machine is not 'a booking clerk in disguise' with whom the driver can negotiate a contract.[33] Rather, the ticket machine communicates an offer from the parking garage owner to

---

'conscious will.' The Restatement's authors are likely appealing to the distinction, familiar from the criminal law, between voluntary acts and involuntary bodily movements. We can apply this distinction to the workings of AI programs without considering whether such programs have a phenomenological experience of willing.

[29] *Courteen Seed Co v Abraham,* 129 Or 427, 427 (1929) states 'Contracts in general are reached by an offer on the one side and acceptance on the other'. See also*, Uniform Commercial Code* (n 28) § 2-206 'Unless otherwise unambiguously indicated by the language or circumstances an offer to make a contract shall be construed as inviting acceptance in any manner and by any medium reasonable in the circumstances.'; American Law Institute, *Restatement (Second) of Contracts* (n 28) § 22, 24 defines an offer as 'the manifestation of willingness to enter into a bargain, so made as to justify another person in understanding that his assent to that bargain is invited and will conclude it'.

[30] See Arthur Corbin, 'Offer and Acceptance and Some of the Resulting Legal Relations' (1917) 26(3) *Yale Law Journal* 169, 171 ('Offer and Acceptance'), which states 'an offer is an act on the part of one person whereby he gives to another the legal power of creating the obligation called contract. An acceptance is the exercise of the power conferred by the performance of some act or acts. Both offer and acceptance must be acts expressing assent.' For the *locus classicus* of the concept of a legal 'power' see Wesley Newcomb Hohfeld, 'Some Foundational Legal Conceptions as Applied in Judicial Reasoning' (1913) 23(1) *Yale Law Journal* 16.

[31] *Thornton v Shoe Lane Parking Ltd* [1971] 1 All ER 686 ('*Thornton*'). The central issue in the case was whether a contractual condition disclaiming liability for personal injuries to the customer could be incorporated into the parking contract by text printed on the ticket issued by the machine. This posed a dilemma, because English common law held that conditions printed on tickets could be binding contractual terms because a person receiving a ticket from a human agent remained free to reject the ticket if he did not like the terms. Lord Denning reasoned that the case of an automatic ticket machine was different because the ticket could not be returned once it had been issued. Thus, whether the terms on the ticket itself were a part of the contract depended on whether the ticket was issued by a human agent or a machine. In *Thornton*, as in *Quoine*, the legal significance of the difference between a contracting agent and a contracting tool can crop up in surprising ways.

[32] Ibid 689.

[33] Ibid.

the driver, and then provides the means by which the driver can communicate acceptance of the offer: by paying for a ticket.

Computers, communicating over the internet, are used in the same way.[34] This kind of computer contracting is not without problems. It can be factually difficult to know whether an online offer has been read and it can be legally difficult to know whether unread online contracts should be enforced.[35] So-called browse-wrap contracts proport to bind consumers to terms in agreements that are disclosed in a hyperlink.[36] So-called 'click-wrap' or 'scroll-wrap' contracts require consumers to click 'I agree' and perhaps to scroll through the agreement before clicking.[37] Evidence suggests that 'consumers seldom read Internet contracts' even though they 'contain many controversial provisions.'[38]

Computers can also run complex algorithms which can themselves be incorporated into contracts. For example, two parties might negotiate and agree on an algorithm to determine how profits from their joint venture are allocated. As Shawn Bayern argues 'there is no conceptual difference, other than perhaps the level of precision, between an agreement that sets up

---

[34] The question of whether electronic means, whether an autonomous AI program or a simple algorithm, can be used to complete contracts was debated by scholars at the beginning of the 21st century. See Tom Allen and Robin Widdison, 'Can Computers Make Contracts' (1996) 9(1) *Harvard Journal of Law and Technology* 25, 49, which argues that 'neither American nor English law, as they currently stand, would confer legal status on all computer-generated agreement,' and highlighting the problem of consent with autonomously or even somewhat autonomously computer negotiated contracts; Anthony J Bellia Jr, 'Contracting with Electronic Agents' (2001) 50 *Emory Law Journal* 1047, which discusses whether electronically completed transactions, whether autonomous or not, are enforceable under the common law and whether they ought to be enforceable under the leading contract law theories: will theory, consent theory, reliance theory, economic efficiency, and trust.

[35] The traditional common law view is that there is a duty to read contracts before signing them. See, eg, *Upton v Tribilcock*, 91 US 45, 50 (1875) where Justice Hunt stated: '[I]f he will not read what he signs, he alone is responsible for his omission'. However, this traditional view has recently come under attack in the context of consumer contracts completed on websites. See, Ian Ayres and Allen Schwartz, 'The No-Reading Problem in Consumer Contract Law' (2014) 66(3) *Stanford Law Review* 545, which argues that mass-market sellers should be required to determine which terms in its user agreements are not expected by consumers and to highlight those unexpected terms so that they are drawn to the attention of consumers; Meirav Furth-Matzkin and RoseAnnie Sommers, 'Consumer Psychology and the Problem of Fine Print Fraud' (2019) 72 *Stanford Law Review* 2, studies the effect on consumers of false or fraudulent statements made by sellers which are directly contradicted by the fine print of a contract; Florencia M Marotta-Wurgler, 'Will Increased Disclosure Help? Evaluating the Recommendations of the ALI's Principles of the Law of Software Contracts' (2011) 78 *University of Chicago Law Review* 165, surveys consumer behaviour to argue that mandating increased contract disclosures—as proposed by the ALI's draft principles—would not meaningfully promote comparison shopping or alter consumer behaviour; Yannis Bakos, Florencia Marotta-Wurgler, and David R Trossen, 'Does Anyone Read the Fine Print? Consumer Attention to Standard-Form Contracts' (2014) 43 *Journal of Legal Studies* 1, 1 which finds that 'only one or two out of every thousand retail software shoppers chooses to access the license agreement, and those few that do spend too little time, on average, to have read more than a small portion of the license text'.

[36] See Ayres and Schwartz (n 35).

[37] See ibid 548.

[38] See ibid 547.

procedures to manage performance and an algorithm.'[39] In other words, because an algorithm is simply a set of rules that produces an output from a set of inputs, we can understand a contract as an algorithm that specifies the rights and obligations of parties given different inputs. We could even use a computer algorithm as a contract, because algorithms and contracts can be functionally equivalent.

This already happens in the case of so-called 'smart contracts.' These 'smart contracts' are computer algorithms that automatically transfer property, such as digital currency, from one party to another according to an algorithm to which the parties have agreed.[40] A 'smart contract' is thus both algorithm and contract at the same time.[41] There is some dispute as to whether these smart contracts are properly understood as contracts in the legal sense.[42] The parties who implement these contracts often see them as a way of avoiding recourse to the legal system and explicitly deny that smart contracts are contracts in the legal sense.[43] More generally, as Kevin Werbach and Nicolas Cornell point out, there is reason to doubt whether self-executing smart contracts are properly understood as contracts because they are self-executing and therefore do

[39] Shawn Bayern, 'Algorithms, Agreements, and Agency' in Woodrow Barfield (ed), *The Cambridge Handbook of the Law of Algorithms* (Cambridge University Press, 2020) 153, 154.

[40] See Kevin Werbach and Nicolas Cornell, 'Contracts Ex Machina' (2017) 67 *Duke Law Journal* 313, 356 (discussing the rise and use of 'smart contracts' and discussing whether they can solve the issues of trust and reliance that the contract law has historically addressed. 'Smart contracts make it the case that promisors will do precisely what they promise, radically strengthening promises. If this is the point of judicial contract enforcement, then it looks like smart contracts offer a superior technology, and smart contracts would leave judicial enforcement essentially obsolete.'). See also 'Community Guides and Resources', *Ethereum* (Web Page, December 2020) <https://ethereum.org/en/learn/>: 'A "smart contract" is simply a piece of code that is running on Ethereum. It's called a "contract" because code that runs on Ethereum can control valuable things like [crypto-currency] or other digital assets'.

[41] See Werbach and Cornell (n 40) 343 ('To sum up, smart contracts are contracts. They are agreements to shift legal rights and responsibilities, no less than an agreement between two parties physically exchanging goods for payment over a counter. Their status as contracts might be obscured by the fact that the parties intend litigation to be impossible, may not make any promise, and may be expressed only in code. We suggest that these details do not alter the fact that smart contracts are, indeed, contracts in the important sense.'); Bayern (n 39) 154 (arguing that algorithms can be both contract and algorithm because the two are functionally equivalent.) See also Scholz, 'Algorithmic Contracts' (n 15) 145–9 discussing the history of Ethereum and smart contracting.

[42] See Scholz, 'Algorithmic Contracts' (n 15) 149 (who states that contracts created on these decentralized platforms are 'probably not enforceable').

[43] Andreas Antonopoulos and Gavin Wood, 'What is a Smart Contract' *Ethereum Book* (Web Page, December 2020) <https://github.com/ethereumbook/ethereumbook/blob/develop/07smart-contracts-solidity.asciidoc#what-is-a-smart-contract> which states: 'In the context of Ethereum, the term is actually a bit of a misnomer, given that Ethereum smart contracts are neither smart nor legal contracts.'

not impose obligations on parties.[44] These are important questions, but they are for a different paper.

This paper focuses on a slightly different use for algorithms in contracting, which raises a distinctive set of problems for conventional contract law theory. I do not focus on the use of algorithms as contracts, or as a part of contracts. The problem I focus on is the use of AI programs to propose or negotiate the terms of contracts and then agree to those terms autonomously. As I will use the term, a contract is 'negotiated' by an AI program when the AI program, rather than the human who programs it, sets the precise terms of an offer to be made or accepted. An AI program consents to a contract when it makes an offer that is accepted, or when it accepts an offer. As several scholars have recently argued, it is unclear whether these contracts are legally enforceable under current contract law doctrines. Yet these contracts are also ubiquitous.[45]

The kind of algorithmically negotiated contract most often encountered by consumers are contracts negotiated by dynamic pricing algorithms. Large companies, such as airlines and online retailers, use new technologies to constantly update the prices of products and services offered for sale on their website.[46] These algorithms may in some cases be quite simple, but sophisticated AI algorithms can be used as well. Dynamic pricing allows companies to rapidly adapt their prices in response to market conditions, inventory, and perceived demand in order to maximise profitability.[47] Retailers with access to personal information can use sophisticated algorithms to customise prices based on the profile of an individual customer.[48] Personalised data can include information about the wealth, credit history, and preferences of the person.[49]

---

[44] To be clear, Werbach and Cornell ultimately concluded that smart contracts are legal contracts. In making this point they also offer and reject some good reasons to think that they might not be.

[45] See Lauren Henry Scholz, 'Algorithms and Contract Law' in Woodrow Barfield (ed) *The Cambridge Handbook of the Law of Algorithms* (Cambridge University Press, 2020) 141, 141 ('*Algorithms and Contract Law*') suggesting that we are moving 'from an economy built on form contracts to an economy built on algorithmic contracts.'

[46] See Robert Weiss and Ajay Mehrotra, 'Online Dynamic Pricing: Efficiency, Equity and the Future of E-Commerce' (2001) 6 *Virginia Journal of Law and Technology* 11, 11 which describes the rise of dynamic pricing and some of the technologies that enable it to be far more sophisticated.

[47] This does not necessarily require AI if the adaptation is simple enough. For example, a vending machine selling cool drinks could be programmed to charge a price based on the ambient temperature. See ibid 11.

[48] See Akiva Miller, 'What Do We Worry About When We Worry About Price Discrimination? The Law and Ethics of Using Personal Information for Pricing' (2014) 19 *Journal of Technology Law and Policy* 41, which discusses various data driven retail strategies and the concerns they should raise for consumers.

[49] See ibid 50. 'Each data broker has its own special recipe for combining data sources and marketing strategies: Acxiom Corporation, one of the largest data brokers, obtains its information from publicly available records, commercial entities, and retailer records of consumers' purchases. It also reportedly collects information on user

This personalisation raises concerns about privacy and about the ability of the most sophisticated retailers to take advantage of consumers.[50]

The use of AI programs to negotiate and complete contracts is already widespread in the securities industry. Many of the uses of AI programs are behind the scenes, allowing traders to conduct research, manage risk, supervise employees, and communicate with clients.[51] However, AI programs are also used to directly initiate and complete trades of financial products on securities markets.[52] These programs are known as autonomous trading systems. For example, Citi-group's ominously named Dagger system 'can stalk through more than 20 markets, public and otherwise—hunting anomalies, buying and selling, prowling through mountains of historical data.'[53]

AI programs allow for high-speed trading of financial products.[54] High-frequency traders use AI programs that can recognise opportunities and execute trades in as little as 10 milliseconds, far faster than a human could possibly react.[55] These traders pay to co-locate with the exchanges on which they trade in order to profit on millisecond advantages in available information.[56]

---

queries and habits from search engines and social media. Credit reporting agency Experian focuses its data-gathering attention on consumer's "life-event triggers"; Equifax, also a credit reporting agency, uses detailed salary and paystub information to profile consumer wealth. Targeted advertising provider Epsilon discovers people's interests through their reading habits and cause donations, among other sources. Datalogix collects vast information on consumer purchases from store loyalty cards and retailer records, including information on health-related purchases.'

[50] See ibid 41; Rory Van Loo, 'Helping Buyers Beware: The Need for Supervision of Big Retail' (2015) 163 *University of Pennsylvania Law Review* 1311, arguing that 'existing institutions are insufficient to prevent harm' to consumers and that the costs of failing to increase regulation and oversight could be significant.

[51] Financial Industry Regulatory Authority (n 19) 5–10.

[52] See Gregory Scopino, 'Do Automated Trading Systems Dream of Manipulating the Price of Futures Contracts? Policing Markets for Improper Trading Practices by Algorithmic Robots' (2016) 67 *Flordia Law Review* 221 (discussing the difficulty of identifying scienter when autonomous algorithms engage in seemingly illegal trading strategies on derivatives markets).

[53] Timothy Lavin, 'Monsters in the Market', *Atlantic* (News Article, June 2010), <http://www.theatlantic.com/magazine/archive/2010/07/monsters-in-the-market/308122/>. See also Scopino (n 52).

[54] Andrew Keller, 'Robocops: Regulating High Frequency Trading After the Flash Crash of 2010' (2012) 73 *Ohio State Law Journal* 1457, 1461–71 which gives an excellent explanation of the need for speed in high frequency trading and several high frequency trading strategies.

[55] See Scopino (n 52) 226.

[56] See Keller (n 54) 1426. 'The shorter the length of the cable connecting traders to exchanges' matching engines, the faster the trade can be executed. Trading firms have recognized this fact, as well as the exchanges themselves. Today, exchanges offer space as close as physically possible to their matching engines, where firms can locate their own servers. This scheme is known as "co-location," and a single rack on which to place a server can cost $10,000 per month.'

This kind of high-speed trading is now estimated to make up more than half of all the trading in American financial markets.[57] Case law on high-frequency trading is virtually non-existent.[58] Yet litigation regarding high-speed trading may be on its way. These high-frequency trading strategies are controversial and have received considerable scholarly and regulatory attention.[59] Such AI programs have been linked to several 'flash crashes.'[60] Some scholars have also argued that the use of AI trading programs undermines the efficiency of capital markets.[61]

High frequency trading systems have been accused of 'wash trading' (buying and selling a derivative at the same price in order to manipulate the market),[62] as well as 'spoofing' (the practice of bidding and then cancelling the bid or offer before execution)[63] and 'banging the close' (buying or selling in the last moments of trading).[64] It is difficult for regulators to police these trading practices because often such trades are only illegal when done with some sort of scienter, or intent. This is hard to prove, because humans who program these algorithms do not intend to enter into any particular trades.[65] In fact, the humans may not even be aware of what trading strategies their algorithms will pursue.

There is a theme running through all these examples of AI contracting. Whether they are contracting with retail consumers or trading on financial markets, companies can use innovations in AI technology to negotiate more profitable contracts for themselves while creating new and unpredictable risks for their counterparties and for the market as a whole. These

---

[57] Charles Korsmo, 'High-Frequency Trading: A Regulatory Strategy' (2014) 48 *University of Richmond Law Review* 523, 574–5. 'Despite the small number of high frequency traders, they account for more than half of all trades on US exchanges (although a slightly lower percentage than a few years ago). Because high frequency trade strategies typically involve unusually high order-to-trade ratios—with the vast majority of orders being cancelled without being executed—high frequency trades account for an even higher proportion of total data traffic on exchanges.'

[58] Scholz, 'Algorithmic Contracts' (n 15).

[59] Merritt Fox et al, 'The New Stock Market: Sense and Nonsense', (2015) 65 *Duke Law Journal* 191, 226–61 which argues against premature regulation of highspeed trading; Wallace Turbeville, 'A New Perspective on the Costs and Benefits of Financial Regulation: Inefficiency of Capital Intermediation in a Deregulated System' (2013) 72 *Maryland Law Review* 1173, which defends the need for greater regulation in response to algorithmic trading; Jerry Markham, 'High-Speed Trading on Stock and Commodity Markets— From Courier Pigeons to Computers', (2015) 52 *San Diego Law Review* 555, 563, which describes algorithmic highspeed trading and defends it as a natural and inevitable advance in trading.

[60] United States Securities and Exchange Commission, *Findings Regarding the Market Events of May 6, 2010* (Report, 30 September 2010) 1–4.

[61] See Yesha Yadav, 'How Algorithmic Trading Undermines Efficiency in Capital Markets', (2015) 68 *Vanderbilt Law Review* 1607, arguing that while 'algorithms help markets make gains on several measures of informational efficiency, they also create costs for their ability to allocate capital productively'.

[62] Scopino (n 52) 229. Wash trading is prohibited under the *Commodity Exchange Act*, 7 USC § 6c(a)(1)-(2) (2012).

[63] This is also prohibited by the *Commodity Exchange Act*, 7 USC § 6c(a)(5)(C) (2012).

[64] Scopino (n 52) 231.

[65] Scopino (n 52) 229–30.

companies can do this in part because of a comparative lack of regulation, which allows the operators of AI contracting algorithms to enjoy the benefits of these programs while passing some of the risks on to others. More fundamentally, these companies are able to profit from the use of AI because AI-negotiated contracts are assumed to be enforceable. As Scholz writes, 'It is contracts that give algorithms the power to change our world.'[66] However, as we will now see, there is some reason to doubt whether AI-negotiated contracts are legally enforceable.

C *The Law of Electronic Contracts*

There is currently a limited body of law that specifically addresses the use of computers to make contracts in Australia and the United States.[67] These laws authorise the use of computers to communicate the terms of an offer and make a record of an acceptance, much like a sophisticated version of the ticket machine in *Thornton.*[68] Just as it is possible to communicate an offer and record a legally binding acceptance by means of an automatic ticket machine, it is also possible to do so over the internet. At the time these laws were passed it was thought necessary to clarify this point, in order to give security to the growing use of computers to communicate contract offers and acceptances.[69] However, even without these laws, contracts communicated and accepted electronically would almost certainly have been upheld under the classic doctrines of contract law.[70] The only substantive effect of these laws was to permit the use of digital written records and signatures as a substitute for paper written records and signatures, as required for some transactions.[71]

---

[66] Scholz, 'Algorithmic Contracts' (n 15) 131.

[67] These laws enact, and are in addition to, international standards for giving effect to contracts communicated and signed by electronic means: *UNICITRAL Model Law on Electronic Commerce of 1996*; United Nations Convention on the Use of Electronic Communications in International Contracts of 2005; the *UNCITRAL Model Law on Electronic Signatures of 2001 and the United Nations*, GA Res 51/162, UN Doc E.99.V.4 (1997, adopted 1997).

[68] *Thornton* (n 35).

[69] See *Uniform Electronic Transactions Act* (1999) ('*UETA*'), Prefatory Note 312/915-0195, which states that '[w]hether the legal requirement that information or an agreement or contract must be contained or set forth in a pen and paper writing derives from a statute of frauds affecting the enforceability of an agreement, or from a record retention statute that calls for keeping the paper record of a transaction, such legal requirements raise real barriers to the effective use of electronic media ... By establishing the equivalence of an electronic record of the information, the *Uniform Electronic Transactions Act* ('*UETA*') removes these barriers without affecting the underlying legal rules and requirements.'

[70] See Corbin, 'Offer and Acceptance' (n 30) 99, which states '[t]he offeror is the creator of the power and at the time of its creation he has full control over both the fact of its existence and its terms.'

[71] See, eg, *Uniform Commercial Code* (n 28) § 2-201(1) which provides that '[e]xcept as otherwise provided in this section a contract for the sale of goods for the price of $500 or more is not enforceable by way of action or defence

The most prominent laws authorising the use of computers to make contracts in the United States are the *Electronic Signatures in Global and National Commerce Act of 2000* ('*E-Sign Act*') and the *Uniform Electronic Transactions Act* ('*UETA*'). The *E-Sign Act* is a federal statute which specifies that '[n]notwithstanding any statute, regulation, or other rule of law' a signature or contract 'may not be denied legal effect ... solely because it is in electronic form'[72] The *UETA* is a model law drafted by the National Conference of Commissioners on Uniform State Laws in 1991. It provides that 'a contract may not be denied legal effect or enforceability solely because an electronic record was used in its formation.'[73] Together, the *E-Sign Act* and the *UETA* give electronic records and signatures the same legal status as paper records and signatures. The equivalent Australian law is the *Electronic Transactions Act 1999* (Cth) ('*ETA*'), which similarly ensures that 'a transaction is not invalid because it took place by means of one or more electronic communications.'[74]

Neither the *E-Sign Act*, the *UETA*, nor the *ETA* proport to alter the basic requirements of contract law.[75] They allow for contracts to be made by electronic means, but they do not and could not make every electronic communication into a contract. An electronic communication makes a contract only if it represents an agreement between parties to alter their legal rights and obligations. As the *UETA* notes in its definitions section, an agreement is a 'the bargain of the

unless there is some writing sufficient to indicate that a contract for sale has been made between the parties and signed by the party against whom enforcement is sought or by his authorized agent or broker.'

[72] *Electronic Signatures in Global and National Commerce Act* , Pub L No 106–229 § 101(a) (2000) ('*E-Sign Act*'), 114 Stat 464. The *E-Sign Act* pre-empts state law requiring a written record of a transaction, except for state laws enacting the Uniform Electronic Transaction Act or otherwise allowing the use of electronic signatures and records to make contracts. See *E-Sign Act* § 102(a)(1) which says that '[a] State statute, regulation, or other rule of law may modify, limit, or supersede the provisions of section 101 with respect to State law only if such statute, regulation, or rule of law constitutes an enactment or adoption of the Uniform Electronic Transactions Act' or specifies alternative requirements for the validity of electronic signatures and records not inconsistent with this title or 'if enacted or adopted after the date of the enactment of this Act, makes specific reference to this Act.' See generally Stephanie Lillie, 'Will E-Sign Force States to Adopt UETA?' (2001) 42(1) *Jurimetrics* 21, 23.

[73] *UETA* § 7. The *UETA* has been adopted by nearly every state. The few exceptions (Illinois, New York, and Washington) have adopted other statutes to allow for the use of electronic signatures and records. See Scholz, 'Algorithms and Contract Law' (n 45).

[74] *UETA* § 8(1).

[75] See *E-Sign Act* § 101(b) which says that '[t]his title does not (1) limit, alter, or otherwise affect any requirement imposed by a statute, regulation, or rule of law ... other than a requirement that contracts or other records be written, signed, or in nonelectronic form'; *UETA*, Prefatory Note which says '[i]t is important to understand that the purpose of the *UETA* is to remove barriers to electronic commerce by validating and effectuating electronic records and signatures. It is NOT a general contracting statute—the substantive rules of contracts remain unaffected by *UETA*' (emphasis in original); ETA, *Definitions* which says that '[t]ransaction includes: (a) any transaction in the nature of a contract, agreement or other arrangement.'

parties in fact, as found in their language or inferred from other circumstances.'[76] The comments clarifying this definition refer explicitly to the basic requirements of contract law as described in the Restatement (Second) of Contracts and the Uniform Commercial Code.[77] In the preface, the drafters note that 'the substantive rules of contract remain unaffected ... [w]hether a record or signature is attributed to a person is left to law outside this act.'[78] Thus, while the *E-Sign Act*, the *UETA*, and the *ETA* guarantee equal status to electronic and paper contracts, they do not alter the basic requirements of contract law.[79] As we will see, there is reason to doubt whether contracts negotiated and agreed by autonomous AI programs satisfy these basic requirements.

III THE PROBLEM OF CONTENT AND THE AGENCY LAW SOLUTION

Some scholars argue that contracts negotiated and completed by AI programs may not be legally enforceable, at least under existing contract law doctrines.[80] It seems clear that AI contracts ought to be enforceable. They are widely assumed to be enforceable and this assumption underpins large segments of the economy. So, the problem for contract law is not so much that these contracts are doctrinally unenforceable but rather that contract law doctrines must adapt to explain why these contracts are enforceable. To remain relevant, contract law theory must keep up with contract law practice.

---

[76] *UETA* § 2(1).
[77] *UETA* § 2 Comment.
[78] *UETA*, Prefatory Note.
[79] The *UETA* does specify 'that the actions of machines ("electronic agents") programmed and used by people will bind the user of the machine, regardless of whether human review of a particular transaction has occurred.' *UETA*, prefatory note. The kind of electronic contracting without "human review of a particular transaction" which the act envisages is simply the internet equivalent of the ticket machine in Thornton. For example, orders placed through a website may be binding although they are not individually reviewed. In the same way, a human does not review the tickets issued by the ticket machine in Thornton each time one is issued. I agree with Scholz that the *UETA* does not anticipate the use of autonomous AI programs to make contracts, and does not address the unique challenges that these programs pose for contract law doctrines. See Scholz, 'Algorithms and Contract Law' (n 45) 145–6 which cites the history and text of the *UETA* to demonstrate that 'algorithmic contracts formed using machine-learning algorithms are outside the scope of the *UETA*.' The Australian ETA contains a similar provision. See ETA, 15(C) which states that '[a] contract formed by: (a) the interaction of an automated message system and a natural person; or (b) the interaction of automated message systems; is not invalid, void or unenforceable on the sole ground that no natural person reviewed or intervened in each of the individual actions carried out by the automated message systems or the resulting contract.' Like the equivalent provision of the *UETA* this language gives certainty to the enforcement of contracts formed by the computer equivalent of the ticket machine in Thornton. This language does not resolve the problem of consent that is presented by more sophisticated AI contracting problems. As we will see, the objection to the enforceability of AI-negotiated contracts is not made 'on the sole ground that no natural person reviewed' each transaction. Rather, the objection is that there is no contract at all because the basic requirement of consent to the terms of the contract have not been met.
[80] See n 82.

To solve this problem, several theorists have recently proposed that AI programs negotiating and completing contracts should be considered the agents of the legal persons (usually corporations) who operated them.[81] In this section, I lay out the problem and agency law solution. In the next section, Part IV, I argue that this agency law solution is a mistake.

A *The Problem of Consent*

AI-negotiated contracts are increasingly common and ought to be legally enforceable. But are these contracts enforceable according to our best theories of contract law? Several leading scholars have recently argued that they are not.[82] They claim that 'there is a very strong argument for the unenforceability of algorithmic contracts.'[83] This is not a merely theoretical problem. Even the potential that these contracts might not be enforced if challenged in court could create uncertainty in large segments of the economy.

In order to enter into a legally enforceable contract, both parties must manifest an intention to be bound by the terms of the contract.[84] The intent manifested cannot simply be an

[81] See n 93.

[82] See Scholz, 'Algorithmic Contracts' (n 15) 128, which says that '[w]hen algorithms take on a role in contract formation analogous to that of human agents, they should be considered constructive agents for the purpose of contract formation. The company consenting to the contract can be said to have authorized or ratified the contract formed on its behalf by the algorithm. This approach explains easy cases while also showing why algorithmic contracts, even many black box algorithmic contracts, are enforceable.'; Tom Allen and Robin Widdison, 'Can Computers Make Contracts?' (1996) 9(1) *Harvard Journal of Law and Technology* 25, 49, which states '[s]ince the program changes over time, without any human intervention, it would be very difficult to characterize it as the embodiment or expression of human intention. Hence, the doctrine as it now stands would deny validity to agreements generated by an autonomous computer.'; Chopra and White (n 15) 29, which says that '[t]he status of these contracts, however, is problematic, for artificial agents are considered mere things, not persons, in the law's ontology.'; Ian R Kerr, 'Spirits in a Material World: Intelligent Agents as Intermediaries in Electronic Commerce' (1999) 22 *Dalhousie Law Journal* 190 which argues that 'electronic agents' lack the legal status and capacity to contract and that "it is difficult to conceive of any such transaction as achieving the fundamental traditional prerequisite to contract formation, viz the parties' formation of a meeting of the minds'. See also, Scholz, 'Algorithmic Contracts' (n 15) 141 which further defends her agency theory of AI contracting.

[83] See Scholz, 'Algorithmic Contracts' (n 15) 151.

[84] See *Lucy v Zehmer*, 84 SE 2d 516, 521 (Va, 1954) where the Court held that '[a]n agreement or mutual assent is of course essential to a valid contract but the law imputes to a person an intention corresponding to the reasonable meaning of his words and actions.' In some cases equitable doctrines impose a contractual obligation on a party that has not actually consented to a contract. For example, the equitable doctrine of promissory estoppel imposes a contractual obligation on a party that made a promise on which another reasonably and detrimentally relied. See, eg, *Restatement (Second) of Contracts* (n 28) § 19 which states that '[a] promise which the promisor should reasonably expect to induce action or forbearance on the part of the promisee or a third person and which does induce such action or forbearance is binding if injustice can be avoided only by enforcement of the promise.'; *Drennan v Star Paving Co*, 333 P 2d 757 (Cal, 1958) which holds that a general contractor could enforce an offer made by a sub-contractor on which the contractor had detrimentally relied when making a bid for a construction contract. Nevertheless, consent remains the foundational principle of contract law from which doctrines such as promissory estoppel mark an exception.

intention to enter into some kind of contract.[85] The party must manifest an intention to be bound by the specific terms of the contract in question. The agreement must be 'sufficiently definite to enable a court to give it an exact meaning.'[86] Thus, in order to create a binding contract both parties must objectively manifest consent to be bound by the sufficiently specific terms of that contract.

We can use an example to illustrate why contracts negotiated and completed by AI programs may not satisfy these basic requirements of an enforceable contract. Suppose that Annie runs an online antique store, 'Annie's Antiques.' To maximise her profits, she uses a sophisticated AI algorithm to set her prices and sell her antiques through a website. At a given moment, Annie does not know what antiques are being sold by her website, to whom, or for what price. Annie's activities are focused on sourcing antiques and maintaining the algorithm. The contracts negotiate and completed by Annie's algorithm purport to bind Annie. Yet, it seems that Annie does not manifest subjective or objective consent to be bound by the terms of these contracts.[87]

Suppose that Annie's website offers to sell a rare antique widget for one dollar, and that a customer eagerly uses the website to buy the widget at that price. Does Annie intend to sell the widget to that customer for one dollar? Clearly she does not. How could she, since she does not know what terms her website offers or whether they are accepted? There is thus no meeting of

[85] See *Wagner Excello Foods Inc v Fearn International Inc*, 601 NE 2d 956, 964 (Ill App Ct, 1992) where the Court held that '[t]o be enforceable, a contract must show a manifestation of agreement between the parties and be definite and certain in its terms. When material terms and conditions are not ascertainable there is no enforceable contract, even if the intent to contract is present.' See also *ANZ v Frost Holdings Pty Ltd Supreme Court of Victoria (Full Court) [1989] VR 695* where the Court held that mere agreements to agree are not binding.

[86] *Trimmer v Van Bomel* 434 NYS 2d 82, 88 (NY Sup Ct, 1980), quoting 1 Richard A Lord, *Williston on Contracts* (Lawyers Cooperative Publishing, 2000) § 37. See also *Council of the Upper Hunter County District v Australian Chilling & Freezing Co Ltd* (1968) 118 CLR 429, para 10, where the Court held that a contractual term referring to one party's 'costs' was not so vague as to be unenforceable because 'A contract to build a bridge at cost could not, in my opinion, be held void for uncertainty: it could not properly, in my opinion, be said to be meaningless: nor is it, in my opinion, ambiguous. Endless might be the arguments pro and con as to whether or not in marginal cases some item of expenditure is as claimed a cost, or as to how much of an expenditure is a cost, of the particular activity. But to my mind, generally speaking, the concept of a cost of doing something is certain in the sense that it provides a criterion by reference to which the rights of the parties may ultimately and logically be worked out, if not by the parties then by the courts.'

[87] Using AI programs to negotiate and complete contracts 'entails difficulties in satisfying the requirement the two parties should be in agreement, since in many cases one party will be unaware of the terms of the particular contract entered into by its electronic agent': Chopra and White (n 15) 30.

the minds between Annie and her customer. Annie lacks the specific intent—to enter into a contract to sell that widget at that price—required by a subjective theory of contracts.

The objective theory of contracts does not resolve this problem. According to the objective theory of contracts, it does not matter what Annie actually intends.[88] What matters is whether she performs acts 'which ordinarily accompany and represent a known intent,' such that a reasonable counterparty would believe that she intends to sell the widget for one dollar.[89] A reasonable customer (let us assume) would know that the website is operated by an algorithm and is not supervised in real time. No matter what words the website uses, a reasonable customer would not believe that Annie specifically intends to sell the widget to them for one dollar. For this reason, the words used by the website cannot constitute an objective manifestation of consent by Annie to the specific terms of the contract. Thus, it seems that Annie has manifested neither the subjective nor the objective intent necessary to enter into a legally binding agreement to exchange a rare antique widget for one dollar.

One might think: is Annie's website really different from the parking lot ticket machine in *Thornton*? After all, the operators of the ticket machine set up the machine to enter into contracts on their behalf. They do not know when the machine makes a contract or who it is contracting with. Yet, the contracts created by a ticket machine like the one in *Thornton* are legally binding. Here is the difference. In a case like *Thornton,* the people setting up the machine know what terms they are offering and they put a sign on the machine listing those terms. A reasonable customer, approaching the ticket machine, would believe that the people who set up the ticket machine intend to sell access to their parking lot on the terms listed on the machine. The ticket machine is merely the means by which the owners of the parking lot communicate with the customer. In this way the ticket machine is unlike Annie's algorithm, because the ticket

[88] See generally Joseph M Perillo, 'The Origins of the Objective Theory of Contract Formation and Interpretation' (2000) 69(1) *Fordham Law Review* 427. Perillo argues, contrary to popular understanding, that the objective theory of contract was not a 19th century invention but a longstanding feature of the common law.

[89] *Hotchkiss v National City Bank*, 200 F 287, 293 (SDNY, 1911) in which Learned Hand J famously asserts that '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.'; See also *Air Great Lakes Pty Ltd v KS Easter (Holdings) Pty Ltd* (1989) 2 NSWLR 309, where the Court adopts the objective theory of contracts by holding that the existence of a contract depends on the intentions conveyed by the words and actions of the parties rather than their actual subjective intent.

machine does not decide what terms to offer. It just communicates and records an offer and acceptance.

One might still object that the difference between Annie's AI algorithm and the ticket machine is just a matter of degrees. An algorithm is a set of rules that produce an output. A ticket machine might follow a simple algorithm: on weekdays sell tickets for $10, on weekends sell tickets for $8. Annie's website uses a much more sophisticated algorithm. Why should this difference in complexity be legally significant?

The difference matters because it is not merely a difference in complexity. It is a difference in the kind of input provided by the operators of the algorithm. If the algorithm is simple and deterministic, then the person programming the machine decides in advance the terms to be offered in each circumstance. The machine simply communicates that offer, or set of offers, to the customer. However, if the algorithm is programmed with general goals and parameters, as is the case with an autonomous AI program, then the program it is not merely communicating a set of offers.[90] Such an algorithm cannot merely communicate a set of offers because the information it receives from its programmer is not specific enough to constitute a legally binding offer or set of offers.

The problem can be illustrated graphically like this:

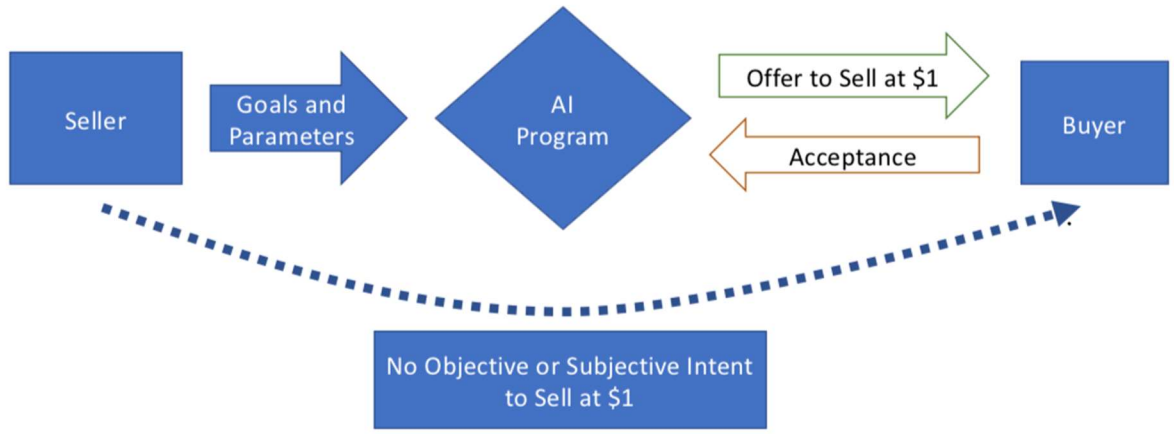

[90] As Scholz, 'Algorithmic Contracts' (n 15) 153 writes, 'if the instructions given to an algorithmic-agent by its principal are vague, they cannot be considered the level of objectively manifested intent necessary to ground a contractual promise.'

The offer to sell at one dollar is communicated by the AI program to the buyer. However, the AI program does not receive an instruction to that effect from the seller. The seller only provides vague goals and parameters. Thus, the AI program is not merely communicating an offer from the seller to the buyer; it is making the offer to sell at one dollar by itself. The seller, meanwhile, does not subjectively intend to sell at one dollar and does not manifest objective intent to sell at one dollar. Without subjective or objective intent to be bound by the specific terms of the agreement, there can be no legally enforceable contract between the seller and the buyer.

Can it really be the case that contracts negotiated by AI programs are unenforceable?[91]

Any judge would be tempted to avoid the massive disruption and uncertainty that refusing to enforce AI-negotiated contracts would create. Yet the manifestation of consent to be bound to the specific terms of the contract is not a merely formal requirement. It is the central principle of contract law.[92] So we seem to face a genuine puzzle: we want to enforce these contracts, but they seem not meet the most basic requirements for an enforceable contract.

B *The Agency Solution*

Several scholars have begun advancing a potentially elegant solution to this problem: treat AI programs as legal agents of the people who operate them.[93] This solution seems plausible. An AI program that negotiates and completes a contract on behalf of a person does the work that might

[91] Bayern (n 39) 155 suggests that 'it is hard to see a serious objection' to AI-negotiated contracts, and references the fact that these contracts are widespread. He writes: 'I don't believe they are problematic in practice, and there appear to be many examples where clearly enforceable legal agreements result from the ongoing operation of algorithms in ways that are today considered almost mundane.' I agree with Bayern that these contracts can be enforced within existing contract law theory. Yet I also agree with the agency theorist that there is a difficult problem here for contract law theory to solve. Bayern elides contracts that are themselves algorithms and contracts negotiated by algorithms. Even if both these uses of algorithms result in legally enforceable contracts, they pose very different issues from the perspective of contract law theory.

[92] See Randy E Barnett, 'A Consent Theory of Contract' (1986) 86 *Columbia Law Review* 269 which argues that, in spite of the objective theory of contracts and the legal protection of reliance, consent remains crucial to contract law because the best explanation of contract law is that it enforces consensual transfers of property rights and entitlements.

[93] See Chopra and White (n 15); Scholz, 'Algorithmic Contracts' (n 15); Ian R Kerr (n 82). See also Lehman-Wilzig (n 26) 451 arguing that agency 'is the most comprehensive and germane with regard to both the essence and function of such a creature'; Matthew U Scherer (n 26) considering whether AI programs could be granted, or have already been granted, legal personhood and arguing that 'agency law provides the most effective and flexible legal analogue for artificial systems.' Scherer discusses contract law to the extent of considering whether AI programs have the capacity to enter into contracts in their own right and whether AI programs can be given legal personhood by contract through an LLC. He explicitly sets aside the question of whether AI programs can be used negotiate and consent to legally binding contracts.

otherwise be done by a human employee. Annie might hire a salesperson and give her the same vague goals and parameters she gives to her AI program. Yet when the salesperson enters into a contract on Annie's behalf, that contract is clearly enforceable because the salesperson is Annie's agent.

Agency law allows one person, an agent, to act on behalf of another person, a principal. An agency relationship can be created in several ways. The principal can give an agent actual authority, by agreeing with the agent that the agent will act in some matter on the principal's behalf.[94] The principal can also give an agent apparent authority, if they lead others to reasonably believe that the agent has authority to act on their behalf.[95] A principal can also retroactively adopt the acts of another by ratification.[96]

These conditions might be satisfied by AI programs. In the simplest case, the people operating the AI program might be said to give the program apparent authority by inducing other people to reasonably believe that the AI program is authorised to act on their behalf.[97] We can see how this might be the case in Annie's antique shop. By creating her website Annie leads her customers to reasonably believe that the website controlled by her AI program is authorised to act on her behalf. An agency relationship does not need to be created intentionally.[98] Crucially, a

[94] American Law Institute, *Restatement (Third) of Agency* (2006) (*'Restatement of Agency'*) § 3.01 provides that '[a]ctual authority, as defined in § 2.01, is created by a principal's manifestation to an agent that, as reasonably understood by the agent, expresses the principal's assent that the agent take action on the principal's behalf.'

[95] Ibid § 2.03 cmt (c) provides that '[a]pparent authority holds a principal accountable for the results of third-party beliefs about an actor's authority to act as an agent when the belief is reasonable and is traceable to a manifestation of the principal. ... A third party's reasonable understanding of the principal's conduct will reflect general business custom as well as usage that is particular to the principal's industry and prior dealings between the parties.'

[96] Ibid § 4.01 cmt (b) provides that '[a]s the term is used in agency law, ratification is both an act and a set of effects. The act of ratification consists of an externally observable manifestation of assent to be bound by the prior act of another person ... [W]hen a person ratifies another's act, thelegal consequence is that the person's legal relations are affected as they would have been had the actor been an agent acting with actual authority at the time of the act.'

[97] This is not exactly what proponents of the agency solution suggest, although I think it is an improvement on their view. Scholz, 'Algorithmic Contracts' (n 15) 167 suggests that the operators of AI programs could create an agency by ratification. By ratifying a contract, a person becomes bound to its terms, as if it was negotiated by their agent. This way of applying agency law principles to AI contracts seems especially favourable to the operators of AI. It would allow them to ratify contracts that they like and ignore contracts that they did not like.

[98] *Restatement of Agency* (n 94) § 1.03 cmt (d) provides that '[i]f a principal voluntarily manifests assent or intention, the manifestation is effective although it is made negligently or is otherwise in error.The principal's manifestation of assent determines whether a relationship of agency arises, not the principal's unexpressed intention if otherwise unknown to the agent.'

person can create an agency relationship by their conduct, whether they intend to create this relationship or not.[99]

When an agent enters into a contract on behalf of a principal, that contract legally binds the principal.[100] If Annie's AI program is her agent, then it does not matter that Annie does not herself manifest consent to be bound to the terms of the contracts negotiated by her website.[101] As long as the AI program consents to the specific terms of the contract, that consent is sufficient to bind Annie.

The agency law solution can be represented graphically like this:

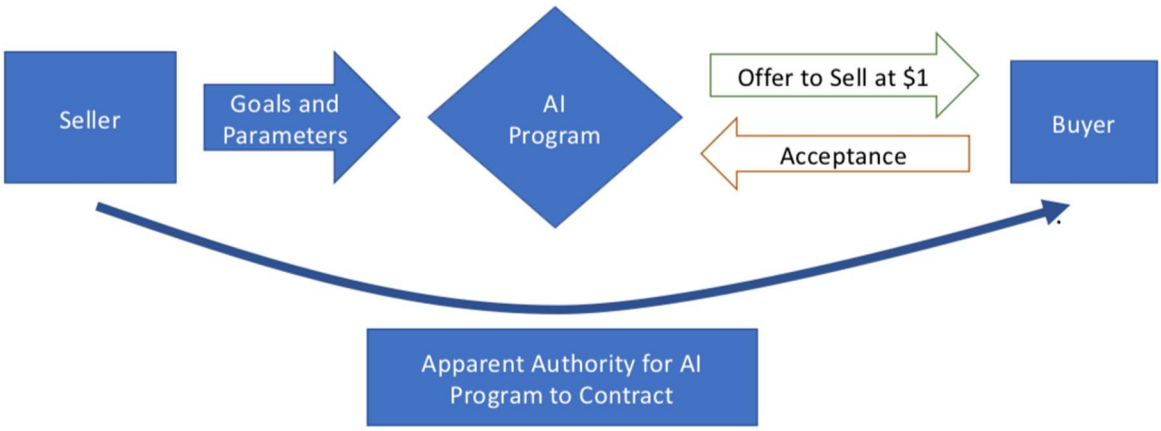

Here we see that the AI program completes the contract with the buyer through an exchange of offer and acceptance. Although the seller does not subjectively or objectively intend to sell for one dollar, the contract to sell for one dollar is binding on the seller because of the law of agency. The law of agency applies because, by setting up the AI program to contract on their behalf, the seller has given the AI program apparent authority to act as their agent.

[99] Ibid § 1.03 cmt (e) provides that '[a] manifestation does not occurin a vacuum, and the meaning that may reasonably be inferred from it will reflect the context in which the manifestation is made. Assent and intention may be expressed explicitly, but often they are inferred from surrounding facts and circumstances.'

[100] Technically this is only the case if the agent is acting on behalf of a disclosed principal. However, this is likely the case in AI-negotiated contracts since the counterparty will know whose AI program they are negotiating with. *Restatement of Agency* (n 94) § 6.01 cmt (b) provides that '[a]n agent has power to make contracts on behalf of the agent's principal when the agent acts with actual or apparent authority.'

[101] Ibid provides that '[a]n agent enters into a contract on behalf of the agent's principal by manifesting assent to an exchange.'

One might object that the agency solution just trades one problem of consent for another far more philosophically challenging problem of consent. If the AI program is acting as an agent and contracting on behalf of the seller, then our focus shifts from the intentions of the seller to the intentions of the AI program. In order for there to be an agreement between the AI program and the buyer, which could bind the seller by virtue of the law of agency, the AI program must subjectively or objectively intend to enter into such an agreement. Is it really possible for a computer program to *intend* to enter into a contract?

One might think that the objective theory of contracts provides an easy answer to this question. On the objective theory, the law has 'nothing to do with the personal, or individual, intent of the parties.'[102] So as long an AI program can give the appearance of intending to be bound, perhaps we do not need to answer the hard philosophical question of whether AI programs have intentions. This easy answer does not work. An AI program can only give people the reasonable belief that it intends to contract if it is reasonable to believe that an AI program can have intentions. Thus, for the agency solution to be plausible, the people who propose it must show that AI programs really can have intentions.[103]

Fortunately, there is good reason to suppose that AI programs can be said to have intentions, at least in some philosophically coherent sense.[104] There are a few ways this could work.[105] One theory of intentions is that they are simply an explanation of the behaviour of other people. This understanding of intentions is advocated by Daniel Dennett.[106] When we explain a

[102] *Hotchkiss v National City Bank* (n 89) 293.

[103] Perhaps it could be reasonable to believe that an AI program has intentions even if an AI program cannot actually have intentions, but this seems like a shaky foundation on which to build the enforceability of AI-negotiated contracts.

[104] When discussing artificial intelligence and intentions it is important to distinguish the concept of an intention in the philosophy of action, which is the relevant kind of intention for the law, from the concept of intentionality in the philosophy of mind. Intentionality refers to the capacity of mental states to represent things in world; for example, the capacity of a perception to represent an object. This is distinct from the more familiar idea of an intention in the philosophy of action, and in the law, according to which an intention is a plan or purpose which guides and controls an act. Cf Elizabeth Anscombe, *Intention* (Oxford University Press, 2nd ed, 1963); John Perry, 'Intentionality (2)' in Samuel Guttenplan (ed), *A Companion Volume to the Philosophy of Mind (John Wiley and Sons, 1996) 386.*

[105] It is sometimes suggested that we can assign intentions to AI in the same way we assign intentions to corporations. However, there are limits to how far the analogy between corporate intent and AI intent can be stretched. Corporate belief and intent are always in some sense reducible to the intent of the humans who work for and on behalf of the corporation. No analogous reduction is possible in the case of an AI program.

[106] Daniel C Dennett, *The International Stance* (MIT Press, 1989).

person's behaviour in terms of intentions we take what Dennett calls the 'intentional stance.'[107] We can take this intentional stance towards an AI program.[108] When a person contracts with Annie's website, it is natural to explain the website's behaviour by concluding that the website intends to enter into contracts. For Dennett, this means that the website can be said to have intentions. If we accept Dennett's theory of intentions—at least for the purposes of the contract law—then AI programs can be said to have the intention to enter into a legally binding contract.[109]

There is another problem that the proponents of the agency solution must overcome. AI programs, even if they can be said to have intentions, lack legal personhood. This is a problem for the agency view because agency is a relationship between legal persons. The Restatement (Third) of Agency explicitly rejects the possibility that a computer program could act as an agent, noting 'a computer program is not capable of acting as a principal or an agent as defined by the common law. At present computer programs are instrumentalities of the people who use them.'[110] This makes sense for simple computer programs, but it may not make sense for more complex AI programs.

The proponents of the agency solution take different approaches to the problem of legal personhood. Some argue that algorithms should be treated as constructive agents, allowing principles of agency law to govern AI contracts without settling the legal status of AI.[111] Others would go further and treat AI programs as legal persons.[112] Whatever approach they advocate,

---

[107] Ibid 17 states that '[t]his is how it works: first you decide to treat the object whose behavior is to be predicted as a rational agent; then you figure out what beliefs that agent ought to have, given its place in the world and its purpose.' See also, Chopra and White, (n 15) 13.

[108] This is proposed by Chopra and White (n 15) 11–19, where they state that:

> An artificial agent could and should be understood as an intentional agent, as acting for reasons that are the causes of its actions, if such an understanding leads to the best interpretation and prediction of its behavior. Such an interpretationist view is taken by the theory of the intentional stance, which renders coherent the ascription of mental predicates or propositional (intentional) attitudes to nonhuman entities.

[109] Notice that one of the common objections to Dennett's theory—that it matters to us that we and others have a phenomenological experience of intending—may not be relevant in the legal context. Afterall, the objective view of contracts specifically disavows any concern for the internal experience of intending to agree, preferring to focus on the outward manifestations from which others draw the conclusion that one intends to agree. In Dennett's terms, an objective manifestation of consent is speech (or other outward manifestation) explainable by the conclusion that the speaker thereby intends to enter a contract.

[110] Restatement of Agency (n 94) § 1.04.

[111] Scholz, 'Algorithms and Contract Law' (n 45) 165 states that 'I will set aside, for the moment, whether or not algorithms can be persons or agents in the traditional sense. Nothing in this analysis suggests that algorithms could or should be considered persons. Algorithms can be agents without legal personality, or quasi-agents for the purpose of understanding the legal obligations of their principles.'

[112] Chopra and White (n 15) 189 argue that an AI program with the right capabilities 'would have a strong case for legal personality… [t]here is no reason in principle that artificial agents could not attain such a status, given their current capacities and the arc of their current development in the direction of increasing sophistication.'

these scholars agree that we should use principles of agency law to explain how AI-negotiated contracts can be legally enforceable. In the next section, I argue that this is a mistake.

IV WHY ALGORITHMS ARE NOT AGENTS

In this section I argue that it is a mistake for judges and lawyers to treat AI contracting programs as legal agents, for two reasons. First, despite what some scholars claim, AI-negotiated contracts are enforceable under current contract law doctrines. In other words, there is no consent problem for the agency solution to solve. Second, there are strong conceptual and practical reasons for not applying the law of agency to the relationship between AI contracting programs and the people who operate them.

A *There is No Consent Problem for the Agency Solution to Solve*

There is no need to adopt radical solutions in order to make AI-negotiated contracts legally enforceable. We just need to change the way we think about the nature of the agreement between the person operating an AI contracting program and their counterparty. When a person sets up an AI program to contract on their behalf, they make an open offer to contract on the terms the AI program offers.[113] An offer need not be made in so many words. An offer can be implied from a person's conduct.[114] The conduct of setting up an AI program to contract on one's behalf clearly communicates an open offer to contract on whatever terms the AI program offers.

This open offer to contract on whatever terms the AI program offers is underdetermined, because its meaning depends on a future event: the decision of the AI program. We are familiar with other kinds of underdetermined offers.[115] For example, a person might offer to sell a widget

[113] Bellia (n 34) 1057–8, 1095 briefly considers whether contracts made by computers, whether or not acting autonomously, can be enforced on a similar open offer theory before concluding that current doctrine 'provides no clear answer' to the question of whether electronic contracts are enforceable. Bayern (n 39) 155 also contemplates an analogy between AI-negotiated contracts and contracts that requirements contracts that contain terms to be filled in by an algorithm, although he does not identify or discuss the distinctive issues that AI-negotiated contracts raise from the perspective of conventional contract law theory. He seems to suggest that algorithms can agree to contracts because algorithms can be contracts, but he does not discuss the issue of consent that arises in the former case but not the latter.

[114] *Formation in General*, Uniform Commercial Code (n 28) § 2-204 states that '[a] contract for sale of goods may be made in any manner sufficient to show agreement, including conduct by both parties which recognized the existence of such a contract'; *Restatement (Second) of Contracts* (n 28) § 19.

[115] See *Godecke v Kirwan* [1973] HCA 38 which holds that a contract is not void for vagueness if it specifies that some of its terms can be filled in at a later time by a third party). See also *Uniform Commercial Code* (n 28) § 2-204(3) which states that '[e]ven though one or more terms are left open a contract for sale does not fail for

for 'market price on Jan 1' or 'whatever terms Jack offers to Sally.' This kind of underdetermined offer introduces risk into contracting, but there can be good business reasons for taking this kind of risk.[116] In the case of AI contracts, the benefit to the seller of taking this risk is obvious: the seller gets the advantage of using an AI program to determine the specific terms of their contracts.

Deliberately underdetermined contracts have long been recognised as an efficient contracting innovation. The casebook example of underdetermined contracts are requirements and output contracts, which respectively entitle a buyer to buy as much of the seller's product as it needs, or require the buyer to buy as much of the seller's product as it produces.[117] Yet underdetermined terms can be built into a contract in other ways. A contract might deliberately use vague terms like 'reasonable best efforts' on the assumption that it is cheaper to have some other entity, such as an arbitrator, give greater specificity to the meaning of those terms if the need ever arises.[118] A recent innovation in this direction are so-called 'ex tempore' contracts, in which the contract delegates the job of specifying the meaning of vague contractual terms to a specially created entity, such as a dispute board.[119] That entity can then specify the meaning of contractual terms 'in real-time' as the need arises during the life of the contract.[120] On the open offer model, AI-negotiated contracts are just the latest and most extreme example of the creative use of underdetermination as a means of limiting transaction costs and achieving other efficiencies in contracting.

One of the great advantages of this model is that we do not need to consider what kind of thing a given AI program is or whether it can be said to have intentions. We can thus bypass

---

indefiniteness if the parties have intended to make a contract and there is a reasonably certain basis for giving an appropriate remedy.'

[116] Bayern (n 39) 155 notes that agreements can give legal effect to algorithms because algorithms can themselves be contracts or because the outputs of algorithms can be explicitly incorporated into contracts. I agree and go further by arguing that we can solve the problem of consent by assuming that an open offer to contract on the terms one's AI program offers is implied by the act of setting up an AI program to contract on one's behalf.

[117] See, eg, *Feld v Henry S Levy & Sons, Inc,* 37 NY 2d 466 (NY, 1975) which interprets the New York Uniform Commercial Code to impose an obligation of good faith on parties to output contracts.

[118] See *Coal Cliffs Collieries Pty Ltd v Sijehama Pty Ltd* (1991) NSWLR 1 which reasons that formal agreements to negotiate in good faith could be binding if they are intended to be binding and have terms that can be given a sufficiently definite meaning, such as a term giving power to an arbitrator to resolve disputes; *United Group Rail Services Ltd v Rail Corporation of NSW* [2009] NSWCA 177 which held that a clause requiring parties to negotiate in good faith was not void for vagueness.

[119] See Andrew Verstein, 'Ex Pro Tempore Contracting' (2014) *55*(5) *William and Mary Law Review* 1869, 1873 which describes this phenomenon as 'ex tempore contracting' and notes that 'this approach allows the parties to have clarity of obligation at or around performance, avoiding costly disputes while economizing on front-end costs.'

[120] Ibid 1872.

the thorniest legal questions raised by AI.[121] In theory any external event, such as market price on 1 January, could fill in the content of an underdetermined open offer. There does not even need to be a particular legal relationship between the person making the offer and the external event. This is an advantage of this view of AI contracting, because it allows courts to enforce AI-negotiated contracts without answering hard questions about the nature of AI programs, their intentions, and their legal relationship to the people who operate them.

If the buyer likes the terms offered by the AI program, then the buyer can accept the seller's offer to sell on the terms the AI program offers. The buyer can then communicate their acceptance to the seller by means of the AI program. For example, on Annie's website, a buyer can accept Annie's offer to contract on the terms her AI program offers by placing an order through the website. In this way the AI program serves two roles in the contracting process: (1) it provides an external event that gives more specific content to the seller's underdetermined open offer, and (2) it communicates the open offer from the seller and communicates an acceptance of that offer from the buyer. When filling the first role, the AI program acts as a content filling tool by providing content to the seller's open offer. When filling the second role, the AI program acts as a simple tool of communication, like the ticket machine in *Thornton*.

This model can be represented graphically like this:

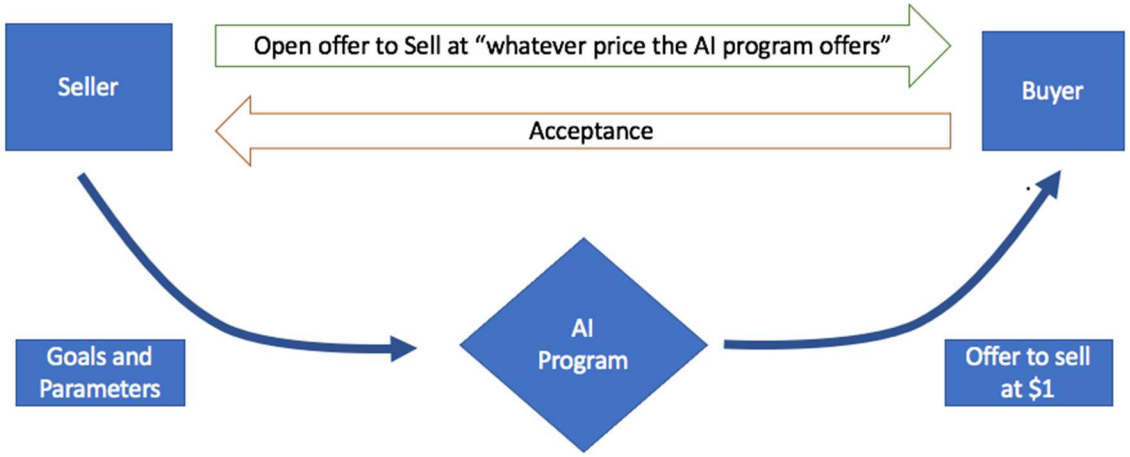

[121] For a discussion of this thorny question see Chopra and White (n 15); Lehman-Wilzig (n 26); Scherer (n 26). Bayern (n 39) 158–9 notes that these problems do not arise when an algorithm is a contract or when the output of an algorithm is incorporated into a contract.

Here we see that the seller makes an open offer to contract 'on whatever terms the AI program offers.' That offer is given precise content by the AI program. If the buyer likes the terms the AI program offers, then the buyer can accept the seller's offer to contract on those terms. Not pictured here is the fact that the AI program can also be used to communicate the offer and its acceptance.

It might be objected that there can be no contract between the seller and the buyer on this model because the offer and the acceptance do not match. The seller offers to sell on whatever terms their AI program offers, but the buyer agrees to buy at one dollar. If the offer and acceptance do not match, it might be argued, there can be no contract. Yet, while the offer and the acceptance may be expressed in different words, they are functionally the same. Once the AI program offers a particular buyer the price of one dollar, the meaning of the offer to sell at 'whatever price the AI program offers' becomes the same for that buyer as an offer to sell at one dollar. An offer and acceptance do not need to use the exact same words for there to be a contract, as long as it is clear that both parties agree to the same terms.

It might also be objected that an open offer to sell to each buyer at whatever price the AI program offers that particular buyer is unacceptably vague. Certainly, when the offer is made, its content is deliberately underdetermined. However, if a buyer accepts the terms the AI program offers, the terms of the resulting contract between the buyer and seller are perfectly clear. Recall that to be enforced a contract must be 'sufficiently definite to enable a court to give it an exact meaning.'[122] Courts will have no problem determining the exact meaning of the contract created when the buyer accepts particular terms offered by the AI program.

Perhaps this is still not enough to avoid the problem of vagueness. In our example, the AI program specifies the price at which to sell each widget. In reality, however, a highly autonomous AI contracting program might be free to make all kinds of complex trades, to buy and to sell or to make no contracts at all. It might be argued that an open offer to contract with each counterparty on whatever terms an AI program happens to offer that counterparty is simply too open-ended an offer to result in a binding contract.[123]

---

[122] *Trimmer v Van Bomel* (n 86) 88. See also *Uniform Commercial Code* (n 28) § 2-204 which states that 'a contract for sale does not fail for indefiniteness if ... there is a reasonably certain basis for giving an appropriate remedy.'

[123] Some advocates of the agency solution anticipate and reject an analogy to output and requirements contracts. Scholz, 'Algorithmic Contracts' (n 15) 155 seems to offer two arguments. First, she claims that unlike requirements

This argument is unpersuasive for two reasons. First, there is no doctrinal limit to the kind of open offers a party can make, as long as the resulting contract is precise enough to be enforced by a court. A basic principle of Anglo-American contract law is that individuals are free to invent new forms of contract to meet their needs. Unless there is reason to believe that the resulting contracts is vague, unconscionable, or contrary to public policy, there is no reason for courts to prevent parties from contracting in this way. Second, there is precedent for this kind of highly open-ended offer. A party might bindingly agree to match whatever price or terms a competitor offers. If a party can bind itself to what a competitor offers, it can bind itself to what its AI program offers.

It might be objected that a highly autonomous AI contracting program could learn from market data and evolve beyond the original expectations of the people who operate it. If an AI contracting program changes itself too much, it might be objected that the program's operators no longer consent to the way the program works.[124] This objection misunderstands the role of the AI contracting program in the open offer model I have proposed. On the open offer model, the AI contracting program simply fills in the content of the operator's open offer to contract on whatever terms the AI program agrees. In theory any exogenous event could fill in the terms of an open offer. For example, the market price on 1 January could fill in the terms of an open offer to sell at whatever the market price happens to be on 1 January. There is no need for the person making the open offer to understand or consent to the forces that control the exogenous event. In our example, there is no need for the person who makes the open offer to understand or consent to the forces that produce the market price on 1 January. In the same way, there is no need for the operator of an AI contracting program to understand or consent to the way that

contracts, in which some term of the contract depends on a later event, 'in the case of black box algorithmic contracts, *the choice to be bound itself is indeterminate*' (emphasis added). This is a misleading because (at least on my view) the 'choice to be bound' is actually made by the offeror when she decides to make an open offer to contract on terms decided by the AI program. This 'choice to be bound' is no more indeterminate than it is in any open offer. Second, Scholz, 'Algorithmic Contracts' (n 15) 155 claims that 'when what the algorithm will agree to cannot be determined at the time the company puts the algorithm into use, the company has not objectively manifested the intent to be bound at a sufficient level of specificity to form an enforceable contract.' This is not true. In *Trimmer v Van Bomel* (n 86) 88, as noted above, it was held that a contract is sufficiently specific to be enforced as long as the terms are 'sufficiently definite to enable a court to give it an exact meaning.' When the meaning of an offer or contract is made to depend on a subsequent exogenous event, such as whatever terms Jack offers to Sally, that offer or contract can be given a definite meaning at a later time even though its meaning was underdetermined at the time the offer was made.

[124] My thanks to an anonymous reviewer at the *ANU Journal of Law and Technology* for raising this possible objection.

program works. All the operator must do is manifest the intent to be bound by whatever terms their AI program agrees.

So far, we have been imagining that the operator of an AI program makes an offer that other people accept. This is what happens in the example of Annie's website. However, it is also possible for AI programs to accept offers made by others. It is even possible, as in *Quoine*, for there to be an AI program on each side of a transaction. In order to successfully explain why AI-negotiated contracts are enforceable, we must explain how AI programs can accept offers made by others. This is not a problem for my theory.

As Corbin explains, 'an offer is an act on the part of one person whereby he gives to another the legal power of creating the obligation called contract.'[125] In other words, the legal significance of an offer is that it gives to another, or to others, the power to create a legally binding contract with the offeror on certain terms. When a person makes an open offer to contract on whatever terms their AI program agrees, they give to others the power to make a legally binding contract with them on any terms the offeree agrees with their AI program. In the correspondence between the AI program and the counterparty, the AI program could make an offer that is accepted by the counterparty, or the counterparty could make an offer that the AI program accepts. Either way, when the counterparty reaches an agreement with the AI program, the counterparty thereby exercises the power, conveyed by the open offer, to create a legally binding contract with the person who set up the AI program on the terms the AI program has agreed.

What if two AI programs contract with each other? I have argued that a person who sets up an AI contracting program makes an open offer to be bound by whatever terms the program agrees. It might be objected that two open offers cannot combine to make a binding agreement. Consider a hypothetical case of duelling offers: Company A sends out a flyer offering to sell one widget to anyone for five dollars. At the same time, on the other end of town, Company B sends out a flyer offering to buy one widget from anybody for five dollars. Have A and B made a

[125] See Corbin, 'Offer and Acceptance' (n 30) 171. See also *Uniform Commercial Code* (n 28) § 2-206 which states that 'an offer to make a contract shall be construed as inviting acceptance in any manner and by any medium reasonable in the circumstances.'

contract? Obviously not, so how can there be an enforceable contract when two AI programs contract with each other?[126]

To answer this question, we have to think about what conduct can constitute an acceptance. When a person makes an offer, that person gets to decide how their offer can be accepted.[127] The offeror will specify how the offer can be accepted either tacitly or explicitly. When Company A and Company B send out their duelling offers each presumably specifies, either tacitly or explicitly, that to accept the offer one must contact the company directly. Sending out a matching open offer does not constitute an acceptance. That is why neither Company A nor Company B successfully accept the other's offer in the hypothetical.

Things are different in the case of two AI programs that reach an agreement with each other. A person who makes an open offer to contract on whatever terms their AI program agrees may explicitly or tacitly indicate that their offer can be accepted by another AI program. In essence, each person setting up an AI contracting program that can contract with other AI programs makes the following offer: 'I offer to contract with any counterparty on whatever terms my AI program agrees with that counterparty. This offer can be accepted either by (1) a legal person making an agreement with my AI program or (2) a legal person setting up an AI program to contract on their behalf which in turn makes an agreement with my AI program.' If two people make this kind of offer, and their two AI programs agree to specific terms, then the terms agreed by the two AI programs constitute a legally binding agreement between those two people.

For these reasons, I believe there is no doctrinal obstacle to AI-negotiated contracts. Perhaps the single greatest virtue of the Anglo-American law of contracts is its adaptability. Parties are able to shape agreements, and invent new forms of agreements, to meet ever-changing business opportunities. Unless there is a very good reason to the contrary, judges and contract law scholars should be willing to embrace new creative forms of contract. Open offers to contract on terms to be specified by the decision of an AI program may be new, but they can and should be enforced.

---

[126] See Bellia (n 34) 1058 who mentions this problem in the context of all electronic contracts, whether autonomous or not, and concludes that current doctrine gives no clear answer. Bellia's treatment of the issue fails to consider the significance of the offeror's power to set the terms of a valid acceptance.

[127] See Corbin, 'Offer and Acceptance' (n 30) 199 who states that '[t]he offeror is the creator of the power and at the time of its creation he has full control over both the fact of its existence and its terms.'

B *Why Algorithms are Not Like Agents*

We have now seen that there is no doctrinal obstacle to enforcing AI-negotiated contracts. The scholars who advocate using agency law as a way to make these contracts enforceable have a solution but no problem. Their solution also has some serious problems. As Anthony Bellia and Shawn Bayern have separately argued '[a]s a matter of legal doctrine, there are several problems with invoking principles of actual or apparent authority to argue that exchanges arranged by bots are enforceable.'[128] Bellia notes that it is difficult to see how an algorithm could consent to act as an agent, as required for actual authority. An appeal to apparent authority does not solve this problem, he argues, because apparent authority presumes the possibility of actual authority. Apparent authority exists when a person leads a counterparty to reasonably believe that some third party has authority to act on behalf of that person. Yet a counterparty cannot reasonably be led to believe that an AI program has authority, if it is not reasonable to believe that an AI program can have actual authority. Bayern notes further complexities. He points out that it could be legally unclear who the proper principal of the AI program is. For example, if a company makes contracts using an algorithm which is managed and updated by a third-party software company it might be unclear who 'controls' the algorithm for the purposes of agency law.[129] These are important concerns, although perhaps they could be addressed by Scholz's suggestion that the acts of AI programs could become binding on their principals by ratification rather than actual or apparent authority.[130]

Even if we concede that AI programs can be treated as legal agents, agency law still creates substantial and unnecessary complications. Treating AI programs as legal agents will sometimes require us to determine what an AI program intended or believed. Contract law requires this inquiry when applying doctrines such as unilateral mistake, mutual mistake, frustration of purpose, and dealing in bad faith. Even if determining the beliefs and intentions of AI programs is theoretically possible it will be a difficult, unpredictable, and expensive process for judges and juries.[131] It is a weakness of the agency view that it could require judges and juries to examine the inner workings of complex AI algorithms.

---

[128] Bellia (n 34) 1060.
[129] Bayern (n 39)160.
[130] Scholz, 'Algorithmic Contracts' (n 15) 167.
[131] This may be particularly difficult if we adopt Dennett's intentional stance. Dennett relies on our ability to determine a person's intentions from observing their behaviour. This will be difficult when the person is alleged

In addition to these practical problems, there are broader theoretical problem with describing AI contracting programs as legal agents. As Bellia points out, a relationship of trust between autonomous people is at the core of the concept of agency. Without that core, what does it mean to call an AI program an agent? '[I]f the bottom line is that a "principal" (user) is simply liable for the actions of the "agent" (bot), calling bots agents appears to add nothing more than a label.'[132] I share this concern, although it may be alleviated by conceding, as Scholz does, that AI programs could merely be considered 'constructive agents' for the limited purposes of contract law.[133]

Yet, even if AI programs are treated merely as constructive agents, the theoretical assumptions of agency law still matter. The core legal and moral assumption of the private law is that each person is always responsible for their own behaviour but not responsible for the behaviour of others.[134] We are each responsible for our own contracts and torts but not responsible for the contracts and torts of others. Agency law operates against this background, by making an exception to this foundational principle in a fairly narrow range of cases. The risk with treating AI programs as legal agents is that their operators will escape liability in cases that fall outside the range of cases in which agency law assigns liability to the principal.

These dangers become more obvious if we move from contract to tort. Tort law holds each of us responsible for harm that we cause by failing in our duties to others. Agency law, and the concomitant tort doctrine of *respondeat superior*, make an exception to this background assumption.[135] Whether a person is liable for the torts of their agent depends on the degree of control they exercise over the agent's behaviour.[136] The more autonomous the agent, the less

to have made a mistake because the mistake may cause them to behave in a way that is inconsistent with their intent.

[132] Bellia (n 34) 1065.

[133] Scholz, 'Algorithmic Contracts' (n 15) 165.

[134] See Arthur Ripstein, *Private Wrongs* (Harvard University Press, 2016) for an excellent development of this classic idea in the context of tort law. Ripstein derives the basic rules of tort law from the principle that no one should be in charge of anyone else.

[135] See *Ira S Bushey and Sons Inc v United States*, 398 F 2d 167 (2d Cir, 1968) which found an employer liable for the drunken acts of one of its employees; *Restatement of Agency* (n 94) § 7.03 cmt (b) which states that '[a]n agent's tort may, separately, subject a principal to vicarious liability ... a principal is subject to liability when an agent who is an employee ... commits a tort while acting within the scope of employment.'

[136] *Restatement of Agency* (n 94) § 7.07(3)(a) states that 'an employee is an agent whose principal controls or has the right to control the manner and means of the agent's performance of work.'

likely it is that the principle will be held liable.[137] This principle seems fair because of our background assumption that people are not responsible for the behaviour of others.[138] Yet, this principle of agency law would create dangerous perverse incentives if artificial intelligence programs are treated as legal agents. It would mean that the operators of AI programs could avoid liability by failing to control the AI programs they operate.

Some scholars have suggested solving this problem by making AI programs legal persons and even giving them their own assets from which to pay damages.[139] This solution would only make things worse. Even if an AI program were given its own resources it would not necessarily care about those resources or be incentivised to avoid harming others the way a corporation or human would be. An AI program could be programmed to care about its own resources, but this would create conflicts of interest whenever the program worked on behalf of someone else. If AI programs were programmed in this way, then we would have artificially created the very conflicts of interest that agency law tries to solve. It is hard to imagine that this would be effective or efficient as a way of regulating AI.

Of course, we could use agency law to regulate AI contracts without using agency law to regulate AI torts.[140] We could also, as Scholz suggests, treat AI programs merely as constructive agents. I am not claiming that it is incoherent or immediately disastrous to use agency law in this way. My point is that it is dangerous for scholars and judges to start thinking about artificial intelligence programs as legal agents. Corporations who create and benefit from AI programs have strong reasons to disclaim responsibility for harms caused by those programs. Conceptualising AI programs as legal agents makes it easier to assume that the programs

[137] See *Restatement of Agency* (n 94) § 7.07 cmt (f) which states that '[n]umerous factual indicia are relevant to whether an agent is an employee. These include: the extent of control that the agent and the principal have agreed the principal may exercise over details of the work ... Also relevant is the extent of control that the principal has exercised in practice over the details of the agent's work.' See also American Law Institute, *Restatement (Third) of Torts: Liability for Physical and Emotional Harm* (2010) § 57 ('*Restatement of Torts*').

[138] *Restatement of Torts* (n 137) § 57 cmt (c) states that '[s]everal rationales support the general rule that those who hire independent contractors are not subject to vicarious liability for harms caused by thenegligence of the contractor. In employer-employee settings, vicarious liability has been seen as consistent with fairness because the rule places on the employer's enterprise the costs of harms that result from activities characteristic of, and of benefit to, the enterprise. By contrast, in hirer-independent contractor settings, the independent contractor is the person or entity that regularly benefits from the risk-creating enterprise.' These rationales are entirely beside the point in the context of harm done by operating an autonomous AI program because autonomous AI programs do not have their own legal interests.

[139] See Chopra and White (n 15) 153–89.

[140] This seems to be suggested by Chopra and White, (n 15) 71–98.

themselves are primarily responsible for the harm they cause and that those who create and benefit from those programs are only vicariously and occasionally responsible. This is not just a 'slippery slope' argument. It is an argument premised on the idea that legal reasoning matters. If courts decide contracts cases by reasoning that AI programs are akin to legal agents, that reasoning will be significant. That reasoning will lead other courts to conclude that AI programs must be akin to legal agents in other contexts.

Given these problems, it is interesting that the agency law solution has become popular so quickly. When we are confronted with something new, such as artificial intelligence, we try to fit it into our existing legal structures. The common law has always done this by analogy. The common law principle of treating like cases alike allows the law to be highly adaptable to new cases. We just need to find the right analogy.

Are AI programs that negotiate contracts more analogous to employees or to mere tools? Of course, AI programs are not exactly like either, but they *seem* more like employees because the most salient feature of AI programs are the features they share with employees. The thing that makes AI new and exciting is its ability to replace the work of human intelligence. Yet this is the wrong analogy for the purposes of contract law. AI programs may be like employees in their ability to reason and problem solve, but they are unlike employees in terms of their legal status and legal relationship to the people who use and benefit from them. Unlike employees, AI programs do not have their own goals or projects and are not legal persons. In this respect they are more akin to mere tools. The analogy to mere tools may initially seem unattractive, because it picks out features of AI programs that are not particularly exciting. Ultimately, however, I believe it offers a better way of thinking about the legal relationship between AI programs and the people who operate them.

V THE OPEN OFFER THEORY AND UNILATERAL MISTAKES

The problems with the agency view return us to where we started: with *Quoine* and the application of the doctrine of unilateral mistake to AI-negotiated contracts. On the surface, *Quoine* neatly illustrates the two problems I have described with the agency view. First, applying the agency view in cases like *Quoine* would require the court to make a difficult inquiry into the internal workings an AI program to determine why it made a mistake. The court would have to make this inquiry because only some types of mistake (eg typos) can be a ground for claiming unilateral mistake while other types of mistakes (eg a failure to investigate the real value of a product) cannot.[141] Second, applying the agency view would allow a company benefiting from an AI contracting program to avoid liability for one of the risks involved with such programs: the risk that the program will make a mistake. For these reasons, it might seem that the decision in *Quoine*, to treat the AI programs as mere machines rather than legal agents, can serve as a good example to courts in Australia and the United States.

Unfortunately, things are not quite that simple. A closer look shows that the reasoning in *Quoine* risks giving an unjustified advantage to AI parties over human parties. Recall the court's reasoning that there could be no reversal for unilateral mistake because the non-mistaken party (B2C2) was not aware of the mistake at the time of contracting.[142] The people at B2C2 were not aware of the mistake because they, too, used an AI program to negotiate their contracts. Their AI program might plausibly have been aware, but that awareness is legally irrelevant on the theory that the program is a mere tool, not a legal agent. The problem with this reasoning is that it would allow a party to invoke the defence of unilateral mistake against a human counterparty (who might have been aware of the mistake) but not to invoke the defence against a party using an AI program. American courts can avoid this problem because their law generally does not require actual knowledge of the mistake by the non-mistaken party for reversal on the grounds of unilateral mistake.[143] However, to see how these courts should apply the doctrine of unilateral mistake to AI contracts, we must take a small step back and consider the basic principles underlying the doctrine of unilateral mistake.

---

[141] See nn 146, 148.

[142] See *Quoine* (n 1) 46–57.

[143] *Restatement (Second) of Contracts* (n 28) § 153(b) which states that 'the contract is voidable [if inter alia] the other party had reason to know of the mistake or his fault caused the mistake.' For simplicity, and due to my ignorance of Australian equitable doctrines, this section refers only to the law of unilateral mistake in the United States.

The doctrine of unilateral mistake is notoriously unclear. One textbook describes it as 'confused beyond reconciliation.'[144] However, it is possible to bring some unity to the cases by recognising that the central issue is consent: some kinds of unilateral mistake undermine the mistaken party's subjective or objective consent to the terms of a contract while other mistakes do not.[145] If the nature of the mistake is a mismatch between a person's intention in entering a contract and what it would be good or prudent for them to do, then there is no basis for a claim of unilateral mistake. Legally, a person can genuinely consent to an agreement even though it is foolish or not what the person would prefer after further reflection. For example, if I sell what I believe to be a cheap knick-knack for $50, not knowing that it is really a valuable antique worth $50,000, that mistake is not a basis for voiding the contract.[146]

However, if the nature of the mistake is a mismatch between the person's intention in entering the contract and what the contract purports to do, then there is a potential basis for voiding the contract due to a lack of consent.[147] For example, suppose I know the antique to be worth $50,000 but due to a typo, I email an offer purporting to sell it for $50. A resulting contract is voidable because, despite typing $50, I did not actually intend to sell the object for $50.[148] To consent to a contract I must manifest an intent to be bound to certain terms. If it is clear from

[144] Addison Mueller and Arthur Rosett, *Contract Law and Its Application* (Foundation Press, 1977) 474.

[145] The account of unilateral mistake I outline here is my understanding of the traditional consent based theory of unilateral mistake. See Andrew Kull, 'Unilateral Mistake: The Baseball Card Case' (1992) 70(1) *Washington University Law Review* 57 which describes the traditional consent-based theory of unilateral mistake. The alternative approach to the doctrine of unilateral mistake, popular in the law and economics literature, sees the problem in a case of unilateral mistake not as an absence of consent but as a failure to disclose information that ought to be disclosed. See, eg, Eric Rasmusen and Ian Ayres, 'Mutual and Unilateral Mistake in Contract Law' (1993) 22(2) *Journal of Legal Studies* 309; Anthony T Kronman, 'Mistake, Disclosure, Information, and the Law of Contracts' (1978) 7(1) *Journal of Legal Studies* 1. Since the distinctive issue raised by AI-negotiated contracts is the presence or absence of consent, rather than the disclosure or non-disclosure of information, I have focused here on consent based understanding of the doctrine of unilateral mistake. Whether the literature on disclosure ought to give special consideration to AI contractors can be left for another paper.

[146] See *Laidlaw v Oregon*, 15 US (2 Wheat) 178 (1817), where it was held that a buyer's superior information about the likely market value of tobacco does not make the purchase of tobacco at a lower price voidable; *Fox v Mackreth* 29 Eng. Rep. 224, 234 (1788), where it was held that a buyer's superior information about the value of a mine under some farm land does not make the purchase of the land at a lower price voidable; Arthur Linton Corbin, *Corbin on Contracts* (West Publishing Company, 2nd ed, 1952). See also, Kull (n 145) 63.

[147] What matters here is whether there is a mismatch with what the contract purports to do *on its face* rather than a conflict with all the potential consequences or implications of the contract. A mistake due to a mismatch between a person's intention in entering into a contract and all the consequences or implications of that contract is not a basis for reversal due to unilateral mistake.

[148] See *United States v Braunstein* 75 F Supp 137 (SDNY, 1947) where it was held that a buyer cannot enforce an offer of 10 cents per 25 pound box when it knows the seller meant 10 cents per pound; *Taylor v Johnson* [1983] HCA 5 where it was held that a buyer cannot enforce a contract to sell 10 acres of land for $15,000 when it know the seller meant to sell for $15,000 per acre. See also Corbin, *Corbin on Contracts* (n 146) 610; Kull (n 145) 52.

the context that I made a typo, and do not really intend to sell for $50, then my email does not, in that context, manifest an intention to sell for $50.

This is true on both the subjective and objective theory of contracts. On the subjective theory, if I do not actually intend to sell for $50 then I do not consent. On the objective theory, if a reasonable buyer would know I had made a typo, and do not really intend to sell for $50, then I have not manifested objective consent to sell at $50. Without an objective or subjective manifestation of consent, there can be no contract. That is why the buyer cannot 'snap-up' my typo offer to sell for $50.

It would be impractical to void every contract in which one party agrees to something they did not intend to agree to. Society, and especially counterparties, have an interest in the predictability and reliability of contracts.[149] A person entering into a contract is responsible for ensuring that the contract says what they want it to and assumes the risk that it does not. This is why the doctrine of unilateral mistake allows a contract to be voided only if (1) the nature of the mistake is such that it undermines consent, and (2) a reasonable counterparty would have been aware of the fact and nature of the mistake.[150] If a reasonable counterparty would have been aware that there was mistake that undermines consent, then the predictability and reliance interests are weak and the counterparty is partially responsible for the fact that the mistake was not corrected.

We can now see how the doctrine of unilateral mistake ought to apply to AI-negotiated contracts. Recall my proposal that a person who sets up an AI program to contract on their behalf makes an open offer to contract on whatever terms their AI program agrees. Now suppose that the AI program makes a mistake and agrees to a bad deal. The person has clearly made a mistake—they should have done a better job programming their AI program—but that is not the kind of mistake that makes a contract voidable. It is a mismatch between what they did and what it would be prudent to do. There is no lack of consent, because the person clearly intended to enter into contracts on whatever terms their AI program agreed, and the resulting contract was on the terms their AI program agreed. If I offer to sell a knick-knack, without knowing what it is

[149] See Kull (n 145) 67, which states that 'the ordinary functioning of society requires that one who makes an offer or a promise be liable to be taken at his word.'

[150] See *Restatement (Second) of Contracts* (n 28) § 153(b), The contract might also be voidable, if there was the right kind of mistake and the non-mistaken party has not acted in reliance. See also Kull (n 145) 63. Moreover, the *Restatement (Second) of Contracts* (n 28) § 153(a) also allows the contract to be voidable if the mistake renders it unconscionable. For a detailed discussion of other limits that have been put on the doctrine of unilateral mistake in common law jurisdictions see *Taylor v Johnson* (n 148).

worth, I cannot void the sale for unilateral mistake just because the knick-knack is worth more than I expected. In the same way, if I make an open offer to contract on 'whatever terms Jack offers to Sally', without knowing what Jack will offer, I cannot void the contract just because the terms Jack offers to Sally are worse for me than I expected.

This simple application of the doctrine of unilateral mistake to AI-negotiated contracts solves the two problems with the agency view's approach. First, while the agency view requires a difficult inquiry into the workings of the AI program, the open offer approach provides a simple default rule: AI-negotiated contracts cannot be voided due to a unilateral mistake by the AI program. Second, while the agency view allows the operators of AI programs to avoid liability, this default rule is fair because it allocates responsibility for mistakes by the AI program to the party that is best able to prevent those mistakes and assess how likely they are to occur.[151] Despite being simple, this default rule is not too draconian. Operators of an AI program can, and do, contract around this default rule: for example, through disclaimers stating that contracts made with the AI program are subject to review.[152]

I do not mean to make any comment on Singapore law: the Court of Appeal knows best. However, as a matter of US law, the *Quoine* court reached the right result for the wrong reason. The right result was to deny a defence of unilateral mistake where the mistake was made by an AI contracting program. The right reason is that a mistake by an AI contracting program is not the kind of mistake that can ground a defence of unilateral mistake, because it does not undermine the consent of the legal person who offered to contract on whatever terms their AI program agreed.

---

[151] See Richard Posner, *Economic Analysis of Law* (Little, Brown and Company, 3rd ed, 1986) 90, which stresses the importance of putting liability for a contracting mistake on the party that could avoid it at least cost. See also Kronman (n 145), noting that the mistaken party is usually the least cost avoider.

[152] For example, the terms and conditions on Amazon.com include a provision purporting to protect Amazon in the case of a mistaken price. It reads 'With respect to items sold by Amazon, we cannot confirm the price of an item until you order. Despite our best efforts, a small number of the items in our catalog may be mispriced. If the correct price of an item sold by Amazon is higher than our stated price, we will, at our discretion, either contact you for instructions before shipping or cancel your order and notify you of such cancellation': 'Conditions of Use', *Amazon* (Web Page, 22 January 2021) <https://www.amazon.com/gp/help/customer/display.html?nodeId=GLSBYFE9MGKKQXXM>.

VI CONCLUSION

As *Quoine* illustrates, contracts negotiated by artificial intelligence present a new challenge for courts tasked with applying contract law doctrines. Contract law scholars can help these courts by providing a model to explain why AI-negotiated contracts are enforceable. Courts need such a model in order to apply doctrines like unilateral mistake. Some contract law scholars have argued that AI-negotiated contracts are unenforceable under existing contract law doctrines and that courts should apply agency law to enforce these contracts. This paper offers an alternative. When American and Australian courts are confronted with challenges to AI-negotiated contracts, they can enforce those contracts under the classic doctrines of contract law, by adopting the open offer interpretation of AI contracting.

The agency solution is not needed, I have argued, because AI-negotiated contracts are perfectly enforceable within existing contract law doctrines. AI-negotiated contracts are enforceable because the people operating AI contracting programs make an open offer to contract on whatever term their program agrees. This open offer model allows courts to adapt to the reality of AI contracting, and apply doctrines like unilateral mistake, without deciding the legal status of AI programs or altering the basic legal relationship between an AI program and the person who operates it.

ACKNOWLEDGMENTS

My thanks to Amy Chua, Caroline Lawrence, and two anonymous reviewers at the ANU Journal of Law and Technology.