AI-infused contracting and the problem of relationality: Is trustworthy AI possible?*

TT Arvind†

I Introduction: The rise of AI-infused contracting

The focus of this chapter is on a relatively new, but rapidly growing, phenomenon which I term 'AI-infused contracting'.[1] AI has increasingly come to be used in contracting, both in the contracts themselves and in the broader transactional processes within which contracts are embedded. These uses can be broadly classified into four archetypes: making transactional decisions, creating self-enforcing contractual mechanisms, managing the contractual lifecycle, and producing contractual terms. It is these processes, used individually or in combination, that I refer to as 'AI-infused contracting.'

AI-infused contracting holds considerable promise for the overall quality and effectiveness of transactions in the commercial and non-commercial worlds. Nevertheless, this chapter argues that its rise is not problem-free. Making the AI systems that underpin it resilient and trustworthy involves non-trivial challenges, which go to the heart of the nature of contract law and contract practice in the present day. Contract law occupies a distinctive place within the broad domain of private law, differing from other branches of private law in ways that are of considerable importance to AI-infused con-

*NB: Given that this is a conference draft, the ideas in the paper may very well change between now and final publication. Please contact me for an updated copy if you want to cite it, and especially if you intend to disagree with it. Any comments will be very gratefully received.

†York Law School, University of York. email:t.t.arvind@york.ac.uk. The work underlying this paper was partially funded by UK Research and Innovation Project EP/V026747/1, 'Resilient Autonomous Socio-cyber agents (REASON)', which is part of UKRI's Trustworthy Autonomous Systems programme. I am grateful to UKRI for its support.

[1] I use the term 'AI' in this chapter to cover not just artificial intelligence in the strict sense, but also the broader category of algorithmic and autonomous systems, including deterministic systems as well as those based on machine learning and stochastic processes. Most of my argument is in fact concerned with algorithmic systems that function at least partially autonomously, and that use a mix of deterministic and stochastic processes, but are not 'intelligent' in a technical sense. I use 'AI' rather than the more accurate 'AS' purely to reflect popular usage.

tracting. Most branches of private law—including tort, equity, and property law—are concerned directly with interests and interpersonal relations. Contract law, in contrast, is not concerned directly with either. It is not that interests and relations do not matter to contract law. It is, rather, that the focus of contract law is primarily on the *medium* through which these interests and relations are created—that is, the parties' contract—and only orthogonally on the interests and relations themselves. As far as contract law is concerned, the medium is very much the message.[2]

This focus on the medium rather than the substance is a distinctive feature of the general law of contract. While bodies of sectoral regulation such as consumer and even construction law can and do focus on the substance of a transactional relationship,[3] the general law of contract does not. This tendency is the result of deeply embedded features of law,[4] and it means that many of the more significant challenges that AI poses for other areas of private law raise at best trivial issues in the law of contract. The more important intersection between contract law and AI lies, instead, in a set of challenges arising out of contract law's focus on the medium rather than the substance of transactions. The transactional heart of contracts lies in the potential they offer for joint maximisation.[5] However, as a growing body of research not just in law[6] but also in management[7] has shown, the actual practice of contracting deploys and depends on techniques of drafting and management that not only fail to further joint maximisation but, frequently, achieve its precise opposite by facilitating opportunistic and extractive behaviour.[8]

These points matter for AI because AI in the real world is not simply a technical system, but a socio-technical system. Its real world use depends not just on the technology underpinning it, but necessarily also on the manner in which and ends towards which that technology is deployed, and on the

[2] The idea of a medium being the message comes from the work of Marshall McLuhan. See M McLuhan, *Understanding Media: The Extensions of Man* (MIT Press 1994) 7-21. McLuhan is discussed in more detail in Section IV.

[3] See e.g. the discussion of construction law in C Ellis, 'Regulating Commercial Contracts: What can we Learn from Part II of the Housing Grants, Construction and Regeneration Act 1996?' in TT Arvind and J Steele (eds), *Contract Law and the Legislature: Autonomy, Expectations, and the Making of Legal Doctrine* (Hart Publishing 2020).

[4] For an in-depth discussion, see D Mac Síthigh, *Medium Law* (Routledge 2018).

[5] For a recent overview of the theoretical literature on joint maximization, see RE Scott, 'A joint maximization theory of contract and regulation' in H Dagan and BC Zipursky (eds), *Research Handbook on Private Law Theory* (Edward Elgar 2020).

[6] In particular, the classic work of Ian Macneil and Stewart Macaulay. See D Campbell (ed), *The Relational Theory of Contract: Selected Works of Ian Macneil* (Sweet and Maxwell 2001); D Campbell (ed), *Stewart Macaulay: Selected Works* (Springer 2020).

[7] For a review of the management literature, see DJ Schepker and others, 'The Many Futures of Contracts: Moving Beyond Structure and Safeguarding to Coordination and Adaptation' (2014) 40(1) Journal of Management 193.

[8] F Lumineau, 'How Contracts Influence Trust and Distrust' (2017) 43(5) Journal of Management 1553.

predispositions, asymmetries, and biases that characterise the specific social contexts in which those ends are pursued. The argument of this chapter is that this makes AI something of a double-edged sword in the domain of contracts. In each of the areas in which it is used, AI has the potential to transform the drafting, management, and implementation of contracts in a way that ameliorates a significant proportion of the problems posed by existing contracting practices. But, equally, it has the potential to materially exacerbate these problems, in part because of the nature of AI itself and more specifically of the heuristics that underpin it, and in part because of the limits of law's regulatory capacity. Addressing this risk, and creating AI that is trustworthy in relation to contracting, will require channeling its use in a new direction, which this chapter describes as 'transactional responsibility'. Achieving and implementing this standard will, in turn, require rethinking not just contract law but also general programming practices and approaches in the field of contract-related AI.

The purpose of this chapter is to begin that process of rethinking, by outlining and analysing the nature of AI's dual potential in relation to contracting, the roots of that duality in the structure of AI, and the challenge it presents for law as well as for the computational design of AI. Part II introduces the argument by discussing the use of AI in the domain of contracting, and the manner in which its processes differ from those used by human transactors. I argue that for the use of AI to be sustainable, it must meet two baseline conditions: the condition of resilience and the condition of trustworthiness. The nature of the socio-technical process underpinning the use of AI, however, means that AI is not naturally given to meeting this conditions, and achieving them will therefore require a high degree of attentiveness to the techniques used to design AI as well as the social context in which AI is used.

Part III turns to this social context, examining the issues with current contracting practice which get in the way of contract law serving as an effective vehicle of joint maximisation. I argue that contracts can have a range of effects on a transaction, not all of which are beneficial, and that the deleterious effects are a particularly strong risk in asymmetric transactions. Part IV builds on this, by demonstrating how and why these issues with contracting practice leads to AI having a dual potential to either ameliorate or exacerbate the problems of contracting.

Part V considers how these challenges can be resolved in a manner that contributes to the resilience and trustworthiness of AI. I outline a set of principles, which I term the principle of transactional responsibility, which I argue has the potential to make a strong contribution to addressing this issue. I conclude by exploring its practical implications for the operation of AI, and examines how the law, too, will need to evolve in order to create conditions in which AI-infused contracts serve ends that are responsive and responsible, rather than extractive and oppressive.

II Contracts and AI: Mapping the terrain

As things stand, there are four uses to which AI is put that are of relevance to the law and practice of contracting. The first, and probably the best known, is the use of AI in making transactional decisions, for example decisions on whether to transact and if so at what price. Examples range from high-frequency trading, in which market actors use algorithms based on complex and typically proprietary computational models to identify and execute large volumes of security trades automatically and rapidly, to dynamic pricing in which algorithms automatically adjust the price of goods or services to take into account market factors such as demand. Somewhat more controversially, it also includes discriminatory pricing, in which enterprises use complex data-derived profiles of customers to make individualised pricing decisions with the result that identical products and services are offered at different prices to different customers based on the inferred characteristics of that customer.[9] Although this particular use of AI is usually considered in the context of pre-contractual decision-making, it is wholly capable of being used in relation to in-contractual decision-making, such as in relation to whether a contractual discretion should be exercised and if so in what manner. Examples of this type of use already exist, particularly in relation to the exercise of discretion under different types of financial lending contracts.[10]

A second use of AI is in the emergence of smart contracts. Smart contracts are, in essence, computerised transaction protocols that, when initiated, are capable of automatically executing some or all of the terms of a contract.[11] In the modern context, smart contracts typically take the form of a script in a programming language that is stored on a blockchain, and that is triggered when a transaction is addressed to it. The content of the script is the functional equivalent of the terms of a traditional contract, and determine precisely what transpires when the transaction is executed. The classic application of a smart contract is the trading of cryptocurrency, but they are also capable of being used in other types of commercial transactions which are capable of being executed electronically. They have, for example, been used to create an automated system for buying and selling excess energy generated by domestic solar panels.[12]

[9] For a critical analysis of this trend in the context of anti-discrimination law, see TB Gillis and JL Spiess, 'Big Data and Discrimination' (2019) 86 University of Chicago Law Review 459.

[10] See e.g. M Hurley and J Adebayo, 'Credit Scoring in the Era of Big Data' (2016) 18 Yale Journal of Law and Technology 148.

[11] This was the definition used by Nick Szabo in the paper that is generally taken to have developed the idea of a smart contract. See N Szabo, 'Smart Contracts' (1994), available at https://www.fon.hum.uva.nl/rob/Courses/InformationInSpeech/CDROM/Literature/LOTwinterschool2006/szabo.best.vwh.net/smart.contracts.html.

[12] K Christidis and M Devetsikiotis, 'Blockchains and Smart Contracts for the Internet

A third, and slightly less-well known use—at least in legal scholarship—is in the area of what has come to be called contract lifecycle management. The complexity of the obligations involved in large-scale transactions creates a need for managerial capacity, which many organisations lack.[13] The idea of a 'contract lifecycle' was devised by management theorists as a way of conceptualising the manner in which an organisation's focus shifts to different aspects of the transaction over its lifecycle[14] and, thus, for better structuring and directing managerial resources within organisations. There is now a significant amount of software that assists organisations with contract lifecycle management, and ongoing work to use machine learning and natural language processing to infuse these tools with a greater degree of artificial intelligence and autonomy. A particularly promising line of work in relation to commercial contracts lies in combining contract lifecycle management with blockchains, in the form of both distributed ledgers and smart contracts. Examples of the former already exist in sectors such as international trade and transport, which involve large numbers of relatively standard and well-understood contracts, documents, risk mitigation measures, and contractual processes, and therefore lend themselves readily to automation.[15] Similarly, attempts have been made to specify and develop smart contract trees that are capable of handling the entire transactional lifecycle in the context of transactions taking place between members of a decentralised manufacturing network.[16]

A fourth and final potential use of AI lies in the field of the actual drafting of contracts. Contract drafting software already exists, but the ones in common use remain slightly more sophisticated versions of the document assembly engines that emerged in the 1990s. As such, they continue to be based on a deterministic processes, and use questionnaires to trigger logic trees which lead to particular clauses being selected from a bank of predrafted clauses.[17] Although no non-deterministic systems exist in the wild as yet, it is generally accepted that 'predictive contracting' approaches that

of Things' (2016) 4 IEEE Access 2292, 2298.

[13] See e.g. TL Brown and M Potoski, 'Contract-Management Capacity in Municipal and County Governments' (2003) 63(2) Public Administration Review 153, analysing the limits on the capacity of local government bodies to manage outsourcing contracts.

[14] In one frequently used scheme, these are described as involving, successively, transactional architecture, negotiation, operation, and regeneration. See S Cullen, *The Contract Scorecard: Successful Outsourcing by Design* (Routledge 2016) 96-97.

[15] F Munari, 'Blockchain and smart contracts in shipping and transport' in B Soyer and A Tettenborn (eds), *New Technologies, Artificial Intelligence and Shipping Law in the 21st Century* (Informa Law 2020) 9-11.

[16] J Leng and others, 'Makerchain: A blockchain with chemical signature for self-organizing process in social manufacturing' (2019) 234 Journal of Cleaner Production 767, 773-774.

[17] For an overview, see KD Betts and KR Jaep, 'The Dawn of Fully Automated Contract Drafting: Machine Learning Breathes New Life into a Decades-Old Promise' (2017) 15 Duke Law and Technology Review 216, 218–224.

are based on machine learning, and that draw both on prior decisions in relation to the effects of particular types of clause and on the cumulative sum of past experience in relation to the types of risks that have been known to eventuate in the context of transactions of that type have the power to transform the practice of contracting,[18] and the recent successes of predictive systems for generating programming code, such as GPT-3,[19] suggest that the creation of predictive systems for generating contracts is far from inconceivable, and neither is the possibility of combining them with smart contracts and contract lifecycle management.

Nevertheless, there are two obvious and crucial baseline conditions that must be met if these developments are to underpin a sustainable and socially acceptable approach to contracting. These are, firstly, that the systems and process that underpin any autonomous system used in contracting must be resilient and secondly, that they must be trustworthy. Resilience requires systems to be capable of dealing effectively with the very wide range of requirements, interests, and contexts in which contracting is used; of avoiding, withstanding, or responding sensibly to unexpected and uncertain circumstances; and of identifying potential points of failure where their in and responding appropriately to them. Trustworthiness is more complex. Trustworthiness involves an evaluative judgement not just in relation to the technical dimensions of the functioning of an autonomous system, but more fundamentally in relation to the social aspects of its operation: the outcomes it has a propensity to produce, the interests it has a propensity to prioritise, and the impact these propensities have on different categories of social actors. In other words, as a baseline condition, trustworthiness requires not just that the AI be reliably able to produce outcomes under a range of conditions, but also that those outcomes conform to a particular normative standard or set of standards. And whilst trustworthiness can be broad-based or confined to a narrower community, for the type of systems that are the focus of this chapter it is crucial that they be broad-based. An AI that reliably produces a one-sided outcome—ensuring, in Llewellyn's evocative phrase, that 'what Big Fist wants, he gets'[20]—may well meet the baseline condition of resilience, but it is extremely unlikely to meet the condition of trustworthiness. Nor should it.

Creating autonomous systems that are resilient and trustworthy involves three sets of challenges. The first relates to the technical side of AI. Although it is common to use human metaphors in relation to AI processes, these metaphors are neither accurate nor appropriate. The AI systems that we currently have and those that are likely to be capable of being deployed in

[18] S Williams, 'Predictive Contracting' [2019] Columbia Business Law Review 621.

[19] TB Brown and others, 'Language Models are Few-Shot Learners' (*arXivorg*, 22 July 2020) ⟨https://arxiv.org/abs/2005.14165⟩ accessed 10 July 2021.

[20] K Llewellyn, 'The normative, the legal, and the law-jobs: The problem of juristic method' (1940) 49(8) Yale Law Journal 1355, 1376.

contracting in the near future are not intelligent in any sense comparable to human intelligence. They do not think, and they exercise neither judgement nor discernment. What algorithmic systems do have is a high level of competence, far exceeding that of humans, at certain types of tasks that require an ability to layer ideas or representations, to map the conceptual territory covered by those ideas and representations, and to formulate, identify, and execute actions on the basis of those maps.[21] To put it in the language of analytical philosophy, algorithmic systems have significant strengths in the field of practical reasons, but are inherently incapable of framing new grounds for epistemic reasons beyond those set out in hard-coded instructions and identified through pattern recognition. In consequence, a system will only meet the baseline condition of trustworthiness if it has either been hard-coded with directives that have a propensity to produce the types of outcomes that trustworthiness, or been trained on materials that give it an ability to identify and build into its workings the types of patterns that are associated with those outcomes. The nature of the heuristic processes on which the functioning of algorithmic systems is predicated makes it very easy for these systems to lapse into contract law minimalism, behaving like the proverbial 'steely-eyed utility maximisers', unless the system has been explicitly designed not to do so.

The second challenge relates to the legal consequences of the character of the processes that underpin the functioning of AIs. Contract law plays a strong regulatory role in relation to contract:[22] law does not just facilitate the practice of contracting and give effect to contracting, but also creates structures and scaffolding whose express purpose is to channel contracting in particular directions and to impose limits on the types of social relations and interests that contract can be used to order and govern.[23] The tools law uses, however, assume human decision-making processes. The law sets limits on opportunism, for example, by implying terms that require the parties to conduct themselves in a particular co-operative manner, by limiting the way in which and ends for which a contractual discretion may be validly exercised, by structuring the remedies available for breach in a manner that incentivises particular types of conduct, and so on.[24] The tests and concepts on which these tools rely are, however, are framed in a language that, whilst capable of easily being applied to humans, poses considerable difficulties in relation to AI. The effect is to significantly limit the regulatory capacity of

[21] A Kremer, 'Computers do not think, they are oriented in thought' (2021) 36 AI and Society 401.

[22] H Collins, *Regulating Contracts* (Oxford University Press 1999).

[23] S Hedley, 'Two Laws of Contract, or One?' in TT Arvind and J Steele (eds), *Contract Law and the Legislature: Autonomy, Expectations, and the Making of Legal Doctrine* (Hart Publishing 2020).

[24] D Campbell, 'The relational constitution of remedy: Co-operation as the implicit second principle of remedies for breach of contract' (2005) 11 Texas Wesleyan Law Review 455.

contract law in relation to AI-infused contracting. Would, for example, a person with no ability to read programming code be able to avail of the *non est factum* defence in relation to a contract implemented in the form of a computer programme?[25] If parties misunderstand the way in which an algorithm is designed to function—for example, in dynamically devising new terms to deal with emerging circumstances—does that trigger a remedy of mistake? Should a party who is in a superior position to understand the manner in which an algorithm functions be under a duty to disclose or correct a misunderstanding that the other party has (unlike the current position at common law)? These questions do not represent a mere doctrinal querulousness. It is doctrines such as these, and remedies such as rectification, that enable the courts to exercise regulatory power over contract; and if those doctrines are no longer operational a different system of assurance will have to be found for AI to meet the baseline condition of trustworthiness, and for the overall system of contracting to meet the baseline condition of resilience when the AI fails to function as expected.

Underpinning both of these is a third challenge, which goes to the heart of the impact of AI on contracting. As the introduction argued, the functioning of AI is underpinned not just by its technical design but also by the social context within which it operates. As the literature has long noted, this gives AI a marked propensity to inherit biases and limitations from that social context. AI itself is morally neutral as a matter of definition if it has not been programmed to apply a particular moral framework, but that does not mean that it will operate in a morally neutral effects. To the extent an AI replicates functions that, in humans, constitute 'judgement', those judgements are wholly derivative of the human-constructed systems that enable an AI to exercise those functions.[26] An AI that emerges from a social context that displays particular biases and predispositions, and that has not been designed to avoid those biases and predispositions, is therefore likely to have a propensity to replicate them.

The fact that the applications towards which the development of AI-infused contracting has been directed continue to follow the pattern of focusing on the medium of the contract rather than the substance of the transaction instantiates the extent to which the functionality of AI is shaped by the worldviews implicit in the human-constructed systems that underpin their design. But the social context also exerts a deeper shaping influence, which can be illustrated with reference to the analogy between contracts and narratives which is frequently deployed in the literature on contracting.[27] If

[25] The defence in its current form is generally taken to require legal incapacity. See e.g. *Saunders v Anglia Building Society* [1971] AC 1004 (HL) and *Ford v Perpetual Trustees Victoria Ltd* [2009] NSWCA 186, (2009) 257 ALR 658.

[26] J Malpas, 'The necessity of judgment' (2020) 35 AI and Society 1073, 1074.

[27] See e.g. LM Ingram and LS Jensen, 'The utility of narrative voices in the federal procurement contract' (2018) 4(1–2) Journal of Strategic Contracting and Negotiation 58.

a contract contains a transactional narrative, then an AIs strengths—its particular 'competence', in the terminology used above—lies in its ability to identify where we are in the narrative, what possible endings exist, and what path needs to be followed to get to a particular ending (or, at least, to increase the chances of arriving at that ending). An AI does not, however, have an independent ability to form a views on what constitutes a optimal ending, or an optimal way of dealing with an unexpected contingency, in the context of a specific transaction. Its 'views' are wholly shaped by the material it was given, which in turn are shaped by biases and predispositions that arise from the social context in which it is designed.[28] To achieve the baseline conditions of trustworthiness and resilience, it is essential that the processes by which AIs are designed, and the manner in which they are legally regulated, are oriented towards addressing these biases and predispositions. It is, accordingly, to examining the types of biases that the present contracting environment creates that we now turn.

III The many presents of contract: social and commercial contexts

Since the formulation in the 1960s of what has come to be termed 'relational contract theory', it has been clear that there is a real, and growing, disjunction between the transactional and legal aspects of contract. In the last twenty years, a considerable body of empirical and theoretical work on contracts and contracting has given us new and important insights into the nature of this disjunction and the impact it has on transactions. This disjunction and its impact are a fundamental part of the social context against whose backdrop AI is used in the field of contracting, and three aspects of it are of particular importance in considering the impact of AI on contracting and identifying the challenges that creating trustworthy and resilient AI is likely to pose for the law as well as for the practice of AI design.

Firstly, and in contrast to the older literature on relational contract,[29] recent empirical studies have shown that interorganisational relations tend to involve both relational and contractual governance, with the two playing complementary roles rather than being substitutes for each other.[30] Empirical work has also shown that combining relational and contractual governance has a positive effect on transactions, improving outcomes for both parties and reducing opportunistic behaviour.[31] This work, in turn, has led

[28] D Varona, Y Lizama-Mue, and JL Suárez, 'Machine learning's limitations in avoiding automation of bias' (2020) 36 AI and Society 197.

[29] S Macaulay, 'Non-contractual relations in business: a preliminary review' (1963) 28(1) American Sociology Review 55.

[30] L Poppo and T Zenger, 'Do Formal Contracts and Relational Governance Function as Substitutes or Complements' (2002) 23(8) Strategic Management Journal 707.

[31] Z Cao and F Lumineau, 'Revisiting the interplay between contractual and relational

to a more nuanced view of the diverse range of functions that contracts serve in transactions. Apart from the safeguarding function familiar to lawyers, contracts also serve as devices to structure transactional adaptation in the face of changing circumstances,[32] to structure communication and coordination during the lifetime of a transaction (for example, by creating steering groups),[33], as well as to support internal management and medium-term planning in the context of a transaction.[34]

The manner in which contracts are drafted, however, is not oriented towards enabling them to discharge these functions in a systematic or reliable way. Macneil criticised traditional approaches to drafting contracts for their intensification of presentiation, and recent research suggests that little has changed since his day. Legal scholars working within a broadly relational and empirical tradition have argued that contracts require a more 'proactive' element if they are to serve as an effective transactional management tool, but notwithstanding this they tend instead to be drafted reactively in a manner that prioritises safeguarding the parties' interests in a litigated dispute,[35] and focuses on risk allocation rather than on the strategic dimensions of contracting.[36] The effects of this on contract performance have not been happy. Research has shown, for example, that control-oriented provisions tend to be favoured by lawyers, even though contracts which rely on control-oriented provisions tend to exacerbate conflict in environments that are subject to rapid change.[37]

Much of this literature has focused on contracts where the parties are in a position of relative equality, and are thus free to bargain in pursuit of their interests. Yet this is only rarely true in the real world. The practice of contracting is marked by asymmetry and inequality of bargaining power far more frequently than it is by symmetry, and it is in this area—where the failure to structure contracts to further joint maximisation combines with

governance: A qualitative and meta-analytic investigation' (2015) 33-34 Journal of Operations Management 15, 30.

[32] R Klein Woolthuis, B Hillebrand, and B Nooteboom, 'Trust, Contract and Relationship Development' (2005) 26(6) Organization Studies 813.

[33] Y Chen and others, 'Understanding the multiple functions of construction contracts: the anatomy of FIDIC model contracts' (2018) 36(8) Construction Management and Economics 472.

[34] A Hurmerinta-Haanpää and S Viding, 'The functions of contracts in interorganizational relationships: A contract experts' perspective' (2014) 40(1) Journal of Management 193, 107–109.

[35] G Berger-Walliser, 'The past and future of proactive law: an overview of the development of the proactive law movement' in G Berger-Walliser and K Østergaard (eds), *Proactive Law in a Business Environment* (DJØF Publishing 2012) 23.

[36] LA DiMatteo, GJ Siedel, and H Haapio, 'Strategic contracting: examining the business-legal interface' in G Berger-Walliser and K Østergaard (eds), *Proactive Law in a Business Environment* (DJØF Publishing 2012).

[37] O Schilke and F Lumineau, 'The Double-Edged Effect of Contracts on Alliance Performance' (2018) 44(7) Journal of Management 2827.

terms that are heavily tilted towards one side in terms of their safeguarding function—that the transactional impact of contracting practices becomes particularly problematic. In recent work, Jenny Steele and I have devised a framework which, drawing on Mary Douglas's grid-group cultural theory, identifies four distinct approaches or perceptions of markets within legal regulation.[38] Figure 1 presents an adaptation of that framework which charts four different ways in which the drafting of contracts can affect a transaction.

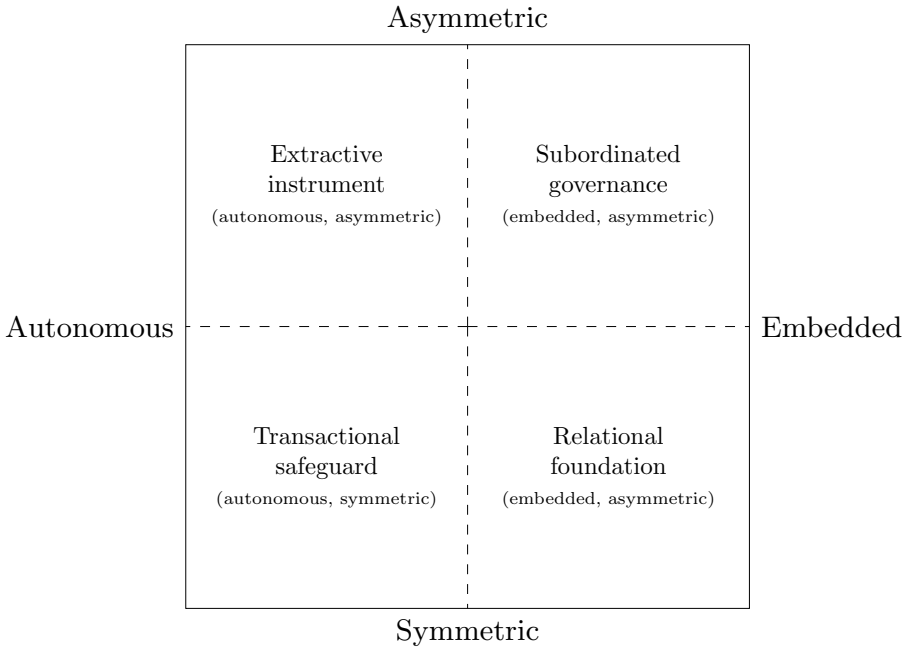

Fig. 1: The transactional impact of contracts: Four configurations

At the bottom end, transactions are largely symmetric. Positions are not fixed, and parties may be a seller one day and a buyer the next. At the top end, transactions are asymmetric, and positions are fixed: a consumer of payday lending services, for example, is unlikely to become a provider of such services, and the head contractor on a construction project and an electrical subcontractor are unlikely to ever find themselves in the opposite position. Similarly, at the left extreme contracts are treated as being autonomous of the commercial transaction: there is a sharp distinction between the 'paper deal' and the 'real deal'. At the right extreme, there is a close and embedded relationship between the contract and the transaction: the paper deal at least partially furthers the commercial purposes of the real deal.

[38] TT Arvind and J Steele, 'Remapping Contract Law: Four Perceptions of Markets' in TT Arvind and J Steele (eds), *Contract Law and the Legislature: Autonomy, Expectations, and the Making of Legal Doctrine* (Hart Publishing 2020) 439-445.

As the figure shows, at the bottom end, the effect of contracts on transactions is relatively benign. In transactions that are genuinely autonomous and symmetric—such as a commercial sale of commodified industrial goods—a contract which follows a pattern of transactional safeguarding will support the transaction relatively well: its provisions efficiently communicate to the parties how they need to reorient conduct and channel their expectations to avoid transactional conflict.[39] Similarly, a contract which relates to a symmetric transaction and is sufficiently well-specified to embed the key requirements of the transaction—for example, a well-drafted relationally oriented joint venture agreement—will operate efficiently in its role of providing a relational foundation for the transaction, which serves to define, plan, and organise the parties efforts around a positive and constructive shared goal.[40]

As we move higher along the vertical dimension into more asymmetric transactions, however, this begins to change. A contract embedded in asymmetry will create a hierarchical system of governance, in which one party—and its interests and goals—is subordinated to the other. The context and specificities of the transaction infuse the contract's terms, but the ultimate say—the right to determine how transactional governance operates, how unforeseen surpluses and losses are allocated, and how the transaction is adapted to deal with emergent circumstances—is vested in the superior party who has considerable discretion in relation to how much and in what circumstances regard will be had to the interests of other parties.[41]

The situation is not much better at the left extreme. Here, contracts serve as extractive instruments, with the contract terms becoming one-sided tools whose purpose is primarily to enable one of the parties to extract as much value as feasible from their counterparty. The paper deal displaces the real deal, because the stronger party's interests are better served through disembedding the transaction from broader social expectations. The party in a weaker position has little ability to influence the terms of the contract, giving contracts a systemic propensity to operate in a manner that is contrary to the other party's interests, while the party in a superior position has little interest or incentive to have regard to the other party's interests save to the extent necessary to prevent the transaction from disintegrating in a manner that harms the superior party's commercial interests.[42]

These propensities are real and they are deeply entrenched in the fabric

[39] There is a strong parallel between this transactional function of contracts and the law-job which Llewellyn termed 'channelling' (Llewellyn (n 20) 1376–1380). I have previously discussed the parallels between grid-group cultural theory and Llewellyn's law-job theory in TT Arvind, S Haliday, and L Stirton, 'Judicial Review and Administrative Justice' in J Tomlinson and others (eds), *The Oxford Handbook of Administrative Justice* (Oxford University Press 2021), on which the account in this chapter is largely based.

[40] In this, it closely resembles the law-job Llewellyn termed 'net drive'. See Llewellyn (n 20) 1387–1391.

[41] Cf the law-job which Llewellyn terms 'the say': ibid, 1383–1387.

[42] Cf Llewellyn's discussion of the law-job he terms 'the trouble-case': ibid, 1375–1376.

of modern contracting. And whilst AI has the potential to ameliorate these propensities, it also has the ability to exacerbate them. As the discussion in Section II has shown, when viewed as a socio-technical system, the processes that create AIs have a strong tendency to replicate biases inherited from existing social practices, and to the extent the purposes for which they are deployed remain subject to the same asymmetries as existing contracting practices, those biases will also be reflected in their use. The next section considers this issue in more detail, with reference to the four types of use to which AI is actually put. As I show, the solution lies in a new standard of transactional responsibility, which must be placed at the heart of the processes by which AI systems in the field of contracting are conceptualised, designed, built, and regulated.

IV The promise and perils of AI-infused contracting

In principle, AI-infused contracting should be readily able to overcome the biases outlined in Section III. The biases and challenges in question are to a very significant extent a product of the drafting and negotiating practices of lawyers and, in particular, their reliance on 'legalese' in drafting and on a heavily legalised and legalistic approach to negotiation.[43] Both of these make sense if viewed as an example of a 'fast and frugal' heuristic in action.[44] Using familiar, tested forms that are known to reduce the chance of bad outcomes can make sense as a heuristic even if those forms are unhelpful in terms of promoting good outcomes or a positive transactional culture.[45] AI, in contrast, is not subject to these limitations, and the superior competence (discussed in section II above) which AI-infused contracting exhibits in relation to layering and mapping a conceptual field and connecting those maps to patterns of outcomes provides a very different heuristic that has a clear and obvious potential to avoid the limitations of the default heuristic that underpins standard legal drafting practices. In their work on 'self-driving contracts', Casey and Niblett have argued that technology creates the potential for a wholly new type of contract oriented around dynamically generated 'micro-directives' that translate a general objective into a specific set of actions that are likely to achieve the objective, and can be programmed to do so on the fly thus eliminating both the need for gap-filling and the

[43] P Hietanen-Kunwald and H Haapio, 'Effective dispute prevention and resolution through proactive contract design' (2021) 5 Journal of Strategic Contracting and Negotiation XX ⟨https://doi.org/10.1177/20555636211016878⟩.

[44] See G Gigenrenzer and PM Todd, 'Fast and Frugal Heuristics: The Adaptive Toolbox' in G Gigenrenzer, PM Todd, and the ABC Research Group (eds), *Simple Heuristics that Make us Smart* (Oxford University Press 1999). For a discussion of its application to law, see C Engel and G Gigerenzer, 'Law and heuristics: An interdisciplinary venture' in C Engel and G Gigerenzer (eds), *Heuristics and the law* (MIT Press 2006).

[45] CA Hill, 'Why contracts are written in "legalese"' (2001) 77(1) Chicago-Kent Law Review 59.

disputes that arise when gaps are not properly filled.[46] Similar arguments have been made in relation to predictive AI which, it has been suggested, can eliminate a lot of the legal uncertainty around the question of what precisely a given term or set of terms requires parties to do under a particular set of circumstances.[47]

Realising these possibilities in practice, however, is likely to face two issues which, while not necessarily true of all conceivable forms of AI are nevertheless inherent in the manner in which AI-infused contracting is currently approached. Firstly, from a technical standpoint, although AI-infused contracting does indeed offer the potential to significantly improve contracting through the dynamic generation of 'micro-directives' and a more rigorously evidence-based approach to the formulation of contract terms—whether in the form of code-based micro-directives or more conventional terms—the practical utility of these terms will be diminished unless those terms, and the contract lifecycle management processes the AI follows, takes due account of the role played by relational governance in contractual transactions. Although there was at one stage a significant strand of the literature which took the view that reliance on relational techniques such as trust and flexibility was a response to transaction costs, the empirical work discussed in section III has unambiguously demonstrated that relational governance plays a distinctive role that is wholly independent of techniques to respond to transaction costs.[48] Eliminating the problem of contract gaps will not, therefore, eliminate the need for relational governance; and given the distinctiveness of relational governance and its importance to the quality of transactional outcomes, it is unlikely that AI-infused contracting will meet the baseline condition of resilience unless it is able to develop, deploy, and embed techniques of relational governance in circumstances where they are appropriate.[49]

Obtaining the data needed to create an AI capable of dealing with relational governance is far from straightforward. Not all relational governance techniques are codified in a contract, and even where they are the precise nature of the type of relationality they are intended to facilitate may not be apparent from the wording of the clause. There is, in consequence, a non-trivial risk of a conflict between the relational world of contract and

[46] A Casey and A Niblett, 'Self-driving Contracts' (2017) 43 Journal of Corporation Law 1.

[47] Williams (n 18).

[48] For a recent overview, see B Petersen and K Østergaard, 'Reconciling contracts and relational governance through strategic contracting' (2002) 33(3) Journal of Business and Industrial Marketing 265, 265–267.

[49] The qualification is important, as empirical work has shown that there are circumstances in which an excessive reliance on relationality negatively affects transactional outcomes. See JH Dyer, H Singh, and WS Hesterley, 'The relational view revisited: A dynamic perspective on value creation and value capture' (2018) 39 Strategic Management Journal 3140.

the logical world of AI, in which an AI's focus on 'micro-directives' risks shrinking the room available for relational governance.[50] Equally, although predictive contracting can build on a good understanding of the *legal* consequences of specific types of provisions based on a machine analysis of the text of judicial dicta, the data on their relational and transactional consequences is far less choate, making these consequences considerably harder to estimate through traditional machine learning techniques.

Secondly, and from a social standpoint, AI-infused contracting runs a non-trivial risk of exacerbating the problems discussed in section III. The costs of AI, the non-transparency of the (typically proprietary) algorithms that underpin it, and the legal and practical hurdles associated with acquiring access to the large quantities of data required to design functional AI, cumulatively give it an inherent propensity to accentuate pre-existing knowledge- and power-asymmetries. This is almost certain to be the case in relation to an AI trained on material which disproportionately represents the preferences of a subset of transactors—for example, larger enterprises—without that material being counterbalanced by broader material. If transactors in the 'subordinated governance' quadrant outlined in Figure 1 lack the ability to exercise a proportionate say in the making of contractual determinations, it is difficult to see how they will be able to exercise any form of say in relation to the particular type of algorithmic system that should be used to govern a contract or the propensities and predilections that system should be designed to have. This is even truer of transactors who would normally be required to accept a contract in the 'extractive instrument' quadrant as a condition of transacting. Indeed, in both cases, the proprietary character which the technology underlying the algorithmic system is likely to have will, if anything, further erode the capacity of transactors in an asymmetric transaction to influence or even have full knowledge of the nature of the predispositions that the algorithm in question exhibits. And yet, it is hard to see AI-infused contracting can meet the baseline condition of trustworthiness unless transactors who would be placed in these quadrants by existing contracting practices do in fact have such an ability.

McLuhan's tetrad of effects of technologies provides a useful way of conceptualising and pulling together the full range of social consequences that a system of AI-infused contracting subject to these limitations is likely to have. McLuhan's work, like this chapter, was concerned with aspects of technology that seek to extend the capabilities of the physical human body or mind, and it was motivated by the insight that the effects of technology were mediated, firstly, by the fact that technologies were always characterised by at least some element of social distrust in relation to the manner

[50] Cf D'Acquisto's discussion of the (closely-related) conflict between the ethical principles that underpin human action and the logical principles that underpin the operation of AI, in G D'Acquisto, 'On conflicts between ethical and logical principles in artificial intelligence' (2020) 35 AI and Society 895.

in which they were actually used and, secondly, by the learning processes through which the know-how underpinning the technology is communicated and transmitted.[51]

Based on this, McLuhan suggests that technologies are never neutral or passive. Rather, they tend to have a 'tetrad' of four effects. Firstly, they enhance, intensify, accelerate or make possible some type of human action or some aspect of the human situation. Secondly, they simultaneously also displace or render obsolete other types of action or aspects of situations. Thirdly, they retrieve or bring back older forms of action that may have previously lapsed into obsolescence. Fourthly and finally, they also have a reverse potential, of inverting their original characteristics when pushed to their limits. These effects are complementary: they operate simultaneously and when taken together and read cumulatively, they provide a means to examine how a given technology acts and affects human action.[52]

Type of effect	Manifestation in AI-infused contracting
Enhances	The ability to create proactive contracts customised to the specific circumstances of a transaction
Obsolesces	Boilerplate, standard form, reactive contracts drafted according to routinised professional heuristics
Retrieves	The return of a role for inaccessible, technical processes in taking action that produce legal effects
Reverses into	A return from contract to status

Table 1: The effects of AI-infused contracting

Table 1 summarises the manner in which these effects play out in AI-infused contracting. The first two effects—enhancement and obsolescence—are both positive. AI-infused contracting enhances the ability of parties to, through AI, create contracts and transactional management systems that are neither dependent on nor influenced by the fast-and-frugal heuristic that currently underpins virtually all of the contract lifecycle that is informed or influenced by the legal aspects of the transaction. For much the same reason, it also renders obsolete the apparatus of boilerplate contracts, as well as the reliance on complex, legalistic, and reactive documents that characterises current contract practice. These effects have been discussed at length in the literature, and therefore do not require an expansive discussion.

[51] M McLuhan and E McLuhan, *The Laws of Media: The New Science* (University of Toronto Press 1988) 93–97.

[52] ibid, 98–99.

The other two effects, in contrast, require a more detailed explanation. AI-infused contracting operates through stochastic processes that differ fundamentally from the processes that underpin ordinary human reasoning. In consequence, the reasoning underpinning the decisions made in AI-influenced contracting inherently and necessarily recedes from the level of accessibility that characterises ordinary human decision-making. The consequence is a return, as far as the processes necessary to give legal effect to actions are concerned, to something not dissimilar to Sir Edward Coke's description of the law as 'an artificial perfection of reason'. To Coke, this 'artificial reason' was critical to understanding the evaluative judgements embedded in legal rules, and the implications of those embedded evaluative judgements for the manner in which the legal system would treat individual cases. However, it was not accessible to the ordinary subject of the law. It could only be properly acquired through 'long study, observation, and experience' rather than through the application of the natural reason that all persons possess.[53] This way of thinking about the law was characteristic of the late mediaeval and early modern common law, and there are a number of pieces of popular literature from the period depicting the struggles of ordinary people to navigate the system, written from perspectives sympathising with the system[54] as well as from perspectives sympathising with ordinary people.[55] The parallel with the position of a party, particularly parties in an inferior position in asymmetric transactions, is obvious, and it suggests a strong need to pay close attention to the potentially deleterious social consequences of AI-infused contracting.

This is also true of the fourth effect, described in Table 1 as a return from contract to status. Henry Maine, famously, described the movement of progressive societies as having been 'from Status to Contract',[56] and his statement points to the important role self-determination through contract played in the emergence of modern social thought. A fundamental premise of this school of thought is the idea that markets operate as an information-communication framework, which enables individual market participants to improve their position by learning from market signals. As I have argued elsewhere, this function is threatened at a fundamental level by the growing use of data-driven autonomous algorithms which, in an asymmetric social environment, operate to limit both the amount of information communicated by market mechanisms to non-dominant participants and the ability of non-dominant participants to influence the opportunities open to them.[57] To the

[53] Co Litt 97b.

[54] See e.g. 'The complaints of the people of Stoughton', MS. Bodl. 57, f. 191v.

[55] See e.g. 'A Satyre on the Consistory Courts', BL MS Harley 2253, f. 70v, available at http://www.bl.uk/manuscripts/FullDisplay.aspx?ref=Harley_MS_2253.

[56] HS Maine, *Ancient Law: Its Connection with the Early History of Society, and its Relation to Modern Ideas* (John Murray 1861) 170.

[57] TT Arvind, 'Personalisation, Markets, and Contract: The Limits of Legal Incremen-

extent AI-infused contracting relies on data-driven algorithmic systems, it is likely to have precisely the same propensity to producing social effects of this type as other data-driven algorithms.

V The principle of transactional responsibility

How, then, can these issues be resolved, and how can AI-infused contracting be brought closer to satisfying the baseline conditions of resilience and trustworthiness? The answer, as I argue in this section, lies in developing a relationally informed principle of transactional responsibility, and embedding that principle into the processes and systems by which AI-infused contracting systems are designed, implemented, deployed, and regulated.

Until recently, the focus of contract law on the medium of contracts rather than the substance of transactions exercised a high degree of influence over the manner in which the normative ends of contracting have been conceptualised. The ends of the law were largely seen in terms of giving effect to the message contained in the medium of the contract, by providing a forum in which the requirements of that message could be authoritatively determined and in which breaches of contract could be remedied.[58] The normative core of contract law was, accordingly, conceptualised in terms that were closely related to the medium of contracts rather than the substance of transactional relations: as being about the promises made by the parties to each other,[59] or about the duties assumed by the parties to each other,[60] or about the fact and consequences of the parties' mutual consent,[61] and so on. That the law sometimes intervened to alter or compel a deviation from the contractual framework was not denied, but it was compartmentalised: as a consumer-welfarist exception to contract law's default market-individualist orientation, or as a policy-based, distributive exception to the corrective and vindicatory goals that were considered to be the normative core of contract.[62]

Recent work has begun to challenge this picture, and demonstrate that contract law can only provide a sustainable basis for contracting if it builds on a relational understanding of the tasks and ends of contract law, and of the types of outcomes that an effective law would promote. This insight, of

talism' in U Kohl and J Eisler (eds), *Data-Driven Personalisation in Markets, Politics and Law* (Cambridge University Press 2021) 114–116.

[58] This is the position taken by the corrective justice account of contracts. See esp. E Weinrib, *The Idea of Private Law* (Harvard University Press 1995).

[59] C Fried, *Contract as Promise: A Theory of Contractual Obligation* (Harvard University Press 1981).

[60] B Coote, *Contract as Assumption: Essays on a Tmeme* (Hart Publishing 2010).

[61] RE Barnett, 'Some Problems With Contract as Promise' (1992) 77 Cornell Law Review 1022.

[62] See e.g. RA Epstein, 'In defense of the contract at will' (1984) 51(4) University of Chicago Law Review 984.

course, lay at the heart of the work of the original relational contract theorists, but more recently their insight has been taken up by scholars working within a much broader set of approaches. Wielsch, for example, has put forward a relational legal analysis which argues that all legal rights have social dimensions, and the social institutions and systems on which these dimensions depend have a normativity of their own which the law must provide the room to consider comprehensively.[63] Similarly, Tan has recently put forward a helpful distinction between 'macro' and 'micro' justice, and has argued that although contract theorists have traditionally assumed that contract is focused solely on a corrective conception of micro justice, contract law embeds a distinctive, relationally oriented and relationally constrained conception of micro justice which underpins a range of positions taken by contract law in areas ranging from employment law to pre-nuptial contracts.[64] By bringing these insights into dialogue with recent work on the distributive implications of AI and the challenges of making AI more socially responsive, we can begin to construct a framework to define and bring into operation a principle of transactional responsibility which, if it underpins AI-infused contracting systems, can help make those systems resilient and trustworthy by avoiding the biases and predilections discussed in sections II and IV. In the remainder of this section, I describe the five core elements on which such a framework will be built.

Firstly, in place of the focus on the medium of contract which the current approach to AI-infused contracting has inherited from the conceptual frameworks on which contract practice is currently based, a transactionally responsible approach will have its heart the understanding that Macneil termed the 'solidary belief.' The solidary belief reflects a common belief in continued future interdependence, which is sufficiently deep and extensive to ensure that no participant in a contractual system has the power to unilaterally appropriate an undue share of the surplus generated by a transaction. Macneil argued that the solidary belief was of fundamental importance to contractual systems of private ordering. Contracts frequently incorporate terms that give one party the ability to cause disproportionate harm to another party, and if there is a widespread belief that those terms will be used by systemically stronger parties to further their interests at the expense of systemically weaker parties, the ultimate effect will be to threaten the viability of the contractual system.[65] That AI-infused contracting carries the potential to threaten the solidary belief should be obvious from the discussion in sections III and IV, and task of making it trustworthy therefore requires that it is the preservation of the solidary belief, rather than inher-

[63] D Wielsch, 'Relational Justice' (2013) 76 Law and Contemporary Problems 191.

[64] ZX Tan, 'Where the Action Is: Macro and Micro Justice in Contract Law' (2020) 83(4) Modern Law Review 725.

[65] IR MacNeil, *The New Social Contract: An Inquiry into Modern Contractual Relations* (Yale University Press 1980) 102–104.

ited ideas such as gap-filling or outcome-prediction, that must be its central function.

Secondly, in terms of the principles that inform the design of algorithmic systems in the field of contracting, it is vital that the systems be designed to be not just contractually and legally aware, but also socially, contextually, sectorally, and transactionally aware. Much of the work on designing systems for AI-infused contracting has focused on developing their ability to analyse contractual texts to extract lifecycle-relevant information, and on developing predictive capabilities in relation to the impact of legal rules on the contract's provisions. Yet, as the discussion above has demonstrated, a reslient AI system must also be able to work constructively with relational dimensions of contracts, including relational dimensions or aspects of relational governance that commonly inform transactional practice but are not codified in the actual terms of contracts. This does not require an AI to exercise judgement or discernment. Within the frameworks of non-monotonic logic that are used in programming autonomous system, it rather requires the system to be programmed to have an awareness of when relational considerations have *priority* over other considerations. This does not mean that creating sytems with this awareness is straightforward: as section IV has discussed, there are challenges in actually ascertaining what these relational practices are, and when they are triggered and thus acquire priority, these challenges are not insurmountable. There are established methods for ascertaining relevant social norms that are regularly deployed in relation to other autonomous systems, such as care robots,[66] which can with some adaptation be applied to the design of systems for AI-infused contracting.

Thirdly, in terms of the process that underpins the design of systems for AI-infused contracting, it is vital to institute systems of governance that ensure that the system is designed to have due regard to, and give due weight to, the interests of all categories of transactors who are likely to be subject to a particular algorithmic system. This is of particular importance in systemically asymmetric transactions, where neither transactional power nor transactional positions are equally distributed. As the discussion above has shown, whatever claims may be made in relation to the moral neutrality of an algorithm, the actual pattern of outcomes produced by AI-infused contracting will reflect the biases and predispositions of the social circles involved in its design, which could stretch as far as tolerating the propensity of an algorithmic system to cause disproportionate harm to certain categories of transactors. One possible way of creating such systems lies in the system of the 'negotiated economy'. The negotiated economy as a concept was advanced to explain features of governance that are common in Nordic

[66] See e.g. N McBride, 'Developing socially inspired robotics through the application of human analogy: capabilities and social practice' (2020) 35 AI and Society 857 for an example.

countries. In place of the emphasis on individual choice and economic rationality that characterises traditional approaches to contracting, the negotiated economy emphasises structuring governance processes in a manner that enables all key interest groups to be represented, and to achieve mutually acceptable outcomes through a process of negotiation and persuasion.[67] By elevating the negotiations to the level of interest groups, rather than individual transactors, the institutions of the negotiated economy deal effectively with the problems of asymmetry discussed in section III. At the same time, the direct involvement of the user groups themselves also deals effectively with the problems of limited regulatory capacity—a problem that is particularly acute in relation to algorithmic systems—and regulatory capture. It is easy to see how this system can be adapted to meet the needs of AI-infused contracting, and a robust design system will address a significant proportion of the obstacles to its trustworthiness.

Fourthly, transactionally responsible systems of AI-infused contracting will require robust systems of assurance. Assurance as a concept has a technical meaning in the field of system safety. Assurance requires the production of:

> a structured argument, supported by evidence, intended to justify that a system is acceptably assured relative to a concern (such as safety or security) in the intended operating environment.[68]

As the quotation suggests, assurance cases are typically associated with physical properties of a system, such as safety. Both the methodology and the techniques used, however, are not limited to these concerns, and are capable of being applied well beyond the immediate domain of physical safety. Any property that is capable of observation or inference from observation is capable of being assured for, and it is therefore possible to extend existing techniques and systems of assurance to assure a system for AI-infused contracting for its social, relational, and transactional awareness, as well as its distributive propensities. Achieving this will require a high level of transparency about the risks which the system has been designed to mitigate, the safeguards that have been put in place to mitigate those risks, the evidence for the effectiveness of those safeguards at actually mitigating the risks, and the manner in which the accuracy of the claims underpinning the assurance case can be tested. Assurance cases make a significant contribution to a system's actual ability to avert the risks that are the subject of the concern (and, thus, to its performance and resilience) as well as to confidence in the

[67] K Nielsen, 'The Mixed Economy, the Neoliberal Challenge, and the Negotiated Economy' (1992) 21 Journal of Socio-Economics 325.

[68] A Piovesan and E Griffor, 'Reasoning About Safety and Security: The Logic of Assurance' in E Griffor (ed), *Handbook of System Safety and Security* (Syngress 2017).

system's design (and, thus, to its trustworthiness). They are, therefore, an essential component of systems of transactional responsibility.

Fifthly and finally, actually achieving transactional responsibility will require reworking the role of law, so that it provides a scaffold for the design of transactionally responsible systems of AI-infused contracting, as well as a backstop to deal with the consequences of situations in which AI systems fail for one reason or another. The first of these will entail the creation of a statutory framework to ensure that systems for AI-infused contracting are in fact designed in conformity with the principle of transactional responsibility. The second will require a fundamentally revised and updated approach to the manner in which the law regulates contracts. The law of contract will need to specify requirements in relation to disclosure and transparency requiring the provision of accessible explanations of how a system of AI-infused contracting functions, the basis on which it identifies the specific micro-directive it formulates to deal with a situation and the basis on which the allocative decisions implicit in that micro-directive are made, the data and processes that underpin decisions on the exercise of contractual discretion, the systems of assurance used to measure and evaluate an algorithm's propensities, and so on.[69] Equally, the law will also need to evolve techniques for assessing whether an AI has in fact functioned as expected, and remedies to deal with situations where it has not so functioned—including, where appropriate, rectification, annulment, restitution, and so on. The growing literature on contract visualisation suggests ways in which this might be done, but the purpose of this chapter is not to provide a definite answer to this question as much as to highlight the importance of incorporating it into our research programmes on AI and private law.

VI Conclusions: Towards (artificially) intelligent contract design

The practice of contracting in its current form is far from problem-free, and the rise of AI-infused contracting on its face offers a powerful tool with which to begin addressing some of these problems. Nevertheless, as this chapter has sought to show, there are real and non-trivial risks that systems of AI-infused contracting will replicate, or even exacerbate, these problems. As this chapter has shown, while these risks can be mitigated, doing so will require a significant shift in the manner in which AI-infused contracting is currently approached in legal scholarship.

This chapter has suggested two baseline conditions which AI-infused contracting must satisfy for its use to be socially acceptable: a condition of resilience and a condition of trustworthiness. I have argued that meeting these conditions requires leaving behind the long shadow cast on AI by traditional contracting practices, and anchoring AI in a new principle

[69] Some of these are already provided for under the GDPR, albeit in a limited way.

which I have called the principle of transactional responsibility. As I have shown, closer attentiveness to the five elements I have identified as central to transactional responsibility can go a long way to making AI less prone to reproducing the negative social effect of existing asymmetries, and to perpetuating the problems in the current approach to contracting that make contract law a frequent source of, rather than remedy for, transactional friction. In making these points, my purpose has not been to criticise the rise of AI-infused contracting, or the elements of the scholarship that have been supportive of—and, indeed, enthusiastic about—its rise. My purpose has, rather, been to place the focus on another set of issues to which somewhat less attention has been paid, but which are nevertheless of considerable importance to the success of AI-infused contracting.