Self-Driving Contracts

Anthony J. Casey[*] & Anthony Niblett[†]

Ex post *gap filling is a central function of contract law. This is about to change. Predictive capabilities created by big data and artificial intelligence increasingly allow parties to draft contracts that fill their own gaps and interpret their own standards without adjudication. With these self-driving contracts, parties can agree to broad objectives and let automated analytics fill in the specifics based on real-time contingencies. Just as a self-driving car fills in the driving details to get its passenger to a designated end point, the self-driving contract fills in the contract details to achieve the parties' designated outcome.*

This development suggests a new focus for the doctrine and theories of contract law. Our primary goal in this Article is to introduce and develop that new focus. For example, self-driving contracts are both complete and incomplete. They are complete *in that they specify actions for every contingency. This reduces the likelihood of breach and renegotiation. It also means that notions of efficient breach and* ex post *hold-up will be of reduced importance in contract law. At the same time, self-driving contracts are also* incomplete *in ways that render current notions of definiteness and mutual assent irrelevant or at best misleading.*

Perhaps most importantly, with contracts being interpreted by their own internal software, contract law will have to focus on where that software comes from and how it operates. Markets will arise for third-party vendors who either certify or provide independent contract programming. In some cases, these will be new markets; in others, they will evolve from existing markets such as the market for contract arbitrators. Law will play a role in supporting and overseeing these markets. We explore that role, and how it will differ in markets for contracts between sophisticated parties and in markets for consumer contracts.

I. INTRODUCTION

This Article explores the theoretical implications of self-driving contracts. We define a self-driving contract as a contract that writes its own terms or fills its own gaps. To be more precise, a self-driving contract has three key features. It is an agreement where (1) the parties set only broad *ex ante* objectives; but (2) the contract uses machine-driven analytics and artificial intelligence to translate the general *ex ante* objective into a specific term or directive at the time of performance; where (3) those terms are based on information gathered after the parties execute the initial agreement.[1] We draw an analogy to self-driving or autonomous cars. Just as a passenger in a self-driving car relies on the car to determine the optimal means (direction, speed, lane choice) to travel between two locations and to update its determination to account for real-time contingencies (traffic, weather, construction), the parties to a self-driving contract agree to a shared goal and trust

---

1. We call contractual provisions that automatically translate general objectives into specific directives at the time of performance "contractual micro-directives." *See infra* Part II.A. The concept is drawn from our earlier work on micro-directives in public law. *See* Anthony J. Casey & Anthony Niblett, *The Death of Rules and Standards*, 92 IND. L.J. 1401, 1410 (2017) [hereinafter Casey & Niblett, *Death of Rules and Standards*].

in the contract to direct them on precisely how to achieve that goal in light of real-time contingencies.

   Stated in the abstract, this idea may sound like radical science fiction. But the first generation of these contracts already exists, primarily in the form of self-pricing contracts. The most familiar example can be found in the auto-insurance industry, where parties agree to price terms that adjust automatically based on computer-driven analytics.[2] Similar pricing terms can be found in dental insurance, in short-term rental agreements,[3] and in transportation services.[4]

   As technologies that allow for predictive accuracy and ubiquitous monitoring and communication advance, self-driving contracts will proliferate. Predictive technologies will provide increased information that allows parties to more precisely choose the actions that will benefit their mutual interests. Parties will have the ability to predict with high confidence that, given scenario X, action Y is the optimal course of action to achieve their agreed upon goal. Monitoring technologies will provide access to the information necessary to determine whether scenario X has in fact occurred. And communication technologies will transmit the directives resulting from the analysis to the relevant actors.[5] These technologies pave the way for self-driving contracts.[6]

---

   2.   Nearly a decade ago, insurers introduced programs that change the price (through a discount or surcharge) for insurance coverage after the contract is executed based on data gathered from devices placed in people's cars. These programs collect information on how fast, how far, and when a person drives, as well as her braking patterns to determine the optimal price. The parties essentially agree that the price term in the contract will be set *ex post* by computer-driven analytics. *See Drivewise From Allstate®,* ALLSTATE, http://www.allstate.com/drive-wise.aspx (last visited on Oct. 31, 2017) (discussing Drivewise rewards for Allstate customers); *Snapshot® Plug-in Device Terms & Conditions*, PROGRESSIVE (May 11, 2017), http://www.progressive.com/auto/snapshot-terms-conditions/ (explaining the Snapshot program, data, and device). Note that they are primitive in that they currently set prices going forward rather than in real time and let you opt out. But that is likely to change.

   3.   *See How do I turn Smart Pricing on or off?*, AIRBNB, http://www.airbnb.com/help/article/1168/how-do-i-turn-smart-pricing-on-or-off (last visited on Oct. 31, 2017) (explaining how to set prices to automatically go up or down based on similar listings).

   4.   *See Prime Time for drivers*, LYFT, http://help.lyft.com/hc/en-us/articles/214586017 (last visited on Aug. 31, 2017) (explaining Primetime for drivers); *What is Surge?*, UBER, http://help.uber.com/h/e9375d5e-917b-4bc5-8142-23b89a440eec (last visited on Aug. 31, 2017) (explaining surge pricing).

   5.   There is evidence that contract parties are eager to enter into such contracts as the technology becomes available. Contract trends over the last several decades demonstrate a willingness among contract parties to accept the broad information exchange necessary as well as to accept the notion of entering into a contract where they are informed about the general purpose and objective but not about the specific provisions and requirements, *see infra* Part II.A. and II.B.3. Monitoring and information exchange are common for sophisticated business and financial transactions and in consumer transaction examples abound. Consumers have installed monitoring devices in exchange for insurance discounts. Uber drivers and passengers have voluntarily submitted to extensive monitoring in exchange for transportation services. More broadly, user agreements for various software applications often include voluntary information transfers that allow counterparties to monitor behavior that is unrelated to the exchange at hand. There may be concerns that this consumer willingness to share information is undesirable. When that is the case, government oversight will be particularly important. *See infra* Part IV.C.

   6.   We have explained elsewhere how such technology may—years from now—automate the filling of gaps in the public laws that government entities promulgate. *See* Casey & Niblett, *Death of Rules and Standards*, *supra* note 1, at 1419–20 (explaining how technology may one day fill the gaps in public law now filled by government entities); Anthony J. Casey & Anthony Niblett, *Self-Driving Laws*, 66 U. TORONTO L.J. 429 (2016) [hereinafter Casey & Niblett, *Self-Driving Laws*]. But the progress there is likely to be slower than with contracts. While predictive technologies may be equally available for use in public and private law, the same is not true of communication technology. People are more skeptical of ubiquitous government monitoring.

The emergence of this form of contracting will have a significant impact on the way we think and talk about contracts. First, it will reveal that the existing language of contract theory is deficient. Self-driving contracts blur the distinctions between rules and standards, between *ex ante* agreements and *ex post* dispute resolutions, and even between complete and incomplete contracts. For example, while some will view self-driving and related contracts as hopelessly incomplete, the opposite view—that they are among the most complete contracts—equally consistent with existing literature. This indeterminacy in the language of contract theory and law has led scholars (and will lead courts) to struggle in interpreting and enforcing these contracts.[7] Notions of assent, definiteness, agreements to agree, unconscionability, mutual mistake, renegotiation, and even efficient breach cannot be cleanly transported to the world of self-driving contracts. Analyzing these questions through the language of self-driving contracts may provide important clarity that is currently missing.

Similarly, some recent proposals for resolving contractual ambiguities (such as using survey technologies) can be understood more clearly once one recognizes that they are just particular forms of self-driving contracts.[8] Thinking about those proposals through the lens and language of self-driving contacts will reveal their costs and benefits—and how the law should treat them—in ways that traditional contract theory does not.

Most importantly, we identify the questions of who provides the self-driving contract and how to govern the providers as the central challenges for the law. In that context, we will explore first-party providers, third-party providers, and government providers. We draw distinctions between how the law views these alternative providers in sophisticated-party transactions, on the one hand, and consumer transactions, on the other.

This Article, thus, has two primary aims. First, we begin developing a new language of contract theory, so scholars and lawyers can fully grasp what is at stake when disputes about self-driving contracts arise. Second, we explore some fundamental implications of the emergence of self-driving contracts.[9]

We proceed in three parts. In Part II, we describe self-driving contracts and the technology behind them. In Part III, we explore the challenges that self-driving contracts pose for the language and theory of contract law. We briefly lay out some doctrinal and

---

7. *See, e.g.*, Lauren Henry Scholz, *Algorithmic Contracts*, 20 STAN. TECH L. REV. (forthcoming 2017) (showing indeterminacy leading to a struggle in interpreting and enforcing contracts).

8. *See, e.g.*, Omri Ben-Shahar & Lior Strahilevitz, *Interpreting Contracts via Surveys and Experiments* (U. Chi. Coase-Sandor Inst. for L. & Econ., Working Paper No. 791, 2017), https://ssrn.com/abstract=2905873 (showing examples of proposals for resolving contractual ambiguities with survey data).

9. Others have noted the possibility of self-driving contracts in passing. *See* Gillian Hadfield, *Producing Law for Innovation*, *in* RULES FOR GROWTH, 26–27 (2011) [hereinafter Hadfield, *Producing Law for Innovation*]. Few have grappled with the larger implications. A notable exception is an emerging draft project by Lauren Scholz. Her primary focus is on self-enforcing contracts and she provides important insights on how the law should view their formation and deal with accountability for externalities. We view self-enforcing smart contracts as a separate and narrower topic, outside the scope of this article. But Scholz does, at times, lump automated gap-filling in with her broad category of algorithmic contracts. To the extent she implies that terms generated by automated algorithmic gap-filling are not or should not be enforceable, we emphatically disagree. Our analysis below shows that these terms are no different—for enforceability purposes—than standards that a judge adjudicates or unread terms in a hundred-page contract. We also do not think rules regarding third-party effects need to be any different for a self-driving contract than for a traditional contract. Our focus is instead on the implications for foundational theories about complete contracts and renegotiation that self-driving contracts raise and on the market structures that will emerge to facilitate these contracts. *See* Scholz, *supra* note 7.

policy implications related to those challenges.

In Part IV, we show that the most important challenge for the law of self-driving contracts relates to the markets that will provide these self-driving contracts. We discuss the likely rise of a market for third-party vendors providing certified computer code to govern contractual relationships. We predict that arbitration services and insurance providers will evolve and likely fill the space as the providers of self-driving contracts. Other platform providers like Uber or Airbnb are also likely candidates to fill this market. Here, we also discuss the important distinctions between the markets for self-driving contracts in consumer transactions, on the one hand, and for self-driving contracts in sophisticated business and finance transactions on the other hand.

Before we proceed, a quick word on what this Article is not. This Article does not address self-enforcing contracts and similar technologies like blockchain.[10] These contracts—often referred to as "smart contracts"—innovate on a different plane than self-driving contracts. They contain automated provisions that allow for enforcement of specific provisions without resort to adjudication. In this way, they are extensions of escrow technologies like the letter of credit. The innovation, however, is that the self-executing contracts can trigger automated enforcement without a third-party intermediary—that is, without the escrow agent or the bank which issues the letter of credit.[11]

Self-driving contracts are a completely different category of contracts. Self-driving contracts require judicial enforcement.[12] But they automate the creation and interpretation of the *terms* that are to be enforced. The key is that the contract uses technologies like artificial intelligence and learning algorithms to update the terms governing the parties' relationships based on contingencies as they arise.[13]

II. THE CONCEPT OF THE SELF-DRIVING CONTRACT

The self-driving contract, in its pure form, is a simple concept. Parties agree to an outcome they want to achieve and rely on machine-assisted analytics to direct them toward that outcome. The parties might agree to achieve a given outcome, X, and divide the surplus. In the extreme, there are only two terms to which the parties need to agree: (1) the desired outcome and (2) how to divide the surplus. Later, when it comes time to perform, the parties can rely on the machine to set all other terms of performance. The machine

---

10.   *See* Trevor I. Kiviat, *Beyond Bitcoin: Issues in Regulating Blockchain Transactions*, 65 DUKE L.J. 569, 606 (2015) (discussing blockchain transactions and automatically executing contracts). Other technological enhancements to contract are also beyond the scope of this paper. *See, e.g.*, Anthony J. Bellia Jr., *Contracting with Electronic Agents,* 50 EMORY L.J. 1047 (2001) (discussing the ability to use electronic agents to enter into contracts); Harry Surden, *Computable Contracts*, 46 U.C. DAVIS L. REV. 629 (2012) (discussing the ability to express specific contract terms in computable code).

11.   *Id*. at 605–06; *see* Joshua A.T. Fairfield, *Smart Contracts, Bitcoin Bots, and Consumer Protection*, 71 WASH. & LEE L. REV. ONLINE 35, 45–47 (2014) (emphasizing the that the main innovation of smart contracting is the elimination of intermediaries).

12.   The combination of these technologies might pose other interesting issues worth future study.

13.   Self-executing, self-driving, and other machine-assisted contracts are sometimes described collectively as "algorithmic contracts." *See, e.g*., Scholz, *supra* note 7. We will generally avoid that term. Every contract—not just one that incorporates machine technology—is an algorithm. An algorithm is a set of ordered steps or operations that one follows to achieve an outcome. All but the simplest algorithms specify a number of "if-then" statements (if X, then do Y) to take into account uncertain external inputs. For example, an algorithm for keeping dry might include this step: "If it rains, then use an umbrella." A contract performs the same function. It is merely a set of operations, which we call obligations, that each side must perform to complete the contract.

gathers information[14] about the present state of the world and makes predictions about what actions the parties should take to best achieve outcome X and then directs the parties to take those actions. It also identifies the surplus that exists and directs the parties to divide that surplus according to their agreed objectives. To do these things, the machine monitors both the external world and the parties' behavior. As facts in the world change, the machine updates its directives, and breach is defined as failure to comply with those directives.

While fully self-driving contracts do not currently exist, early partial forms are emerging in flexible-pricing arrangements. The auto-insurance example is the most salient. A driver can install a monitoring device that gathers information on the driver's skill and habits, and the pricing terms of the contract automatically update.[15] The essence of this transaction is one where the parties have agreed up front about the division of surplus, but certain facts about the driver were unknown to them at the time of agreement.[16] As a result, they do not know what the surplus will be or how to maximize it. After the information is gathered, the price changes. That price change has two effects. First, it divides the surplus that is now known to exist. Second, it changes the driver's incentives in a way that could maximize the surplus. The driver may seek to change her behavior—by driving less or taking classes or leaving for work earlier—if the price change makes these behavioral changes more attractive.

Of course, a more advanced self-driving contract would react more quickly. It would tell the driver what to do to reduce risk *while she was driving*. We are not there yet. But the technological steps to get from here to there are small. And once we are there, things get interesting.

The advanced version of this technology could be expanded to other flexible-pricing agreements. For example, long-term rental agreements do not specify the rental price in the event of renewal. Instead, these agreements use terms like "reasonable" or "market price." Filling in the substance of those terms requires negotiation or litigation. But a self-driving contract could simply calculate the joint-welfare maximizing price every year—or every second for that matter. As a result, no negotiation or arbitration would be required.

Beyond pricing, parties to sales of goods contracts could use machine analytics to determine when a party could refuse delivery. A contract aimed at maximizing joint welfare of the parties could take into account changes to external factors, opportunities to mitigate, incentive effects on the parties, and the parties' preferences for risk allocation.[17] The same could be true of service contracts that require "reasonable efforts." Perhaps, if the data become available, even merger agreements could use self-driving contracts to fill the specifics of the notoriously vague "material adverse change" clauses.

As we have said, we are not there yet. Self-driving contracts are in their infancy, but it will be a short infancy. To show why, the remainder of this Part will focus on the technology behind self-driving contracts. We start by explaining the concept of contractual

---

14. The parties might or might not specify the allowable sources from which the machine gathers the information.

15. *See supra* note 2 (showing that the pricing terms of the contract automatically update if a monitoring device is installed for Allstate's Drivewise program).

16. One might object that the driver did know the facts. That is a frivolous objection. Drivers are famously bad at understanding their own skill level. The driver may have some basic information about her driving skills but she does not fully understand how that information translates into driving risk.

17. These are of course the same things that arbitrators take into account, but the machines will have and understand more data and will provide greater consistency. More on this *infra* Part III.A.

micro-directives, which are the building blocks behind self-driving contracts. We then briefly explain the specific technological advances that make them possible.

A. Contractual Micro-Directives: The Building Blocks of Self-Driving Contracts

A perfectly written contingent contract term would reflect full information about the state of the current world and about the state of all possible future worlds.[18] Moreover, the contract would include mechanisms for *ex post* assessment about which of the possible worlds actually did materialize, but parties to a contract never have full information about present or future states of the world. They must choose between (1) acquiring more information, (2) contracting over more contingencies, or (3) leaving gaps to be filled later. After executing the contract, parties also face the costs of observing and verifying information for enforcement purposes. Where those costs are too high, some terms can never be enforced. As many economists have pointed out, contract drafting, therefore, requires a trade-off between the costs of incompleteness and the costs of negotiation and information gathering.[19]

A contractual micro-directive, however, collapses that tradeoff. Generally speaking, a micro-directive is a form of law that uses machine-based algorithms to translate a general *ex ante* objective into a specific *ex post* directive.[20] In essence the micro-directive is a law that has the benefits of a standard (low information and negotiating costs) and the benefits of a rule (certainty, commitment, and predetermination).[21] With contractual micro-directives, the parties agree to a stated objective. The machine-driven analytics later translate that objective into a specific directive and communicate it to the parties. The parties must follow the directive to perform (that is, avoid breaching) the contract. In this way, the contractual micro-directive can take the place of a vague or ambiguous contractual standard, a gap, or a cumbersome list of specific rules.

A contractual micro-directive is, then, a self-driving term that (at least from the human perspective) fills its own gaps and writes its own terms after the parties have executed the agreement. It arbitrates its own ambiguities. From a more technical sense, however, it might be an immensely detailed contract term that has virtually no gaps to fill—a long form

---

18. *See* Steven Shavell, *Damage Measures for Breach of Contract*, 11 BELL J. ECON. 466, 466–67 (1980) (discussing perfect, fully contingent contracts and the benefits of a party contracting around every conceivable circumstance).

19. *See* Richard T. Holden & Anup Malani, *Renegotiation Design by Contract*, 81 U. CHI. L. REV. 151, 151–54 (2014) (describing the incomplete nature of contracts and the value of spending more time negotiating versus *ex ante* bargaining); Philippe Aghion & Richard Holden, *Incomplete Contracts and the Theory of the Firm: What Have We Learned over the Past 25 Years?*, 25 J. ECON. PERSP. 181, 194 (2011) (discussing how contingent contract negotiation depends on the negotiator's ability to conceptualize contract risk); Avery W. Katz, *Contractual Incompleteness: A Transactional Perspective*, 56 CASE W. RES. L. REV. 169, 183 (2005) (suggesting that parties may negotiate for future arbitration, cooperation, or exit strategies rather than attempt a perfect contingent term).

20. *See* Casey & Niblett, *Death of Rules and Standards*, *supra* note 1, at 1436–40 (explaining micro-directives).

21. In prior work, we have explored the details of micro-directives and their potential future impact on public law regulation. While micro-directives in public law are a potential development in the future, their use in and impact on private ordering—in the form of self-driving contracts—is immediate and certain. There is some well-grounded skepticism about the likelihood of a successful effort to legislate or regulate through micro-directives. Such doubts are not present in private ordering. Parties have already demonstrated a desire to use micro-directives as a means to contracting in many circumstances. *See id.* at 1404–05.

contract to which the parties have agreed to be bound, with knowledge of their joint contractual purpose but without knowledge of the specific terms to which they have agreed.[22]

An example can illustrate how a contractual micro-directive works. Parties to a merger contract might desire that each party take certain actions to facilitate closing of the deal depending on what contingent situations arise. One way to write that obligation into a contract would be to provide a complete list of all "approved" actions for every scenario. The contract would have to specify that if X event occurs, then Y action would be approved, but if X event does not occur then Y action would not be approved. As the parties identify more and more possible contingencies, the list would become enormous. Because no human parties can write (or negotiate) such a list, they will consider alternatives.

First, they could use a blunt rule that is not well calibrated. That could take the form of a list of all the things that are approved, noting that the list is exhaustive. Thus, everything not on the list would not be approved. This provides commitment and certainty—the parties know exactly what is and is not required—but it is blunt in application. When unexpected contingencies arise (or even expected ones that the parties did not have the time or energy to include in the list), the rule will misfire and produce results not intended by the parties. This is the classic problem with rules.

Second, they could just write incomplete lists of actions that are approved and not approved. Courts would fill in the gaps in the list later. Parties do this all the time. It allows for *ex post* calibration by the courts. But it lacks predictability and requires the parties to incur high *ex post* litigation costs. Third, they could use a vague standard that also requires courts to fill in the details after the fact. This could be a clause that requires something like "reasonable efforts," "best efforts," or "commercially reasonable efforts."

Each of these options brings different costs and benefits. The parties must therefore decide among them based on considerations such as their estimates of the likelihood of variance in outcomes (rules are better for low variance outcomes) and the difficulty in *ex ante* information gathering and negotiation as compared to the difficulty in *ex post* information gathering and litigation (or renegotiation).

The contractual micro-directive provides an alternative solution. The parties could agree to their general objective: let us assume that objective is to close the deal if doing so increases the parties' joint surplus. Then they could allow the machine-driven algorithm to process all relevant data and instruct the parties on which actions are cost justified and which are not.

On its face, this is not so different from a standard that is later adjudicated.[23] It just uses a machine to fill the gaps as they become relevant rather than a judge. The machine updates, just as a judge would, based on *ex post* information that is fed into the predictive analytics. In this way, the contractual micro-directive has the flexibility of a standard or a gap in the contract.[24]

Yet with contractual micro-directives, the parties have greater certainty about what

---

22.   More on this distinction *infra* Part III.

23.   Indeed, the parties adopting a standard could also use the language of "increasing joint surplus" as the standard for the judge to interpret.

24.   A micro-directive could be designed to maximize surplus based purely on machine driven analysis of surplus or it could be designed to mimic what judges would do when faced with contracts requiring joint-surplus maximization. The substantive outcomes may be different, but the process and the implications for our inquiry are the same.

they have agreed to and who will determine what is required. Indeed, for any given hypothetical state of the world, they can ask the machine what instruction would result. Moreover, the answer is predetermined and the parties are committed to that answer, which in a very real sense exists within the algorithm's architecture from the moment the contract becomes effective.

The machine (and the parties) are pre-committed to a decision principle without knowing the exact *ex post* distributional consequences of that principle. The machine does not change its answer based on *ex post* distributional biases (unless of course the parties programmed it to do that). This pre-commitment is not possible when a judge is asked to fill gaps. The judge may apply the standard to maximize joint surplus, or she might insert other factors into her decision and those factors might include distributional biases to which the parties would have never agreed. The micro-directive prevents this *ex post* rationalization. It operates as if the parties agree *ex ante* to an outcome for every possible dispute and sealed the answers in various envelopes—each one marked "open if X occurs," where X defines a precise state of the world. [25]

In this way, the contractual micro-directive provides the certainty and commitment of a hard rule, but with more precision. The micro-directive can account for contingencies that arise in a way that human-drafted rules cannot. Humans are simply not capable—because of limitations on cognition and drafting speed—of writing enough rules to cover anywhere near as many contingencies as a computer-driven algorithm can cover. Traditional contracts—human-driven algorithms—are too simple and rigid for that level of specificity.

A skeptic might think that parties would be unwilling to enter into a contract where the specific terms are predetermined but unknown to the parties. We draw on two observations—beyond the fact that the first generation of self-driving contracts are emerging—to dispute this skepticism. First, the increasing length of contracts generally[26]—and to some extent, the ubiquity of form contracts—suggests that in many circumstances parties may desire to be bound by detailed context-specific terms, even when they lack the cognitive capacity to comprehend all of those terms, if those terms are targeted to achieve the agreed upon purpose of their relationship.[27] In that sense, a self-driving contract is just the logical extension of a current trend toward contracts that span hundreds

---

25. There are interesting questions about predetermination when we are talking about learning algorithms and artificial intelligence. For example, a machine that is designed to learn from experience might produce different outcomes in similar cases over time—not unlike the way courts update precedent. This process looks different from a conventional notion of certainty or predetermination. That said, if you define the context in which the machine acts to include the information that caused the machine to learn, then the machine's outcome was in fact predetermined. To put it another way, the machine is committed to that outcome and does not have its own free will. That is, in part, why we refer to machine intelligence as "artificial." One can open an additional philosophical can of worms by challenging the notion of *human* free will, suggesting that machine adjudications will be no different (in any way) from human ones. We need not, and do not, delve into these questions here. Our focus is to identify the trend toward self-driving contracts and discuss the practical implications for law. We leave the general existential questions raised by emergent artificial intelligence to those trained in philosophy.

26. John C. Coates, IV, *Why Have M&A Contracts Grown? Evidence from Twenty Years of Deals*, (Harvard Law Sch. John M. Olin Ctr., Discussion Paper No. 889, 2016), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2862019.

27. This is true in sophisticated transactions at least. We discuss the different treatments of self-driving contracts that might be necessary for unsophisticated parties below in Part IV.

of pages.[28]

Second, *ex post* adjudication may have precisely the same effect. Parties who leave gaps in a contract—or who rely on a broad standard—are implicitly agreeing to have the gap (or the details for the standard) filled in later by algorithms based on the specific state of the world that has materialized. The algorithm they rely on may be found inside the arbitrator's head. But it is no less an algorithm. In that sense, a self-driving contract is just the logical extension of judicial or arbitral gap filling.[29] Thus, parties are willing to agree to hidden terms in a contract. Those terms might be found in a long contract, in a judge's head, or in the code of a machine-driven algorithm. But parties, nonetheless, agree to be bound by their force. The difference with self-driving contracts is that they cover more contingencies than a long contract, and, unlike the judge or arbitrator's mind, the characteristics are committed and predetermined at the time of contracting.[30]

B. The Technology Behind Contractual Micro-directives

We have explained elsewhere how prediction and communication technology could support micro-directives in public law.[31] The ideas are essentially the same for contractual micro-directives. We briefly outline the concepts here, pointing out any distinctions that arise in the private contracting context.[32]

Humans have cognitive limitations that make it difficult to identify, understand, and predict current and future states of the world.[33] As a result, parties to a contract face constraints on how context specific they can make their agreements, but those are shrinking. As legal and other technologies improve, parties can write contracts that cover more contingencies and specify their intent over more possible states of the world. While the improvement thus far has been gradual, an inflection point is at hand: recent advances are facilitating contracts that fill their own gaps. The technologies at work here fall into three categories, each associated with reducing a different cost of contracting: predictive technology, communication technology, and monitoring technology.

---

28. One might counter that some of these contracts are just framing points for relational or norm based interactions. That is not inconsistent with the use of a self-driving contract, which can allow space for norms to control certain aspects of the relationship. *See* Stewart Macaulay, *Non-Contractual Relations in Business: A Preliminary Study*, 28 AM. SOCIOL. REV. 55–67 (1963); Lisa Bernstein, *Opting Out of the Legal System: Extralegal Contractual Relations in the Diamond Industry*, 21 J. LEGAL STUD. 115 (1992); Lisa Bernstein, *Beyond Relational Contracts: Social Capital and Network Governance in Procurement Contracts*, 7 J. LEGAL ANALYSIS 561, 562–65 (2015); Randy E. Barnett, *Conflicting Visions: A Critique of Ian Macneil's Relational Theory of Contract*, 78 VA. L. REV. 1175, 1176–82 (1992).

29. *Cf.* Andrew Verstein, *Ex Tempore Contracting*, 55 WM. &. MARY L. REV. 1869, 1899 (2014) (presenting evidence that parties intentionally leave gaps to be filled in by delegated third parties, thus binding themselves to the *ex tempore* contract terms provided by those third parties).

30. Again, one could make all sorts of philosophical arguments about whether the future outcome from an unknown judge's mind is fixed and inevitable. These are questions beyond the scope of this Article. Our point is simply that in one scenario you rely on a human mind as an algorithm to fill in gaps and in the other you rely on a computer program that is programmed *ex ante*.

31. *See* Casey & Niblett, *Death of Rules and Standards*, *supra* note 1, at 1423–33.

32. For full detail, we refer our readers to Casey & Niblett, *Death of Rules and Standards*, *supra* note 1; Casey & Niblett, *Self-Driving Laws*, *supra* note 6.

33. William M. Grove & Paul E. Meehl, *Comparative Efficiency of Informal (Subjective, Impressionistic) and Formal (Mechanical, Algorithmic) Prediction Procedures: The Clinical-Statistical Controversy*, 2 PSYCHOL. PUB. POL'Y & L. 293 (1996); *see* Jean Tirole, *Cognition and Incomplete Contracts*, 99 AM. ECON. REV. 265, 266–67 (2008) (discussing cognition of contracting parties).

1. Predictive Technology and Drafting Costs

First, advances in predictive technology—particularly in the form of big data and learning algorithms—allow parties to predict, analyze, and direct the proper response to a vast array of contingent states of the world.[34] This technology reduces the costs associated with *drafting* a contingent contract.

Drafting costs can be separated into three sub groups: *thinking costs* (the costs associated with comprehending the possible contingencies and their effects), *writing costs* (the costs of fully specifying each contingency and the duties associated with it), and *negotiation costs* (the costs of bargaining over which contingencies are important and how the contract should deal with them). Analytics and big data reduce all of these costs.

With regard to thinking costs, machine-driven analytics can collect data on previous transactions and use that data to identify, account for, and predict contingencies for a vast universe of possible outcomes. A machine's capacity for this greatly exceeds that of a human.[35] Moreover, machine analytics can continue to update and improve after they are put in place. This is the cutting edge of machine analytics embodied in artificial intelligence and learning algorithms.[36] With regard to writing costs, machines can specify and store massive lists of possible contingences and the obligations associated with each. Again, the difference between machine and human capacity here is enormous.[37]

Finally, predictive technology can lower negotiation costs by identifying previously unrecognized mutually beneficial outcomes and actions. For a contract to be fully specified, the consequences of each scenario and the concomitant obligations would need to be thought through and bargained over. When parties can process data on how such clauses operate in thousands of cases, the value of particular clauses to each party can be more easily estimated. This increase in certainty reduces the negotiation to one issue, how to divide the surplus rather than how to create it.[38]

2. Communication Technology and Performance Costs

Predictive technologies only go so far in reducing the costs of contingent contracting.

---

34. *See* DONALD MICHIE & RORY JOHNSTON, THE CREATIVE COMPUTER: MACHINE INTELLIGENCE AND HUMAN KNOWLEDGE 94–114 (1984) (showing how advances in predictive technology allow people to predict, analyze, and respond to many contingencies in the world).

35. As discussed in our previous article:

> Computation power is growing at exponential rates. The consistent trend of the last fifty years suggests that that power will, by the end of this century, be more than one trillion times greater than what it is today. With even a fraction of that processing power, tomorrow's computers will be able to gather and analyze more facts than any human lawmaker or judge.

Mark Lundstrom, *Moore's Law Forever?*, 299 SCI. 210, 210 (2003); *see* Casey & Niblett, *Death of Rules and Standards*, *supra* note 1, at 1411; *see generally* Gordon E. Moore, *Cramming More Components Onto Integrated Circuits*, 38 ELECTRONICS 114, 114 (Apr. 19, 1965) (introducing Moore's law).

36. *See* Gideon Lewis-Kraus, *The Great A.I. Awakening*, N.Y. TIMES MAG. (Dec. 14, 2016), https://www.nytimes.com/2016/12/14/magazine/the-great-ai-awakening.html (discussing the impact of innovative artificial intelligence).

37. *See id.* (discussing expanded use of artificial intelligence and its benefit).

38. There may, of course, be instances where certainty of this sort may increase bargaining friction but on the whole reduction is more likely. Robert Cooter, *The Cost of Coase*, 11 J. LEGAL STUD. 1, 4 (1982); *see* Oliver E. Williamson, *Transaction-Cost Economics: The Governance of Contractual Relations*, 22 J.L. & ECON. 233, 261 (1979) (arguing that bargains with high transaction costs tend to be those with high levels of uncertainty).

Any contract that contained a vast catalog of specific instructions about precise obligations under every contingency would be enormous, unwieldy, and incomprehensible. Locating the relevant rule in such a contract at any given point in time would be difficult. This is where technologies for communicating notice come in.

Communication technologies allow parties to more cost effectively access and understand complex information. As these technologies continue to improve, they will facilitate the understanding of and compliance with the directives provided by the predictive technology. Communication technology facilitates information sharing between the contract and all parties. One might imagine that the contract is stored on one's phone that constantly updates and lets each party know of relevant, specific obligations at any given time. These technologies reduce the costs of *performing* a contingent contract.

3. Monitoring Technology and Enforcement Costs

In order to enforce context-specific micro-directives that are communicated to the parties in real time, the contract must also observe the context—the relevant state of the world—and the actual behavior of the parties. Improvements in monitoring technologies reduce these enforcement costs and make contingent contracts more complete.[39] Smart devices today can record and transmit massive data about the states of the world and the parties' response to them.[40] In the consumer sphere, smartphones, fitness trackers, and even smart toothbrushes already collect enormous amounts of information relevant to contractual relationships.[41] Collections of such data have, rightly, given rise to questions of invasions of privacy.[42] Much has been written about the potential for exploitation that may occur in consumer transactions. Dealing with those same concerns will be an important part of how the law deals with self-driving contracts. We discuss these concerns

---

39. The inability to observe and verify the actions of parties provides another motivator for leaving gaps in contracts. Even if all the contingencies and actions could be fully specified, there are costs of monitoring the behavior of the other party. And even if one contract party could observe the effort levels of the other, it may not be possible for a third party (i.e., a court) to verify this information. There is less reason to (other than for moral signaling or establishing focal points) include precise rules in a contract if compliance with the rules cannot be observed or verified. The parties may not bother to contract over these issues at all or they may use weak proxies for these non-observable or non-verifiable metrics. Or, they may elect to simply include vague terms—such as "reasonable" or "material"—in order to address these concerns at a later time. These standards allow for the possibility that new proxies might be discovered *ex post*.

40. *See* Scott R. Peppet, *Regulating the Internet of Things: First Steps Towards Managing Discrimination, Privacy, Security & Consent*, 93 TEX. L. REV. 85 (2014) (consumer products with electronic sensors collect large amounts of information).

41. *See Our Technology*, FITBIT, http://www.fitbit.com/technology (last visited on Aug. 31, 2017) (discussing how Fitbit has transformed lives by keeping track of important bodily information); Douglas Macmillan, *Uber's App Will Soon Begin Tracking Driving Behavior*, WALL ST. J., https://www.wsj.com/articles/ubers-app-will-soon-begin-tracking-driving-behavior-1467194404 (last visited Aug. 31, 2017). For example, some dental insurance contracts update pricing based on the brushing habits of the insured party. We discuss below in Part IV.D, the question of whether such monitoring is too invasive and interferes with privacy.

42. *See, e.g.*, Lisa M. Austin, *Privacy and the Question of Technology*, 22 LAW & PHIL. 119 (2003); Omer Tene & Jules Polonetsky, *Privacy in the Age of Big Data: A Time for Big Decisions*, 64 STAN. L. REV. ONLINE 63 (2012); Paul M. Schwartz, *Information Privacy in the Cloud*, 161 U. PA. L. REV. 1623 (2013); Paul Ohm, *Broken Promises of Privacy: Responding to the Surprising Failure of Anonymization*, 57 UCLA L. REV. 1701 (2010); Richard A. Posner, *Privacy, Surveillance, and Law*, 75 U. CHI. L. REV. 245 (2008) (showing examples of the questioning of collections of data).

in greater detail below in Part IV.

In more sophisticated business and financial relationships, the possibilities are broader. Ubiquitous monitoring technologies allow parties to agree to instantaneous verification of compliance (or lack thereof) with a micro-directive and, thus, reduce the cost of *enforcing* contingent contracts.[43] We are less concerned about invasions of privacy where sophisticated parties knowingly enter into business transactions. Some may argue that parties may be reluctant to enter into contracts where the self-driving terms depend on sensitive information, such as costs or the details of other contractual deals. We envision algorithms developed by third parties that rely on sensitive information in order to develop the contract terms but do not reveal the content of that sensitive information to other parties. On the other hand, vague standards that are currently subject to arbitration or other means of gap filling would likely turn on the same information and may, in the process, be revealed to the other party to the contract.

The three technologies discussed here meld to enable a self-driving contract. Instead of agreeing to specific terms, the parties to a self-driving contract agree to a broad *ex ante* general objective that the machine-driven contract later translates into a specific directive when it comes time for the parties to perform. First, the contract takes the objective of the parties and, using predictive analytics, generates a vast catalog of potential obligations under different contingencies. Second, the contract observes and updates, incorporating information about the state of the world and the actions each party has taken. The contract then identifies where in the vast and complex catalog of contingencies each party is. It then determines which obligations are applicable for that particular context. Third, the contract communicates the obligations to the parties, updating the information instantaneously.

III. Implications for the Language and Theory of Contracts

A. The Language of Contract Theory in a World of Self-driving Contracts

The emergence of the self-driving contract exposes a deficiency in the language of contract theory. To illustrate this deficiency, we show that self-driving contracts can be characterized as either of polar opposite categories under current terminology of contracting. The self-driving contract can be thought as both highly complete and highly incomplete. It is akin to an extremely detailed form contract while also sharing many characteristics with an agreement to agree. It provides *ex ante* certainty and highly specific instructions before the contracting parties perform, and at the same time fills gaps *ex post* through a mechanism similar to arbitration.

These views are seemingly contradictory. We find it helpful to categorize a self-driving contract as exhibiting aspects of each, but not fully described by any. The existence of this language deficiency is one we must address first, as a common language will be necessary in discussing other potential implications.

Commentators may view the self-driving contract as a highly incomplete contract.[44] Because parties to a contract find it hard to negotiate and write such a contract that accounts for all contingencies, many contracts are incomplete. This incompleteness manifests itself

---

43. We recognize that monitoring is a form of communication technology. But we separate the application of notice communication from monitoring communication for analytical purposes because they are targeted at different costs of contracting.

44. *See, e.g.*, Scholz, *supra* note 7 (discussing accountability of self-enforcing contracts).

in three ways: silence, ambiguity, or vagueness. Incomplete contracts create a particular set of problems for contracting parties.[45]

Silence is the most straightforward form of incompleteness. There are some events for which the contract does not specifies any "if X, then Y" statement. In that case the parties—either intentionally or unintentionally—have left the decision on those operations for determination at a later time. The silence may be interpreted as a gap to be filled or as a decision to embrace a status quo default.[46]

Alternatively, the contract might contain ambiguous terms. Again the ambiguity might be unintentional (for example, the parties may not be aware that there are two ships called *Peerless*)[47] or one or both parties might have fully intended the ambiguity.[48] Finally, the contract might use vague or general language like "reasonable" or "material" to invoke a standard that provides some guidance but not the exact operations for a specified contextual scenario.

Each of these forms of incompleteness requires content to be filled in after the fact. So how do micro-directives and self-driving contracts deal with incompleteness? When we discussed micro-directives above, we suggested that they can automatically fill gaps in the contract. Indeed, we suggested that the contract would be written at a very high level implying that it contains an enormous gap.

On that view, the self-driving contract is highly incomplete. It simply dictates a general objective without providing specifics. It leaves the details for later resolution by the computer program, which acts as a judge or arbitrator. Parties do not specify the details at the time of the contracting. Instead, the machine produces them just before performance. From the parties' perspective at signing, the contract is riddled with gaps when drafted. The contract then fills in those gaps, itself, as soon as they become important. The contracting parties do not directly contract over the substance that will fill the gaps but they do agree that the machine-driven algorithm will do the filling. If one views self-driving contracts this way, then the technology is acting like a powerful arbitrator that fills in all of the terms after the contract is entered into.

The opposing—but equally plausible—view would hold that a fully self-driving contract is perfectly complete. It is an algorithm that produces instructions on what each party is to do under every contingency. For every "if X" the contract produces a specific "then Y." The contract has no gaps. After all, each party has agreed and assented to be bound by the directive and neither can find an incompleteness that would allow them to pursue a strategy of opportunistic behavior. Put another way, the self-driving contract is more akin to a highly detailed form contract.[49] Like form contracts, there is the question of

---

45. At least one scholar has argued that there are benefits to incomplete contracts. *See* Wendy Netter Epstein, *Facilitating Incomplete Contracts*, 65 CASE W. RES. L. REV. 297, 321–22 (2014) (self-driving contracts can be viewed as preserving those benefits without the costs associated with opportunistic behavior that arise from the lack of commitment).

46. In formal terms, a status quo default is just one particular way of filling the gap. Russell Korobkin, *The Status Quo Bias and Contract Default Rules*, 83 CORNELL L. REV. 608, 609 (1998).

47. *See* Raffles v. Wichelhaus (1864) 159 Eng. Rep. 375, 2 Hurl. & C. 906 (showing an example of an unintentionally ambiguous contract).

48. *See* United Rentals, Inc. v. RAM Holdings, Inc., 937 A.2d 810, 813 (Del. Ch. 2007) (explaining the use of intended ambiguity in contracts through a reference to Greek mythology).

49. *See* Randy E. Barnett, *Consenting to Form Contracts*, 71 FORDHAM L. REV. 627, 634–36 (2002) (explaining that form contracts are about "manifesting consent to be legally bound" rather than making a promise based on something a party cannot be expected to have read).

whether a contract is complete if the parties have not read the terms.[50] This question is purely semantic and reveals the limitation of the incompleteness language. If the parties do not read rules in a contract (or a statute) but trust that they advance the general objective of the contract, they treat them as standards. But the courts will treat them like rules. Thus, a very long unread form contract is both complete and incomplete at the same time. Micro-directives just push that contradiction further.

A third way might be to think of self-driving contracts as contracts that use "*ex ante* arbitration." We recognize the internal conflict in that term. But that conflict captures our point. The contractual micro-directive performs arbitration services in the sense that—from the parties' perspective—there is a gap that is not filled until they bring their dispute to the machine after the contract is executed. But the contractual micro-directive is performing this function *ex ante* in two senses.

First, the outcome is predetermined at the time of execution. The machine can be said to have internally arbitrated all possible questions at the moment of contracting. The machine waits to be asked. Hence, *ex ante* arbitration. The parties agree to a broad objective, the details of which are arbitrated by a system whose outcomes are pre-determined but unknown to the parties at the moment of contract execution. The parties and the machine analytics have committed themselves to the principles that will decide the dispute. There are no subsequent negotiations or arguments about the bases for the decision,[51] but they do not know what specific rule will govern any particular context. As we have noted, it is as if the parties at the time of agreement placed all possible arbitration outcomes in sealed envelopes marked "open if X occurs." This commitment device goes a long way to prevent opportunistic behavior and to prevent the introduction bias.

Second, the answer is *ex ante* because it is provided at the time of performance rather than after performance as is usually the case with arbitration or litigation. This makes *ex ante* arbitration different from conventional arbitration. There is greater commitment and therefore greater certainty at the time of performance and there are no new transaction costs to seeking the answer.

In essence, the distinction between the negotiation of *ex ante* terms and *ex post* arbitration collapses. The idea that the machine analytics are resolving disputes *ex post* cannot be correct when the outcome is predetermined at the time of agreement and the answer is given before performance. On the other hand, the idea that the parties have *ex ante* agreed to the rule that governs their behavior is not correct either, as the parties do not have knowledge of any specific term.

A pricing example may provide clarity. Parties will no doubt use machine-learning algorithms to provide the optimal prices in some long-term rental agreements. The landlord and the tenant might agree to let the algorithm set the rent in a long-term lease after five years. But once the algorithm is good enough to determine the optimal rental price, why would the parties wait five years to use it? They could likely benefit by letting the algorithm set the price after one year rather than five. But why stop there? Why do the parties even need to negotiate the first year of rent? If the machines are accurate enough, the parties can save *ex ante* negotiation and *ex post* renegotiation costs by just letting the algorithm set the

---

50. *See id.*; *see also* Stephen J. Choi & G. Mitu Gulati, *Contract as Statute*, 104 MICH. L. REV. 1129 (2006) (examining the use of boilerplate contracts between sophisticated parties).

51. This is true even though the machines may be updating and learning.

price continuously from the start.[52] But it is very difficult to say precisely when the price term was set, when the parties agreed to it, and when the pricing in the contract became complete.

Why does the deficiency of language matter? From a practical point of view, the complete-or-incomplete question is important for how the law thinks about self-driving contracts. The two framings imply different analogies through which to analyze contracts. The doctrinal response will likely turn on how the contracts are viewed. For example, some might view the contracts as so incomplete or indefinite as to render them unenforceable on the grounds that the parties have failed to set out the essential terms of the agreement or that there was no meeting of the minds. Such an argument would be frivolous from the complete-contract view of things. The parties not only executed the contract but they agreed to an enormous catalog of provisions that covered millions of contingencies. On the completeness side, theories of renegotiation and efficient breach have to be rethought in a world where a contract is complete enough to reduce the instances of both while still providing some uncertainty about exactly how specific terms will play out. We examine these two examples more closely in the next two subsections.[53]

B. Implications for Contract Formation: Lack of Assent, Mistake, and Indefiniteness

Commentators and courts will, no doubt, attempt to pigeonhole self-driving contracts into existing frameworks of thought. When they do so, the language deficiency described above will lead to confusion. We illustrate this by examining the tension of trying to describe self-driving contracts within existing contracts doctrines. We start by exploring doctrines relating to the formation of a contract.

1. Lack of Mutual Assent

Courts refuse to enforce agreements where the parties do not have an objective meeting of the minds at the time of contracting. One may thus question if courts should enforce self-driving contracts where the parties do not know exactly what they are agreeing to at the time of formation. This skepticism may take two different forms. On one view, the self-driving contract is the ultimate long-form contract that no one reads. Every obligation is spelled out with incredible specificity. But no human could know exactly what they are agreeing to. They have not "assented" to the terms. Or one may view the contract as a blank sheet of paper that is filled in only after the parties have entered into the contract.

---

52. An arbiter taking on a role in setting *ex ante* contract terms is not fanciful. This takes place even in a world where contracts are written without the aid of machine learning technology. The American Institute of Architects, for example, has for years both operated as a means of arbitrating disputes between parties in the building industry and as a provider of standard form contracts to be used by contracting parties. *See* AIA, https://www.aiacontracts.org/ (last visited Oct. 31, 2017). Similarly, many existing contracts peg future pricing to indices which change independent of any action of the parties. As we discuss below, agreements that peg a price to an indices are just using an early from of self-driving contracts.

53. It is worth noting that self-driving contracts expose the deficiency of language in contract theory, they do not create it. Andrew Verstein has previously noted one such deficiency: that the *ex post/ex ante* distinction is misleading in all kinds of sophisticated contracts. *See* Verstein, *supra* note 29. Verstein's important insight was that parties often use vague terms and delegate the power to define those terms to non-court parties. In this way, the parties agree to what Verstein calls "*ex tempore*" contracts. Self-driving contracts take this even further because the parties can delegate that power to the non-human contract. This development reveals the important implications for contract law that we examine in the following sections.

It is nothing more than an indefinite agreement to agree.

One would be wise, however, not to overstate the role of assent in modern Anglo-American contract law.[54] Mutual assent has never required a complete understanding of—and assent to—every specific term in a contract by every party to a contract. If it did, the overwhelming majority of contracts would be unenforceable. Indeed, as long as there is a basic common understanding at the time of formation, mutual assent can be found. Take, for example, the case of *Hill v. Gateway 2000*.[55] A customer purchased a computer over the telephone. The contract for the product contained an arbitration clause. The customer was not informed of the arbitration clause at the time of purchase. Judge Easterbrook, writing for the Seventh Circuit, held that a "contract need not be read to be effective; people who accept take the risk that the unread terms may in retrospect prove unwelcome."[56]

Parties to self-driving contracts exhibit this sufficient level of common understanding and assent at the time of formation. Parties will agree to broad objectives and, importantly, they will agree to allow the algorithm to fill the gaps. This is not meaningfully different from agreeing to a contract that contains a standard, such as "material changes" or "reasonable efforts," to be adjudicated later. For this reason, as a general matter, lack of assent should not and likely will not be a reason courts give for failing to enforce self-driving contracts.

2. Mutual Mistake

Parties sometimes claim that they were mistaken when the contract was formed. They ask the court not to enforce these contracts. Self-driving contracts use technology to reduce the likelihood of mutual mistake. The contract updates the parties' obligations to account for factual changes that are discovered after contracting. Courts are already reluctant to strike down contracts for mutual mistake. A key takeaway here is that courts should be even less willing to interfere when relationships are governed by self-driving contracts.

Consider how self-driving contracts would deal with the mutual mistake and lack of assent in cases such as *Raffles v. Wichelhaus*.[57] In that case, the defendant buyer refused to accept delivery of cotton that left Bombay in December on a ship called "Peerless." The defendant argued that they thought they were buying cotton from Bombay that was arriving on a different ship called "Peerless." That ship had departed two months earlier. No contract was found to exist, as there was no meeting of the minds.

The facts of this case illustrate how predictive technology embedded in self-driving contracts will largely nullify the problems of mistake and lack of assent in *Raffles*. A self-driving contract identifies surplus maximizing paths for the two parties. The two parties are informed of how to behave in order to achieve this surplus maximization. An expected level of surplus can be communicated at the time of contracting. The specific ship upon which the cotton arrives should be largely irrelevant.

Similarly, self-driving contracts would have simplified matters in the case of

---

54.   RESTATEMENT (SECOND) OF CONTRACTS § 18 (AM. LAW INST. 1981); *see* Daniel A. Farber & John H. Matheson, *Beyond Promissory Estoppel: Contract Law and the "Invisible Handshake,"* 52 U. CHI. L. REV. 903, 932–34 (1985) (illustrating different methods of manifestation of assent).

55.   Hill v. Gateway 2000, Inc., 105 F.3d 1147 (7th Cir. 1997).

56.   *Id.* at 1148.

57.   Raffles v. Wichelhaus (1864) 159 Eng. Rep. 375; 2 Hurl. & C. 906.

Sherwood v. Walker.[58] There, the contract dealt with a sale of a cow and was not enforced on the grounds that the parties had been laboring under a mistaken belief that the cow was barren. Shortly after a price was agreed upon, but before the date of delivery, the cow was found to be with calf. The seller refused to deliver the cow. The court refused to enforce the contract on the grounds that there was a mutual mistake as to the subject matter of the contract.

If the parties had, however, elected to govern this transaction using a self-driving contract, the agreement would be enforced. The price of the cow may well have adjusted upwards upon revelation of the new information in line with the parties' agreement to maximize and share their joint surplus. If, with this new information, it will no longer be surplus maximizing for the trade to proceed, the contract would provide for a mechanism by which the seller kept the cow and compensated the buyer. But the mere fact that new information was revealed would not be enough for a court to refuse to enforce self-driving contracts.

In essence the doctrine becomes much less important because along this metric, self-driving contracts are far more complete and far more capable than courts of correcting for the parties *ex ante* mistakes about factual circumstances.

3. Indefiniteness and Agreements to Agree

To some, a self-driving contract may appear to run afoul of the indefiniteness doctrine.[59] But this, to us, is not the correct way to characterize a self-driving contract. Indeed, a self-driving contract can actually save agreements that would otherwise fail for indefiniteness. Vague terms like "fair share" are not actually vague when the parties also agree that an algorithm will enforce them. As a result, the law should be hesitant to invoke concerns of indefiniteness when dealing with self-driving contracts.

It is true that courts refuse to enforce agreements where the essential elements of a contract are missing[60] or where the contract fails to specify the contractual matter and a reasonable term cannot be implied.[61] If, for example, the proposed contract is merely an agreement to agree, courts may find no true agreement to enforce.[62] Imagine a contract without a price term. That contract may not be enforceable. But that is true only if the contract has no formulae or no mechanisms to determine a price at a later time.[63]

With self-driving contracts, the machine provides the mechanism to create definiteness. In the same way that contracting parties can use arbitration to fill the gaps on

---

58. Sherwood v. Walker, 33 N.W. 919 (Mich. 1887).

59. *See* Epstein, *supra* note 45, at 321–22 (discussing the role of indefiniteness in contract doctrine).

60. *See, e.g.*, Metro-Goldwyn-Mayer, Inc. v. Scheider, 360 N.E.2d 930 (N.Y. 1976) (holding that contracts including essential elements will be enforced).

61. *See, e.g.*, Stanton v. Dennis, 116 P. 650 (Wash. 1911) (holding that the law implies a reasonable time when one is not stated).

62. *See* Hunt v. Coker, 741 So.2d 1011, 1015 (Miss. Ct. App. 1999) (holding that a contract is unenforceable if its material terms are indefinite); Robert E. Scott, *A Theory of Self-Enforcing Indefinite Agreements*, 103 COLUM. L. REV. 1641, 1657–60 (2003) (discussing why parties still write indefinite agreements despite possessing knowledge of the indefiniteness doctrine).

63. *See* May & Butcher, Ltd. v. The King, [1934] 2 K.B. 17 (holding that no contract existed since an essential term of the contract was not determined); Hillas & Co. Ltd. v. Arcos, (1932) 147 L.T. 503 (holding that a contract existed when price was the only term to be determined); Foley v. Classique Coaches Ltd., [1934] 2 K.B. 1 (holding a contract existed when price can be settled at arbitration).

price, or peg the price to an external index, parties to a self-driving contract rely on the machine to provide the price at any point in time, depending on context. The machine either arbitrates a pricing standard or provides a complex pricing formula. Either way, a price is ultimately specified. And the same is true for other non-price terms as well.

Again, a canonical contracts case can illustrate. Consider how a self-driving contract would have resolved the dispute in *Varney v. Ditmars*.[64] In *Varney*, a plaintiff draftsman worked for a defendant architect during a period of commercial difficulty for the defendant. The plaintiff was promised a raise and "a fair share of the profits" to help the defendant out of his predicament. Over the next nine months, the plaintiff worked overtime, but was dismissed for extreme disloyalty and insubordination upon falling ill.[65] The plaintiff sued for "a fair and reasonable percentage of the net profits."[66] The defendant denied the existence of an agreement on the grounds that "fair share of his profits is vague, indefinite, and uncertain, and the amount cannot be computed from anything that was said by the parties or by reference to any document, paper, or other transaction."[67] The court agreed. The plaintiff's case failed on the grounds that no means of determining the net profits or determining the share that was agreed upon.

This case would proceed very differently if the parties' general relationship were governed by a self-driving contract. In a self-driving contract, the mechanism for determining both the fair share and the calculation for net profit could have been calculated based on the general objective of their relationship. The parties may not have specified anything precise, but if an algorithm was already managing their relationship, it could have told them what share of profits was appropriate. Indeed, the problem never even arises, because any self-driving employment contract would instruct the parties not only on the correct share of profits, but also on the correct amount and type of work.

All of this suggests that the rise of self-driving contracts will likely change what courts and commentators think of as an "essential" term in contracting. A broad objective—such as "maximize joint welfare" combined with a general notion of how to divide the surplus—may well provide sufficient certainty. While these broad objectives almost certainly do not suffice under any current construction of law, the change in the contracting environment where gaps can be filled on the fly will provide the requisite certainty and pre-commitment. As we have noted, this pre-commitment feature of self-driving contracts is a way in which they are more (not less) complete than traditional contracts.

C. Implications for the Time of Performance: Renegotiation, Breach, and Excuse

As we have suggested, self-driving contracts are complete in some ways but not others. On the completeness side, they have the characteristic of very strong *ex ante* commitment. This commitment changes the way in which contracts are performed. Indeed, as we shall see, the term "performance" itself may be misleading in a perfectly specified self-driving contract.

Currently, a poorly calibrated contract may require renegotiation or may incentivize inefficient breach once the state of the world is realized. But a well-functioning self-driving contract re-prices itself in a way to make it irrational for one party to renegotiate. Indeed,

---

64.   Varney v. Ditmars, 111 N.E. 822 (N.Y. 1916).
65.   *Id.* at 823.
66.   *Id.*
67.   *Id.*

renegotiation is a sign that the contractual micro-directives contained errors. Similarly, efficient breach is no longer an option. Because the price changes to maximize the joint welfare of the party, the price will *become* efficient. This means that any breach is inefficient. Thus, breach of a self-driving contract reveals either a market failure (liquidity constraints might be the most likely) or that the contractual micro-directives contained errors.

Our point about renegotiation and efficient breach can be illustrated with a simple numerical example. Imagine a manufacturer and a supplier enter a contract for the production of desks. The supplier can sell desks for $1,000 each. The manufacturer can produce desks for $500 each. A contract is formed with parties splitting the surplus. A price of $750 is initially specified. But then, due to no fault of the manufacturer, the cost of production increases markedly. Perhaps labor costs have unexpectedly increased or a large forest has been destroyed by wildfire. Each desk now costs $800 to manufacture. Under a contract that does not automatically update, the manufacturer may consider breaching the contract if the damages are likely less than $800 or they may need to engage in costly renegotiation. The self-driving contract, observing the increase in cost, updates and may re-price the desks at $900 to preserve the surplus split. There is no need for renegotiation to generate a surplus maximizing outcome.

If the production costs, however, increase beyond $1,000, the self-driving contract will no longer call for production of the desks. What was previously labeled "performance," (i.e., delivering desks) is inefficient in this state of the world. But the terms of the contract have adapted. The manufacturer conforms to the terms of the self-driving contract by *not* producing and delivering desks. Instead the manufacturer pays a non-production price that leaves the supplier no worse off.[68] For there to be an efficient breach, there would have to be a coding error where the contract does not update to take account of the price beyond $1000, or a situation where one party did not have the liquidity to pay the damages instructed by the contract.

The overall effect is similar to a contract with specific performance. The parties have committed to not challenge the efficiency of their relationship (through renegotiation or breach) and will simply perform. The difference is that specific performance in a traditional contract is blunt, whereas in a well-functioning self-driving contract, it adapts to the outcome that the parties would be expected to have renegotiated to in the absence of hold up.[69]

This has important implications for courts. Applying traditional contract doctrine, courts might be wary of enforcing a self-driving contract that is imposing an extreme price on a party who is seeking to get out of an obligation. Courts may view such a price as the equivalent of disfavored specific performance or of a penalty clause. But in the absence of market failures (asymmetric bargaining, manipulation of the computer code, liquidity constraints, or others), the courts should enforce them. To understand why, we now dig a little deeper into theories of renegotiation and efficient breach.

---

68. Of course, with these simple parameters, a traditional contract could cover this example. The self-driving contract, however, can cover thousands of examples and do so more precisely.

69. *See* Holden & Malani, *supra* note 19 (discussing the holdup problem in incomplete contracts); Aghion & Holden, *supra* note 19, at 194 (discussing firm management of incomplete contracts).

1. The Reduced Role of Renegotiation

Self-driving contracts will lead to a reduced scope for *ex post* renegotiation. Renegotiation, from a purely *ex post* view, can be viewed as a positive concept. Under the *ex post* view of renegotiation, if a contractual obligation is inefficient, the parties should change the terms. There is, however, a darker side to renegotiation. From an *ex ante* view, if the possibility of renegotiation is anticipated, renegotiation can temper the incentive for parties to invest in the relationship. If a party needs to make irreversible investments, it may be in the interests of the other party to try to opportunistically renegotiate the contract. If the value of the investments is relationship specific—that is, the investments only have value within the relationship—then the value of the contract can be renegotiated such that the investing party receives less of the surplus. This is called the "hold up problem." A rational party, foreseeing this eventuality, will have a reduced incentive to invest in the relationship.[70] Thus, when parties write an incomplete contract and are unable to commit not to renegotiate the contract, parties under-invest in the relationship.

The underinvestment problem is, however, reduced in a world of self-driving contracts. The key is the pre-commitment. With fewer gaps, reduced discretion, and no need to seek assistance from arbitrators, the ability to renegotiate terms is greatly reduced. The parties have agreed not to litigate a gap. That means they cannot take advantage of those gaps to opportunistically redistribute rents. The reticence introduced here toward renegotiation reinforces the lessons of *Alaska Packer Ass'n v. Domenico*.[71]

To be sure, the machine's arbitration of the gaps will redistribute value. But, because the course is unalterable by humans, there will be no efforts to divert the machine in the same way that there are efforts to divert a court. If the parties do renegotiate the contract, it is a signal that something went awry at the contracting stage. Perhaps a party could divert the machine, if it figured out how to "hack" the algorithms and engage in massive hold up attempts. It could then either directly divert value, or use the threat of diversion to create a negotiation position. The key for contract law then would be a doctrine that prevents manipulation or gaming of the algorithm. This argues in favor of laws that encourage algorithms provided by third parties who do not allow the contract parties access to the code.

Similarly, courts should not balk at enforcing terms that are substantively extreme (high price, for example). The price of the contract will move to make sure the efficient outcome among the parties is reached, and even if it is very high it is likely to be consistent with the parties' agreement. Courts instead should focus on procedure and ask whether

---

70.   *See* Steven Shavell, *The Design of Contracts and Remedies for Breach*, 99 Q.J. ECON 121 (1984) (discussing uncertainty implications for contract design); Sanford J. Grossman & Oliver Hart, *The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration*, 94 J. POL. ECON. 691 (1986) [hereinafter Grossman & Hart, *The Costs and Benefits of Ownership*] (discussing allocation of specific rights and residual rights); Paul A. Grout, *Investment and Wages in the Absence of Binding Contracts: A Nash Bargaining Approach*, 52 ECONOMETRICA 449 (1984) (examining bargaining outcomes in the absence of a binding contract); Patrick W. Schmitz, *The Hold-Up Problem and Incomplete Contracts: A Survey of Recent Topics in Contract Theory*, 53 BULL. ECON. RES. 1 (2001) (examining the hold-up problem when contracts are incomplete).

71.   Alaska Packer Ass'n v. Domenico, 117 F. 99 (9th Cir. 1902). In *Alaska Packer*, a group of sailors agreed to sail a ship from San Francisco to Alaska and to catch and can salmon while there, before returning it to San Francisco. When they got to the plant, they refused to further perform the contract unless plaintiff agreed to pay more money, which it eventually did. The court refused to honor the second agreement with additional compensation.

there is evidence that one party was taking advantage of coding errors or the like. The court might allow parties out of the contract only upon the party showing a high probability that the contract was programmed incorrectly or that the other party was hacking the analytics to produce a price that was artificially high or low.

It is worth noting that the pre-commitment that reduces the role of negotiation will have broad effects in other areas of law. Gaps and incompleteness in contracting have long been linked with the size of firms.[72] Under the property rights approach to contracting, when contracts are incomplete, the party who owns assets that are essential to performance is granted discretion to fill the gap.[73] This insight has been used to explain why firms would choose to vertically integrate rather than contract.[74] But if hold up and renegotiation are less common and parties can commit not to engage in such behavior, vertical integration becomes less important. The discretion afforded to asset owners is minimized. Under this approach, therefore, one might expect smaller firms and more contracting situations.

This is consistent with anecdotal observations about the gig economy and the rise of platforms like Uber and Airbnb that facilitate the freelance provision of goods and services. Indeed, Uber and Airbnb might be viewed as third-party providers of self-driving contracts. In many ways, Uber is providing a contract platform between the drivers and the passengers. The prices are determined algorithmically and update based on various on the ground circumstances. The result is that the vertical integration between driver, vehicle, dispatcher, pricing agent, is broken. The same can be said of Airbnb.

2. The Reduced Role of Breach

The economist's view of breach of contract differs from that of other legal scholars. To an economist, a breach is simply one party exercising an option to pay damages for failing to perform. In the words of Oliver Wendell Holmes, "[t]he duty to keep a contract at common law means a prediction that you must pay damages if you do not keep it,—and nothing else."[75] To others, this view is horrifying; contractual promises should be kept for moral reasons.[76] This view of breach and performance gains greater salience in a world of self-driving contracts—but only indirectly. The concept of efficient breach makes less sense when a contract offers a menu of obligations from which the parties can choose. The prices, however, are set in such a way to make the optimal path the most attractive path for the parties.

In essence, the efficient breach path becomes an explicit non-breach option.[77] On the one hand, this option makes the concept of not performing the initial obligation less morally objectionable; the option is specifically condoned by the contract. On the other hand, if parties opt for anything on the actual breaching path, that will be scrutinized extra closely

---

72. *See generally* Grossman & Hart, *The Costs and Benefits of Ownership*, *supra* note 70 (discussing how firms can allocate efficiently residual rights of control to prevent harmful effects of wrongful allocation).

73. *Id.*

74. *Id.*

75. Oliver Wendell Holmes, Jr., *The Path of the Law*, 10 HARV. L. REV. 457, 462 (1897).

76. *See generally* CHARLES FRIED, CONTRACT AS PROMISE: A THEORY OF CONTRACTUAL OBLIGATION (Harv. Univ. Press, 1981) (arguing that a contractual promise carries a moral obligation).

77. The embedded options become explicit. *See* Albert Choi & George Triantis, *Strategic Vagueness in Contract Design: The Case of Corporate Acquisitions*, 119 YALE L.J. 848, 853 (2010) (discussing the idea of embedded options in contracts); Scott, *supra* note 62, at 1643; George S. Geis, *An Embedded Options Theory of Indefinite Contracts*, 90 MINN. L. REV. 1664, 1669 (2006).

because it will either be inefficient or be evidence that the self-driving contract was not functioning in a manner consistent with welfare maximization.

Consider the effects of contractual micro-directives that already exist in the world. In the dental insurance case, the insured does not breach by failing to brush. The price simply changes. Similarly, with driving insurance, it is not a breach to drive the car at night or if traffic conditions become more perilous. The price simply changes. In soon-to-be-developed micro-directives, warranties can be written as pricing terms that update based on usage.

We should expect complex versions of this in the market for self-driving cars. Insurance will likely be a bundled product with the car where producers will be strictly liable for all accidents. That means they will want to alter the price based on how often and when the car is driven. This use fee will take into account the wear and tear on the car and the risk of accident from increased use. The result will be a car that is "sold" subject to a price that changes after the sale based on usage. This will look like a hybrid combination of sale, lease, warranty, and insurance. The user will make regular payments that change based on the relevant factors.

In farther-out hypotheticals, a Material Adverse Change clause in a merger agreement might simply be a contingent re-pricing. In this way, the contractual micro-directive just makes the embedded options within a contract explicit.[78] Now imagine that every term is a contractual micro-directive. In that world, every option has a price and the price is calculated to maximize joint welfare. When that is true, there is no such thing as efficient breach because all efficient paths are fully contemplated by the contract. Similarly, specific performance is no longer an issue. If brushing is the only efficient path, then the price for not brushing becomes infinite.

3. The Reduced Role of Excuse

The reduced role of breach brings with it an ancillary effect: the reduced incidence of excuse for breach. The likelihood of parties relying on doctrines of frustration or impossibility to excuse non-performance will no doubt also fall. Take the following famous example. In *Krell v. Henry*,[79] a man wished to observe the coronation parade of King Edward VII. He contracted to rent an apartment with an unobstructed view of the procession route for the afternoon. The rental price was far in excess of the usual price, reflecting the increase in demand generated by the coronation. But King Edward VII was old. He fell ill and the coronation was postponed. This contingency was not explicitly covered in the contract. The Court of Appeals held that the parties could not reasonably have anticipated the cancellation of the coronation at the time the contract was made. The illness of the King frustrated the purpose of the contract.

The self-driving contract, however, would have this eventuality already covered. While two humans to a contract may not be able to contemplate every contingency, the machine—basing its predictions of contingencies on vast amounts of data—would be able to identify and cover outcomes for each potential circumstance. In *Krell v. Henry*, the rental of the apartment is no longer surplus maximizing. Performing the prior obligations no

---

78. *See* Robert E. Scott & George G. Triantis, *Embedded Options and the Case Against Compensation in Contract Law*, 104 COLUM. L. REV. 1428, 1430–33 (2004) (describing the idea of embedded options); Choi & Triantis, *supra* note 77, at 851–60.

79. Krell v. Henry [1903] 2 K.B. 740 (Eng.).

longer makes sense. The losses will be divided in accordance with the surplus division rule.

4. Dealing with Idiosyncratic Preferences

One might suspect that self-driving contracts will be less able to factor in idiosyncratic preferences than contracts currently do. In a case such as *Jacob & Youngs v. Kent*,[80] for example, one wonders how a self-driving contract would deal with a homeowner's idiosyncratic preference for a particular type of pipe. In that case, a homeowner insisted upon strict conformance with the contract that had specified a pipe of "Reading manufacture." The house was built and the homeowner refused to pay an outstanding amount on the grounds that the terms of the contract had not been fulfilled. The builder had used a similar type of pipe manufactured by companies other than Reading.[81] Assuming the homeowner actually preferred Reading pipes—and only Reading pipes—would that idiosyncratic preference be factored in with a self-driving contract?

Most homeowners, of course, would not care if the pipes were of Reading manufacture or some other similar brand of equal quality. Thus, one would suspect that a program "unaware" of these idiosyncratic preferences would (like the actual court in *Jacobs & Young v. Kent*) substitute one brand for another if it were efficient. The idiosyncratic surplus destroyed by using a different type of pipe may be difficult to detect.

But this view is pessimistic. We suspect that self-driving contracts will provide an information-forcing role when it comes to idiosyncratic preferences. Homeowners with idiosyncratic preferences must be explicit about these preferences at the time of setting the objectives. If the preferences are genuinely idiosyncratic, they must be identified for the algorithms *ex ante*.

Indeed, one might imagine contracting parties having preferences and objectives that do not always result in pecuniary maximization. Take, for example, the case of *Attorney General v. Blake*.[82] The government of the United Kingdom sought an account of profits when a former British spy fled to the Soviet Union and (many years later) wrote a book, divulging information in breach of his employment contract. By the time of the book's release, all the information was public. Given that there is little cost to the government, any claim for expectation damages would likely be small. Using monetary gains as the metric, efficient breach theory, here, would permit Blake's breach. But the government had other non-monetary objectives.[83] To vindicate those objectives the court disgorged Blake's profits. For a self-driving contract to work in such a case, those objectives would have to be explicit in the algorithm. Otherwise it would deem Blake's action to be surplus maximizing.

This is not very different from how the law works today, but the uncertainty of courts makes it possible for parties to act opportunistically. A party might not reveal its idiosyncratic preferences *ex ante* but try to enforce them *ex post*. This results in the mispricing of the original contract. Or a party might fabricate idiosyncratic preferences *ex*

---

80. Jacob & Youngs, Inc. v. Kent, 230 N.Y. 239, 240 (N.Y. App. Div. 1921).

81. "The plaintiff tried to show that the brands installed . . . were, indeed, the same thing, though manufactured in another place . . . [T]he evidence, if admitted, would have supplied some basis for the inference that the defect was insignificant in its relation to the project." *Id.* at 241 (per Cardozo J.).

82. Attorney General v. Blake [2001] 1 AC 268.

83. "It is not a commercial claim in support of any commercial interest. It is a claim relating to past criminal conduct." *Id.* at 298 (per Lord Hobhouse of Woodborough).

post to extract value from the other side.[84] Because self-driving contracts take human courts out of the gap filling game, they reduce uncertainty, and therefore reduce the incentives for parties to play this game. Thus, they are a firmer information-forcing device than mere precedent.

5. Correcting Machine Errors

We have noted that breach or renegotiation might be a sign of some error in the functioning of the self-driving contract. Especially in early generations, we should expect mistakes. Mistakes are nothing new. Parties commonly make errors in drafting contracts. And they will likely make mistakes in creating self-driving contracts. If the parties admit to the mistake and renegotiate, the outcomes are simple. But parties will sometimes try to take advantage of the mistakes. The law will have to deal with those mistakes. Where we have a doctrine for scrivener's error today, the law of self-driving contracts will need a doctrine for coding error to fix obvious mistakes in coding that cause the contract to glitch.

More complicated will be instances of perverse instantiation.[85] A classic example in computer literature is that if you tell an artificially intelligent machine to make as many paper clips as possible and give it no other constraints it will attempt to turn everything in the universe into a paper clip.[86] Where parties misstate their objective, they will create micro-directives that give rise to perverse instantiation. Courts will have to deal with this problem. The likely solution is setting some high evidentiary bar for a party to prove perverse instantiation. We suspect this doctrine will be akin to the traditional contract law doctrines of mutual mistake and, perhaps, impossibility.

The above analysis is intended to move the theory of self-driving contracts forward. It also reveals how we should think about existing and proposed forms of contracting that are themselves partial or early versions of self-driving contracts. Thus, when parties agree to terms that are filled in by an index or by some other external source of information, we should worry primarily about whether that external source can be manipulated and whether the government should be involved in regulating those sources. Recent scandals about the manipulation of the Average Wholesale Price for pharmaceutical contracts[87] and London InterBank Offered Rate ("Libor") for financial contracts[88] come to mind.

Similarly, where some scholars have proposed using new external sources like surveys to fill gaps in contracts,[89] we should analyze these proposals through the lens of self-driving contracts. Survey results are similar to algorithms. If the parties to the contract have advance knowledge about how the survey will be implemented—what questions it will ask and of whom—they can run the surveys in advance and anticipate the results during the contract formation period. Imagine one party knows how a particular term

---

84. This might explain the homeowner's behavior in *Jacob & Youngs v. Kent*.

85. *See* NICK BOSTROM, SUPERINTELLIGENCE: PATHS, DANGERS, STRATEGIES 120–24 (2014) (describing perverse instantiation).

86. *See id.*

87. *See* Phuong D. Nguyen, *A Review of Average Wholesale Price Litigation and Comments on the Medicare Modernization Act*, 9 QUINNIPIAC HEALTH L. J. 249 (2006) (analyzing fraud concerns surrounding average wholesale price reports).

88. *See* Michael R. Koblenz et al., *Libor: Everything You Ever Wanted To Know but Were Afraid to Ask*, 6 J. BUS. ENTREPRENEURSHIP & L. 281 (2013) (analyzing the LIBOR scandal and its effect).

89. *See, e.g.*, Ben-Shahar & Strahilevitz, *supra* note 8, at 48 (showing an example of using survey data to fill the gaps in contracts).

performs in surveys, and the other party does not. This creates an asymmetry that can result in the original contract being mispriced. It is the equivalent of one party knowing how to outsmart the algorithm of a self-driving contract because they have access to the underlying code. Again, the government must determine whether to mandate or regulate third-party surveys and in general apply the analysis of the preceding section.

These questions about hacking algorithms and regulating providers are the new questions around which contract law must develop.

IV. PRACTICAL IMPLICATIONS OF SELF-DRIVING CONTRACTS: MARKET AND DEVELOPMENT

The emergence of second-generation self-driving contracts will likely coincide with the creation of a robust market for third-party contract providers and other forms of infrastructure to support the use of self-driving contracts. But these developments will not be uniform across all contracting. The infrastructure to support self-driving contracts will look very different for consumer transactions than it does for sophisticated business and finance transactions. Relatedly, there may be certain areas of the economy where the necessary infrastructure does not emerge. This will either hinder the development of self-driving contracts or allow for certain parties to take advantage of the situation to extract rents. In these circumstances, government involvement in the market may be required.

In this Part, we outline these likely developments. We focus on the question of *who provides* the self-driving contracts, noting where we think the markets will be deficient. We argue that the provider question is the most important procedural question for the law of self-driving contracts. We not only want to know what a micro-directive is, but how and where it was generated. In the same way that parties care about an arbitrator's neutrality and expertise when gaps are to be filled *ex post*, they must also focus on neutrality and expertise when the gap filling occurs *ex ante*. Only if there are issues of procedural irregularity should we concern ourselves with substantive irregularity.

There are three broad categories for potential providers of contractual micro-directives: (1) a party to the contract; (2) a private third-party; and (3) a government entity.[90] The implications for law are very different depending on which of these providers is involved.

A. Contractual Micro-directives Created by Parties to a Contract

Self-driving contracts may be written by the parties to the contract, but this generates additional bargaining costs. Consider how the parties would arrive at a self-driving contract. Parties may use machine-driven analytics to generate information about contract terms. They may collect data on previous deals to better understand how vague standards play out in practice, and use big data to better generate clauses that describe the obligations of the other party with far greater specificity than is currently permitted.

Imagine, for example, a party to a marketing agreement wants to assure optimal marketing efforts by the counterparty. The lawyers for the party could propose a vague

---

90.   We group these sources in a deliberately coarse fashion to better understand the benefits and challenges that self-driving contracts offer. But we do not expect that these groups, in reality, will be so neatly and cleanly defined. Indeed, a single contract may consist of several different self-driving terms that use micro-directives from various providers.

standard (for example, "best efforts") or a long list of rules specifying the marketer's obligations. Alternatively, the party could gather data on all similar transactions and use that data along with machine-driven analytics to produce the list of rules.

How will the lawyer from the other side respond? She will likely seek out similar analytics and produce her own list. Lawyers could then swap lists, identify inconsistencies and negotiate over them. This will work for a while. But as the list becomes more complete, they will get too long. What then? One party might propose combining the data to produce a contractual micro-directive that predicts the terms the parties would have negotiated for whenever a conflict arises. But who should provide the software? The lawyers do not trust each other and may feel the need to negotiate every line of code in the software, but that defeats the cost-saving purpose of using micro-directives. Over time, issues of trust might be overcome by reputation. If that happens and the parties are sophisticated, the law should be unconcerned that one party provides the self-driving contract. As a predictive matter, however, we expect that sophisticated parties will opt for third party vendors to solve this problem and that such vendors will emerge to fill that demand.

Things are more problematic in the universe of consumer transactions. We might expect the business party to provide the self-driving contract the same way they currently provide form contracts. But self-driving contracts exacerbate the potential for abuse that has been identified with form contracts.[91] A business might hide certain abusive terms inside the code. Or it might use its exclusive access to the underlying computer code to extract opportunistic gains. We suspect that market forces are unlikely to root out such hidden terms. The law should, therefore, be wary of self-driving contracts provided by the business party to a consumer contract—or any party to a contract in a relationship fraught by asymmetry.[92]

B. Contractual Micro-directives Created by Private Third Parties

As stated in the previous section, we suspect that third parties will be the primary source of contractual micro-directives. Markets for providers of self-driving terms will likely emerge. Private companies will specialize in providing self-driving contracts for contracting parties. They will provide computer code that takes account of their clients' joint goals and then produces directions on how these parties must behave in any given circumstances. The product is a contract in that the parties are bound by its terms, but the specifics of those terms are provided as the relationship progresses. In a well-functioning market, the private firms will compete over how well calibrated their ultimate terms are to the parties' objectives—and how hard it is for either party to game the system.[93]

---

91. *See generally* Xavier Gabaix & David I. Laibson, *Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets,* 121 Q. J. ECON. 505 (2006) (analyzing how firms can suppress information from consumers even in competitive markets); MARGARET JANE RADIN, BOILERPLATE: THE FINE PRINT, VANISHING RIGHTS, AND THE RULE OF LAW (2014) (examining how to better evaluate boilerplate provisions).

92. *See* Katz, *supra* note 19, at 183 (suggesting tactics to protect parties in asymmetric contracting).

93. This market for self-driving contracts is similar to, and takes inspiration from, recent work by Gillian Hadfield. She has suggested in several articles that we are moving toward a private provision of substantive contract law. In the same way that states like California and New York currently "compete" in the market for choice of law clauses, Hadfield suggests that private firms could fill that gap and compete for the right to arbitrate. Hadfield envisions a world where private firms adopt optimal default rules and contracting parties elect for one company's provisions over another. In this world, providers of contractual default rules will compete with each

The evolutionary paths toward such a market are numerous, but not all will produce well-functioning markets. First, new businesses might emerge to become providers. Second, private arbitrators who currently resolve conflicts *ex post* will shift their focus and marketing to *ex ante* protection. Third, insurance companies who currently gather and analyze data to provide risk allocation services will transform the use of that data to provide middleman services in dynamic contractual risk allocation, the result of which will be the provision of contractual micro-directives to manage the parties' relationship. Finally, in consumer transactions, consumer protection agencies might develop competing terms to offset abuses by businesses providing their own contractual micro-directives.

When these paths produce competent, neutral third-party providers, the law will have few substantive challenges. But that will not always be the case. To be sure, among sophisticated parties, the market will disfavor non-neutral providers. That said, there is still room for abuse. As with any provider of services, including arbitrators and judges, there will be a risk of fraud, bribery, and other forms of manipulation. Existing prohibitions on such behavior should therefore be enforced to regulate private third-party providers of self-driving contracts. The regulation should be light where markets are well functioning and parties can be expected to protect themselves. In asymmetric markets—such as those involving consumer transactions—more robust regulation might be necessary.[94] The debates here will be similar to those currently surrounding clauses requiring third-party arbitration in consumer contracts.[95]

Neutral third-party providers of *ex ante* terms are akin to neutral *ex post* arbitrators who fill gaps and arbitrate other disputes. In the same way that contract law insists upon neutral arbitrators to fill the gaps in the event of a dispute, we may very well insist upon neutral algorithm providers. Under such a scheme, in the same way that courts today defer to neutral arbitration, they might defer to self-driving contracts that come from neutral contract providers.

An aside on the evolution of *ex ante* arbitration is appropriate here. The most likely way for this market to emerge, and perhaps the most interesting, is through innovation by firms already in the arbitration business. Imagine a private firm that provides arbitration for contractual disputes. The firm settles many thousands of disputes each year. It has collected data from all these disputes and all the disputants to best determine the best course of action in many given situations. From here, it is a small step for the firm to create computer code that uses that data to provide cheaper and more accurate arbitration services. When a case comes up, the human arbitrator is provided with results from running the relevant data through the firm's proprietary algorithms. The results better inform the

---

other. We take a similar view, although we do not limit our analysis to the provision of *ex post* rules. GILLIAN K. HADFIELD, RULES FOR A FLAT WORLD 249–51 (2016) [hereinafter HADFIELD, RULES FOR A FLAT WORLD]; Hadfield, *Producing Law for Innovation*, *supra* note 9, at 47; *see also* Oliver Goodenough, *Digital Firm Formation*, *in* RULES FOR GROWTH 343–66 (2011) (outlining the process and results of digital firm formation).

    94.   Again, Gillian Hadfield's work on the regulation of private law providers by "super-regulators" is relevant and instructive here. *See* HADFIELD, RULES FOR A FLAT WORLD, *supra* note 93, at 267 (describing types of regulation already in place elsewhere).

    95.   *See, e.g.*, Michael S. Barr*, Mandatory Arbitration in Consumer Finance and Investor Contracts,* 11 N.Y.U. J.L. & BUS. 793 (2015) (explaining various issues with mandatory arbitration clauses within consumer contracts, including the inability for consumers to opt-out or seek collective redress); Brian T. Fitzpatrick, *The End of Class Actions?*, 57 ARIZ. L. REV. 161 (2015) (predicting that "decisions permit[ing] corporations to opt out of class action liability through the use of arbitration clauses . . . [will] eventually lead us down a road where class actions against businesses would be all but eliminated").

arbitrator of how to resolve the disputes fairly and consistently. Arbitrators soon see which algorithms work and which do not. As the algorithms become increasingly trusted by humans, the firm may bypass the human arbitrator altogether (and offer a large discount to the parties who agree to the bypass).[96]

At first such algorithms can be used to decide relatively homogeneous cases that frequently arise and resolve common issues. But as the technology improves and the data about arbitration increases, the firm can apply its technology to less common cases. Evolutionary algorithms will "learn" the types of outcomes that are of greatest benefit to the two parties. The algorithms can also weed out solutions that are not surplus maximizing. Over time, artificially intelligent learning algorithms might even be able to decide strange new cases that are not at all related to the data by predicting and mimicking what a human arbitrator would do in those situations.[97]

There is, however, no reason that parties or the arbitration provider would limit the application of this technology to *ex post* dispute resolution. Once the technology is available and its accuracy is demonstrated, the parties can simply agree to machine decisions from the outset. The arbitration firms will almost certainly repackage the analytic software to be sold to contracting parties *ex ante*, thus eliminating the need for *ex post* arbitration altogether. If there is a dispute, parties defer to the software. The software provides a solution that best fits the parties' aligned objective.

If the software provides this means to resolve disputes, the contract is already taking on features of a self-driving contract. Contingencies will be covered; parties will instantaneously be informed of their obligations. Over time, parties will delegate more and more key points of negotiation to the algorithms within the software. In this way, arbitration software—repackaged as contractual micro-directives—will not only replace *ex post* disputes over gaps, but also *ex ante* negotiation over specific terms.

C. Contractual Micro-directives Developed or Governed by the State

A third possibility is that government entities will provide contractual micro-directives for parties to use in their self-driving contracts. Government provision would arise as a response to the potential abuses discussed in the prior two sections or from a market failure that results in no available providers. We expect a robust debate over the comparative benefits of government oversight, on the one hand, and government provision, on the other. These questions will be especially salient for consumer contracting. In the absence of robust and truly neutral third-party providers, the calls for state restrictions on the use of self-driving contracts will be loud.

It is likely that the government will be better at overseeing or governing providers than it is in becoming a provider itself.[98] One might imagine government subsidies of firms

---

96. We do not doubt that the transition will not be immediate. Indeed, there is an emerging literature exploring adversity to algorithmic decision making, even after it has been proven that an algorithm outperforms human decision-making. *See* Berkeley J. Dietvorst et al., *Algorithm Aversion: People Erroneously Avoid Algorithms After Seeing Them Err*, 144 J. EXP. PSYCH. GEN. 114 (2014).

97. This ability to mimic human intuition is the heart of artificial intelligence and behind recent computing breakthroughs. *See* Choe Sang-Hun & John Markoff, *Master of Go Board Game Is Walloped by Google Computer Program*, N.Y. TIMES (Mar. 9, 2016), https://www.nytimes.com/2016/03/10/world/asia/google-alphago-lee-se-dol.html; *see also* Lewis-Kraus, *supra* note 36 (discussing how recent technological breakthroughs can mimic human intuition).

98. *See* HADFIELD, RULES FOR A FLAT WORLD, *supra* note 93, at 266–68 (describing "superregulation" as

that develop such regulatory powers given the likely public good element of the service. In the absence of government provision, a robust oversight and governance scheme could focus on the neutrality of the third-party provider, the opportunity for asymmetric exploitation of the algorithms, the likelihood that a party "hacked" the algorithm, and potentially the reasonableness of the directives produced.

D. How Should the Law Deal with the Provider Question?

The discussion above suggests that instead of examining the content and completeness of the contract, courts should focus on the *source* of the contractual micro-directives. The procedural mechanism by which parties receive the self-driving terms of the contract should be the subject of greater scrutiny, rather than the terms themselves.

If one were concerned that an arbitrator were biased, a court would be more likely to hold an arbitration clause in a contract unconscionable.[99] Similar principles apply here. If we are concerned that the *ex ante* arbitration unduly favors one party, then we should question whether the contract should be enforced.

Concerns of exploitation are particularly prominent in the consumer context. Insurance companies, for example, may use their massive collections of data to refine the standards in an insurance contract. Insurance companies armed with greater information about their customers' health, status, and behavior will better target risks of exposure. Based on years of experience and using advanced techniques to analyze the data, the insurance company may be able to reduce their legal uncertainty under a contract by specifying what is covered under the contract with far greater specificity.

Here, just one of the parties—the business developing the consumer contract—will likely have access to data. Just one side will develop the algorithm. And just one side will understand the legal obligations and rights under the contract. The complexity of the algorithm will make it impossible for any individual consumer to understand their contractual obligations and rights.

This additional complexity of self-driving contracts exacerbates a concern that scholars have long held regarding consumer contracts. Consumers rarely read long contracts with fine print. Even when important clauses that drastically affect legal rights— such as arbitration clauses—are highlighted with colored, capital, and bold letters, consumers neglect to read the terms, let alone absorb the importance of such terms. As companies add more and more contingencies, the contracts will begin to resemble self-driving contracts. The likelihood of consumers being able to read and understand all terms at the time of contracting will converge to zero. This is essentially a magnified problem of boilerplate.

Such an environment offers enormous opportunities for exploitation. It affords the business the opportunity to include terms that are favorable to the business, but impossible for consumers to identify at the time of contracting. The terms are hidden or shrouded at

---

a system in which private firms would regulate business in a competitive framework while government oversees the framework and approves the private firms who may regulate).

99. *See* David Horton, *Unconscionability Wars*, 106 Nw. U. L. Rev. 387, 396 (2012) (criticizing argument against court intervention on arbitration provisions as preserving arbitrator bias in favor of one party); *see generally* Anthony Niblett, *Tracking Inconsistent Judicial Behavior*, 34 Int'l Rev. L. & Econ. 9 (2013) (explaining results of study showing bias of California judges in ruling on unconscionability of arbitration provisions).

the time of contracting. The consumer may only become aware of these terms once the contingency is revealed and the consumer's additional obligations are triggered, or the consumer's rights have evaporated. Consumers' choice is vastly reduced when consumers are not made aware of the exploitative terms, and competition between businesses has proven to be largely ineffective at sifting them out.[100]

Jurisprudence in contract law and consumer law has developed a number of different tools to combat problems of exploitation. The self-driving contract may blunt some of these instruments. Take, for example, how the doctrine of *contra proferentem* will be affected by the advent of self-driving terms. *Contra proferentem* is a doctrine whereby courts interpret any ambiguity in a manner that is most favorable to the party that did not draft the contract. This doctrine encourages parties who draft contracts to be as clear as possible when writing contracts. Vague, ambiguous, or undefined terms will be viewed in a light that is unfavorable to a contracting party. But in a world where contractual terms are clear, explicit, and leave nothing open for renegotiation or interpretation, the likelihood of a term requiring interpretation is reduced.

Courts may respond with the blunt and perhaps ineffectual rule of only allowing businesses to change the price in a self-driving contract if it is favorable to the consumer. For example, a court would enforce a self-driving dental insurance contract if premiums dropped when the insured brushes her teeth; but courts may not enforce the contract if premiums rose when the insured failed to brush her teeth. A rational economist would view this policy as toothless. This policy would simply result in higher base prices that fall as the insured takes more care. The effect would be neutral. A behavioral economist, on the other hand, may take a different view to this policy. A higher baseline price may present sticker price shocks for consumers, forcing firms to set a lower baseline price.[101] While this would reduce the ability of the firm to exploit consumers, it could also nullify the beneficial effects of self-driving contracts.

A further concern in consumer contracting is privacy. One might object that monitoring technologies are too invasive.[102] By storing information about one's every move, contracting parties are eroding their own privacy. We are less concerned with the use of monitoring technology here than we are with the use of monitoring technology by the state.[103] Entering into such contracts is voluntary. For example, an insured party may choose to use an electric toothbrush with a monitoring device, but they do so with the understanding that they will be rewarded with lower premiums when they brush their teeth. They can, of course, elect not to have to a monitoring toothbrush. They would not, though,

---

100.   Gabaix & Laibson, *supra* note 91, at 509; Florencia Marotta-Wurgler, *Competition and the Quality of Standard Form Contracts: The Case of Software License Agreements*, 5 J. EMPIRICAL LEGAL STUD. 447, 448 (2008).

101.   The "default effect" is a consequence of status quo bias. *See* Daniel Kahneman et al., *Anomalies: The Endowment Effect, Loss Aversion, and Status Quo Bias*, 5 J. ECON. PERSP. 193, 199 (1991); Richard H. Thaler & Shlomo Benartzi, *Save More Tomorrow: Using Behavioral Economics to Increase Employee Saving*, 112 J. POL. ECON. 164, 168 (2004).

102.   *See generally* Laura Evans, *Monitoring Technology in the American Workplace: Would Adopting English Privacy Standards Better Balance Employee Privacy and Productivity?*, 95 CAL. L. REV. 1115 (2007); Waseem Karim, *The Privacy Implications of Personal Locators: Why You Should Think Twice Before Voluntarily Availing Yourself to GPS Monitoring*, 14 WASH. U. J. L. & POL'Y 485 (2004) (showing objections to monitoring technology that is too invasive).

103.   *See* Casey & Niblett, *Death of Rules and Standards*, *supra* note 1, at 1410–12 (discussing "contractual micro-directives").

be entitled to the rewards. Thus, failing to exercise the option of using such technology signals either a strong desire for privacy or it signals that the insured party engages in destructive behavior. With sufficient information, the value of one's privacy can be adequately priced into any contract.[104]

V. CONCLUSION

Self-driving contracts will, no doubt, be greeted with a healthy mixture of skepticism, trepidation, and fear. For some, the fact that obligations will be "determined" by a predictive algorithm presents cause for concern. The removal of human discretion from the contracting process may seem antithetical to the idea that economic actors voluntarily choose what obligations to undertake. Further, the idea that parties to the contract do not know exactly what obligations they are undertaking may lead some to conclude that the assent element of contracting can no longer be met. These are alarmist visions that we do not share.

The cause for concern will be limited to certain contexts. With two sophisticated parties and self-driving terms provided by a neutral third party, the introduction of contractual micro-directives will, for the most part, introduce few hurdles to contract doctrine. Contract law in the Anglo-American tradition is robust. New questions and legal issues can be dealt with in the existing framework of contract law. The law contains many jurisprudential safeguards, already built-in, that will prevent many of the perceived deleterious effects of contractual micro-directives from ever developing. A potential red flag should be raised in the area of consumer law, where businesses may develop their own self-driving contracts. The shrouded nature of the obligations at the time of contracting may yield greater need for state regulation where businesses provide their own terms.

But the development of self-driving contracts presents an exciting opportunity for scholars of contract law. Some questions that have been vociferously debated for decades may suddenly become obsolete; other questions, that have yet to be dreamed up, will steal the limelight. The language of contract law will be forced to change with the ever-changing technology. The dramatic reduction in bargaining costs opens the door to a world that looks more Coasian by the day; a world where beneficial private co-operation is easier to stumble upon and surplus can more easily be maximized. Scholars of the future may look back in wonder, questioning how contracting parties managed to negotiate and write clumsy human-created contracts at all before the rise of the self-driving contract.

---

104. *See* Ariel Porat & Lior Jacob Strahilevitz, *Personalizing Default Rules and Disclosure with Big Data*, 112 MICH. L. REV. 1417, 1467–70 (2014) (discussing the tradeoff in personalization and privacy in default rules); Lior Jacob Strahilevitz, *Toward a Positive Theory of Privacy Law*, 126 HARV. L. REV. 2010, 2021–34 (2013) (discussing challenges when dealing with big data and privacy). There are, of course, distributive questions that are raised and should be considered when things like privacy are traded on markets. *See, e.g.*, Craig Konnoth, *Health Information Equity*, 165 U. PA. L. REV. 1317, 1334–35 (2017) (discussing burdens that arise when dealing with privacy, security, and health records).