

Forthcoming in HARVARD JOURNAL OF LAW & TECHNOLOGY

Version submitted to conference (3.27.2026)

ZOMBIES, AI, AND THE “OBJECTIVE” THEORY OF CONTRACTS

James Toomey*

Abstract

“Agentic” “AI” that can pass the “Turing Test” might behave in ways externally indistinguishable from persons but without our subjectivity or mental states. Under the so-called “objective” theory of contracts, taught as black letter in the common law world for a century, this is purportedly enough for large language models to enter into self-binding contracts. In contrast, this Chapter argues that this possibility offers a reductio in favor the alternative view that takes the “objective” theory as the law’s evidentiary commitment rather than a normative or metaphysical one. For starters, it is not clear that language “agreed upon” by entities without mental states means anything, at least not anything anyone cares about. More generally, there is no normative reason for the law to attempt to enforce the output of probabilistic language models entirely disconnected from subjective intent; the only possible justifications of contract are, in the end, “subjectivist.”

* Associate Professor of Law, University of Iowa College of Law.

INTRODUCTION

“[T]he making of a contract,” Holmes tells us, “depends not on the agreement of two minds in one intention, but on the agreement of two sets of external signs.”¹ “The mental assent of the parties is not requisite for the formation of a contract.”² “A contract has, strictly speaking, nothing to do with the personal, or individual, intent of the parties.”³ These are the shibboleths of the so-called “objective theory of contracts,” taught as black-letter in American law schools, which holds that for all its talk of “autonomy,” “agreements,” “promises,” “meeting of the minds,” and so forth, in fact all the law cares about is whether the parties have gone through motions that, taken “objectively,” appear to “agree” to some content.⁴

If true, the “objective theory” would apparently entail that large language models (LLMs) can enter into binding contracts with one another. Clearly, the output of a contemporary language model, probabilistically stringing together words from human-created corpora, could, read “objectively,” look like an offer—“I can sell you 1,000 widgets for \$1,000,” say. Fed into the input of another model, we could easily get an output that looks, “objectively,” like acceptance—“Sure, send them over!” The objective theory of contracts would have it that this is an enforceable contract—and some scholars appear to have no problem enforcing it.⁵ In contrast, I think a legal rule that would purport to “enforce” this “contract” against entities mindlessly stringing together word sequences in imitation of human language use—philosophical zombies, perhaps as yet without bodies—is rather profoundly out to lunch.⁶

There is an alternative. Lawrence Solan, among others, has argued that—for all the anyway discredited high-strung mid-century behaviorism of its rhetoric—the “objective theory” of contracts is not a committed normative position but rather a straightforward evidentiary rule designed to *better* pursue the *genuine, subjective* agreement of the parties; a shield against the possibility that they may later lie about their mental states at the time of

¹ Oliver Wendell Holmes, *The Path of the Law*, 10 HARV. L. REV. 457, 464 (1897).

² LUCY V. ZEHMER, 84 S.E. 2d 516, 522 (Va. 1954).

³ HOTCHKISS V. NATIONAL CITY BANK OF N.Y., 200 F. 287 (S.D.N.Y. 1911).

⁴ See, e.g., 1 Williston on Contracts § 3:7 (4th ed.).

⁵ SHAWN BAYERN, AUTONOMOUS ORGANIZATIONS 158 (2021).

⁶ To be sure, the concept of zombies in philosophy is usually meant to describe entities that *look* and act identically to humans but without mental states, whereas today’s LLMs don’t have bodies and at best *talk* like humans. But surely you could rig the controls of a robot identical to a human body to an LLM, and the point is the same either way. See, e.g., Robert Kirk, *Zombies*, in EDWARD N. ZALTA & URI NODELMAN, EDS., STAN. ENCYC. PHIL. (2023).

agreement.⁷ After all, without access of our own, it might be very hard to prove that they *are* lying about what their mental states in fact were.⁸

This paper argues that the possibility of LLMs going through the motions of agreement proves that Solan's is the better view of the objective theory of contracts. The law's purporting to enforce probabilistically generated linguistic sequences wholly untethered from *anyone's* mental states is an absurdity—no matter how much the linguistic sequences look like those that human beings with mental states use to enter contracts that the law ought to enforce. For one thing, it's not clear that the outputs of LLMs in this way mean *anything* on their own terms. And even if they do, whatever they mean is certainly not the kind of thing any reasonable legal system has normative reason to burn taxpayer dollars “enforcing.”

To be clear, I do not deny that LLMs can play a useful role in contemporary contract drafting—if there's one thing these models are good at, it's generating boilerplate.⁹ Executive A can agree with Executive B to bind themselves to the language generated by LLM C (presumably, for interpretive purposes, as attributed to some kind of fictional “reasonable” speaker)¹⁰—the law can certainly enforce these sorts of agreements.¹¹ But this is different from the cases I'm talking about—there *are* actual mental states involved here; two entities with mental states (humans) actually agreed on *something*. The cases in which legal enforcement is absurd are those where two LLMs *wholly disconnected* from human mental states (except insofar as they were trained on human-generated corpora) appear to “agree” with one another—technically possible but, as far as I know, *not* a major component of the contemporary economy.

The analysis proceeds in three Parts. First, I briefly summarize the positions of the “objective theory” of contracts and its discontents. Next I argue that to the extent the prevailing understanding of the rule entails enforcement of apparent agreements reached by language models “talking” to one another wholly disconnected from anyone's mental states, it would be quite literally absurd. Finally, I close with some reflections on the

⁷ See, e.g., Lawrence M. Solan, *Contract as Agreement*, 83 NOTRE DAME L. REV. 353, 379 (2007); see also Andrew S. Gold, *A Property Theory of Contract*, 103 NW. U. L. REV. 1, 34 (2009); CHARLES FRIED, *CONTRACT AS PROMISE: A THEORY OF CONTRACTUAL OBLIGATION* 61 (2d ed. 2015); Hanoch Sheinman, *Introduction: Promises and Agreements*, in HANOCH SHEINMAN, ED., *PROMISES AND AGREEMENTS: PHILOSOPHICAL ESSAYS* 3, 32 (2011); Peter Meijes Tiersma, *The Language of Offer and Acceptance: Speech Acts and the Question of Intent*, 74 CALIF. L. REV. 189, 223 (1986).

⁸ See *id.*

⁹ See, e.g., S. Sean Tu, Amy Cyphert & Samuel J. Perl, *Artificial Intelligence: Legal Reasoning, Legal Research and Legal Writing*, 25 MINN. J. L. SCI. & TECH. 105, 125 (2024).

¹⁰ See *infra* ____.

¹¹ Bayern, *supra* note [BAYERN2022], at 23.

implications of this view for contract theory and the role of LLMs in contemporary business transactions.

I. THE OBJECTIVE THEORY AND ITS DISCONTENTS

Before you went to law school, you might have assumed that contract law—that paradigm of voluntary agreement—probably has something to do with enforcing promises or agreements; intentional, voluntary commitments that people choose to enter into, or not.¹² And, of course, much of the language and structure of contract doctrine is consistent with this intuition. Courts routinely use the language of “promise” and “agreement” in contract litigation.¹³ To this day, though many commentators find it confusing, courts speak of seeking a “meeting of the minds.”¹⁴ And it is contemporary black letter that the content of a contract is the meaning attributed to it by both parties.¹⁵ Meanwhile, many contract doctrines—from duress to mistake to undue influence—seem carefully crafted to screen out apparent agreements that were not meaningfully voluntary; those in which, despite superficial appearances, there was no genuine agreement.¹⁶

Wrong, all wrong—law professors, courts, casebooks, and treatises will eagerly tell you.¹⁷ Sure, they’ll concede, contract law *looks* in many ways like it cares about mental states and the intentions of actual human beings. But that’s because what it *does* care about just is *looking like* it cares about mental states and intentions—contract, on this view, is concerned only with the *appearance* of agreement, not anything so “soft” and “mysterious” as whether anyone actually agreed to anything. This is the “objective theory of contracts,” and it purports to be the black-letter law across much of the common law world—the claim is that contract formation and interpretation have nothing, strictly speaking, to do with what was going through the minds of the parties, but are the sole product of what an objective, reasonable third party would understand to have taken place given the parties’ external

¹² See, e.g., CHARLES FRIED, *CONTRACT AS PROMISE: A THEORY OF CONTRACTUAL OBLIGATIONS* (2d ed. 2015); Seana Valentine Shffrin, *Promising, Intimate Relationships, and Conventionalism*, 117 PHIL. REV. 481 (2008).

¹³ See, e.g., Seana Valentine Shffrin, *The Divergence of Contract and Promise*, 120 HARV. L. REV. 708, 709 (2007); Jed Lewinsohn, *Paid on Both Sides: Quid Pro Quo Exchange and the Doctrine of Consideration*, 129 YALE L. J. 690, 693 (2020).

¹⁴ See, e.g., *TURNER-BASS ASSOCIATES OF TYLER V. WILLIAMSON*, 932 S.W. 2D 219, 222 (Tex. Ct. App. 1996); *LUNAS COOPERATIVA DE SEGUROS MULTIPLES DE PUERTO RICO*, 100 So. 3d 239, 241 (Fla. Ct. App. 2012).

¹⁵ Restatement (Second) of Contracts § 201(1).

¹⁶ Restatement (Second) of Contracts §§ 174–177.

¹⁷ See, e.g., ROBERT A. HILLMAN, *PRINCIPLES OF CONTRACT LAW* 39-40 (2004); 1 Williston on Contracts § 4.1; E. ALLAN FARNSWORTH, *CONTRACTS* § 3.6 (4th ed. 2004).

conduct.¹⁸ Contracts are, on this view, formed if the parties *looked like* they agreed; their language and conduct interpreted in some “objective” sense.¹⁹

From this perspective, the “objective theory” is thought to be positioned against a “subjective” alternative—in which contract formation would be governed by whether the parties *actually* agreed; in which the meaning of the agreement is constituted by what *they* meant.²⁰ The story goes that the hard-headed, pragmatist objective theory triumphed over the mushy subjective alternative in American law in the last century, while it still governs French law and the civil law countries that follow them.²¹ It is surely no coincidence, from this perspective, that the triumph of the “objective theory’s” effort to purge mental states from legal analysis roughly coincided with the apex of behaviorism in psychology and distributionalism in linguistics in the American academy.²²

Notwithstanding its ubiquity, however, the objective theory of contracts has always been controversial—both descriptively and normatively. Arthur Corbin, for one, doubted its coherence on the grounds that the “meaning” of agreements can only be what people actually meant by their language.²³ And although both Restatements are often thought to embrace the theory, they are, in fact, all over the place on the question—they tell us both that contracts are founded on mutual assent and to be interpreted according to the meaning of the parties,²⁴ and that by “assent” is meant “what looks like assent.”²⁵

Lawrence Solan, most prominently, sought to recast the “objective theory” of contract not as a normative commitment, but as an evidentiary rule designed to efficiently enforce what is ultimately, normatively, the aspiration of contract law to enforce genuine subjective agreements.²⁶ On this view, the law fails at its own ultimate goal—enforcing genuine, subjective agreements—where it defers to one party’s later fraudulent claim that they

¹⁸ See *supra* notes ____; see also Melvin Aron Eisenberg, *The Emergence of Dynamic Contract Law*, 88 CAL. L. REV. 1743, 1757 (2000); Todd Rakoff, *Washington v. Davis and the Objective Theory of Contracts*, 219 HARV. C.R.-C.L. L. REV. 63, 78-79 (1994).

¹⁹ Restatement (Second) of Contracts § 21.

²⁰ French Civil Code § 1110.

²¹ See, e.g., Joseph M. Perillo, *The Origins of the Objective Theory of Contract Formation and Interpretation*, 69 FORDHAM L. REV. 427, 428-29 (2000); Wayne Barnes, *The French Subjective Theory of Contract: Separating Rhetoric From Reality*, 83 TUL. L. REV. 359 (2008).

²² See, e.g., George Graham, *Behaviorism*, in EDWARD N. ZALTA & URI NODELMAN, EDS., *STAN. ENCYC. PHIL.* (2023); Marietta Auer, *Can Kantians Be Realists?*, 70 AM. J. JURISPRUDENCE (2025).

²³ 1 Arthur Linton Corbin, *Corbin on Contracts* § 4.12 at 628.

²⁴ Restatement (Second) of Contracts § 18; 201(1).

²⁵ Restatement (Second) of Contracts § 2(1).

²⁶ Solan, *supra* note [SOLAN2007], at 356.

meant something other than what they did at the time.²⁷ Given individuals' own privileged access to their prior mental states, it may be very hard to individually adjudicate whether they *are* now lying—and, all else equal and absent evidence to the contrary, probably they *did* mean what a reasonable person would have understood them to mean.²⁸ Moreover, Solan's view notes that if someone was indeed lying at the time of the agreement, there may be good normative reasons to hold them to its content (perhaps sounding more, in that case, in promissory estoppel than contract proper, but either way justifiable).²⁹

Taken together, these considerations can justify something like the "objective theory" of contracts as a legal rule not grounded in a normative view of what contracts are, but as an effort to enforce what the parties actually, subjectively meant, given the risk that they might now be lying, or one might have been engaged in fraud now or then. This view, in other words, recasts the objective theory of contracts not as a theory of *contract*, but as a rule designed to efficiently enforce what contracts actually are; genuine, subjective agreements.

From this perspective, it is illustrative that one of the paradigmatic cases of the "objective theory," *Lucy v. Zehmer*, does not really show that the law is not concerned with mental states, for all its language to that effect (it's a very long opinion—grist for many mills).³⁰ In that case, Zehmer purported to sell his farm to Lucy at a bar—after extensive negotiating that culminated in writing and signing an instrument of sale; without any indication at the time that the transaction was anything but serious.³¹ In court, Zehmer claims he had been joking the whole time.³² The court said it was relying on the objective theory of contracts to dismiss Zehmer's defense on the grounds that it didn't matter if he had been joking.³³

But surely it didn't have to. If Zehmer *wasn't* joking, letting him get out of an agreement on the claim that he had been would be inconsistent with the aspirations of a *subjective* theory of contracts—there would have been an actual agreement the existence of which one party is *fraudulently* denying. And, indeed, there *was no evidence*, other than Zehmer's own self-interested testimony, that he had been joking—probably he *wasn't*.³⁴ At least, the court explicitly thought as much—"If it be assumed, *contrary to what we think the*

²⁷ *Id.* at 366.

²⁸ *Id.*

²⁹ *Id.*

³⁰ 84 S.E. 2d 516.

³¹ *Id.*, at 518.

³² *Id.*, at 521.

³³ *Id.* at 522.

³⁴ *Id.* at 521.

evidence shows, that Zehmer was jesting about selling his farm”³⁵ Zehmer’s defense, in other words, fails on a “subjective” theory of contracts too, and French law (obviously?) doesn’t let people get out of their genuine commitments based on unsubstantiated, self-interested post-facto claims that they had been secretly joking the whole time.³⁶ It’s hard to imagine an efficacious contract regime that would recognize this defense, whether purporting to be “objective” or “subjective.”

This, then, is the debate against which the possibility of LLM-formed apparent agreements entirely independent of anyone’s mental states takes place—that between the prevailing Anglo-American “objective” theory of contract as against a “subjective” one, coupled with the possibility that what passes for the “objective” theory is really an evidentiary posture designed better to enforce a “subjective” one.

II. THE ABSURDITY OF “AGREEMENT” BETWEEN LLMs

LLM A outputs the following linguistic string into the input of LLM B—“I can do 1,000 widgets for \$1,000.” LLM B outputs—“You got a deal!” Take that transcript to a reasonable third party, and they’ll say it looks like these two entities manifested assent—the objective theory tells us we have a contract that the courts ought to spend taxpayer-funded salaries enforcing, by the state’s monopoly on violence if necessary.³⁷ This is absurd.

For starters, it is not clear that this “agreement” in fact *means anything*—and it is obviously impossible for the law to enforce meaningless content.³⁸ Indeed, a long-running position in the philosophy of language holds that the “meaning” of a linguistic utterance *just is* the meaning that the speaker is intentionally trying to communicate, and that language has *no* meaning outside of its actual use by speakers to convey intended meanings.³⁹ On this

³⁵ *Id.* at 521.

³⁶ *See, e.g.*, Barnes, *supra* note [BARNES2008].

³⁷ You may want to resist this conclusion on the ground that it is not reasonable to interpret this language *coming from LLMs* as manifesting assent. That may be—but imagine, then, that these LLMs have been given control of bodies superficially identical to human ones—true philosophical zombies. Clearly this same exchange would form a contract under the objective theory as articulated. (And contracts can clearly be formed by anonymous electronic communication, why is it not reasonable to read this transcript as “manifesting assent”?) If you’re still inclined to resist the conclusion because the existence of philosophical zombies would require reasonable people to reinterpret certain kinds of expressions of assent as no longer good proxies for their mental states—that is a subjective theory of contract law.

³⁸ Restatement (Second) of Contracts § 200.

³⁹ *See, e.g.*, Richard Ekins, *Objects of Interpretation*, 32 CONST. COMMENT. 1, 3 (2017); Noël Carroll, *Art, Intention, and Conversation*, in GARY ISEMINER ED., INTENTION & INTERPRETATION 97, 120 (1995); Larry Alexander & Saikrishna Prakash, *Is that English*

view, the outputs of LLMs have literally no meaning—making efforts to enforce them incoherent.

Contrast enforcing LLM outputs, from this perspective, with enforcing the language of ordinary agreements between human beings. The parties in those cases *did* mean something in the language they used—and courts understand their task to be to discover what it was they meant (indeed, notwithstanding the language of the “objective theory” of contracts, courts have *always* tried to figure out what *the parties* meant by the words they chose, and the Restatement tells us that meaning prevails).⁴⁰ At the outset, in other words, courts *could not possibly* enforce the language output by LLMs in the same way they do language agreed-upon by human beings—they do the latter by trying to make sense of what the human beings who used the language meant; a question to which there is, necessarily, no answer in the LLM case.

To be sure, there are other things courts could do. The view that language has *no* meaning outside of its use is a minority position in contemporary philosophy and linguistics—most theorists believe that well-formed linguistic strings have so-called *literal* meanings, as an abstract matter of the semantic and syntactic rules of a particular language—regardless of how they might be used in particular conversational contexts to convey particular intended meanings.⁴¹ Interpreting the language of this exemplary agreement literally has problems of its own, as we’ll see—but it is worth noting that *everyone* in the debate over whether literal meanings exist agrees that intent plays a necessary role in establishing whether a *communicative act* has any meaning at all.⁴² That is, even the most formalistic defenders of literal meaning agree that if the shape of the clouds, or the sounds of the wind, happen to superficially resemble the English sentence “If you build it, he will come,” that wouldn’t *mean* anything.⁴³

And this is exactly what LLMs *are*—they are probabilistic outputs of linguistic strings without intents, the same as cloud shapes, wind noises, and

You’re Speaking—Why Intention Free Interpretation is an Impossibility, 41 SAN DIEGO L. REV. 967, 977 (2004); see also FRANCOIS RECANATI, LITERAL MEANING 83 (2004).

⁴⁰ See, e.g., James Toomey, *Privately-Crafted Obligations*, 78 FLA. L. REV. (forthcoming 2026); Restatement (Second) Contracts § 201.

⁴¹ See, e.g., JERROLD J. KATZ, LANGUAGE AND OTHER ABSTRACT OBJECTS 128 (1981); Manuel Garcia-Carpintero, *Gricean Rational Reconstructions and the Semantics/Pragmatics Distinction*, 128 SYNTHESIS 93, 106 (2001); EMMA BORG, MINIMAL SEMANTICS (2004),

⁴² See, e.g., Keith Begley, *Katz got your tongue? The metaphysics of words*, 202 SYNTHESIS 1, 11 (2023); JERROLD J. KATZ, METAPHYSICS OF MEANING 278 (1990); Francesca Poggi, *Against the conversational model of legal interpretation: On the difference between legislative intent and speaker’s intention*, 40 REVUS 9 (2020).

⁴³ See *id.*

imitative parrots (indeed, they are programmed to mindlessly imitate human speech, just like parrots); the “utterances” of LLMs are (literally) meaningless even granting the existence of literal meaning, because they aren’t intentional communicative acts.

To be sure, we *could* (assuming that literal meaning exists) tell courts to take the language “agreed-upon” by LLMs in cases like the one sketched above as abstract sentence-types, and enforce their literal meaning, even if this particular “utterance” of the relevant sentence in fact had no meaning. There is, however, no reason I can think of to *want* to do this. As it turns out, the literal meanings of the utterances used by actual human beings in conversation are wildly and pervasively under-determinate, and the way in which we convey determinate meanings to one another in ordinary life depends a tremendous amount not on the literal semantic and syntactic rules of language but on speaker’s intentional reliance on mutually-understood context (“pragmatics”).⁴⁴ “I haven’t had breakfast” literally means *ever*—no one uses it to mean that because speakers, subjectively, know a great deal about what their interlocutors, subjectively, know about the kinds of creatures they both are and their mutually shared practices.⁴⁵

Contract law does not and never has enforced the language of contracts literally—indeed, it is again unclear if it could, as indeterminacy on this scale would run aground on contract’s requirement of definiteness.⁴⁶ (This pervasive under-determinacy, indeed, is part of why some philosophers doubt the existence of literal meaning—and even those who grant it might find that the characteristically “liberal truth conditions”⁴⁷ and “propositional radicals”⁴⁸ characteristic of literal meaning are probably not sufficiently definite for contract law to reliably enforce).

Even granting that in at least some cases the law *could* enforce the highly under-determined literal meanings of strings output by LLMs in apparent agreement with one another, we’d still be facing the possibility of courts enforcing content that *no one* (that is, no actual person with mental states) *ever* wanted to happen in any sense. If you can find some reason why the law ought to be in the business of doing this, let me know. And such an account would owe us, moreover, some reason why we ought to enforce the literal meaning of contractual language *only when* it is generated by an LLM—it is,

⁴⁴ See, e.g., RECANATI, *supra* note [RECANATI2004], at 6; 131; BORG, *supra* note [BORG2004], at 63; RAY JACKENDOFF, FOUNDATIONS OF LANGUAGE: BRAIN, MEANING, GRAMMAR, EVOLUTION 280 (2002); Robin Carston, *Word meaning and concept expressed*, 29 LINGUISTIC REV. 607, 609 (2012).

⁴⁵ See, e.g., Borg, *supra* note [BORG2004], at 92.

⁴⁶ See, e.g., 1 Williston on Contracts § 4:22 (4th ed.); Restatement (Second) of Contracts § 33(1).

⁴⁷ Borg, *supra* note [BORG2004], at 241.

⁴⁸ Carston, *supra* note [CARSTON2012], at 609.

of course, facially bizarre to interpret contracts in one way when they are entered into by human beings and another way when agreed to by LLMs; an argument in favor of doing so owes us both an explanation for why we ought to do enforce language output by LLMs *at all*, and why the posture for doing so ought to be so different than that applied in ordinary contract law.

And, again, there are other things courts could do to enforce these sorts of “agreements.” The law could, for instance, *imagine* that the relevant exchange that took place between LLMs took place between hypothetical human speakers, ask what they might have meant or been understood to mean (by imagined observers of various kinds).⁴⁹ Doing so, of course, would be fictional, and to that extent, the content of the fiction would have to be designed by the law—courts couldn’t purport to be discovering either the intended or literal meaning of the language, both of which are things that (might) actually exist; they’d be generating a new meaning that did not exist before, subject to some set of rules for doing so of their own design. From this perspective, the design of a set of interpretive rules to *attribute* fictional meanings to LLMs, taken as certain kinds of fictional speakers, is a normative project.⁵⁰ And it is once again not clear why the law ought to go through the trouble—and what the normative considerations might be that could justify the tremendous waste of time and resources.

To take things back to the fundamentals, there are *plenty* of reasons why the law has a role in enforcing genuinely agreed-upon content of human commitments—considerations of autonomy, the prerogative to shape one’s legal life, benefits to relationships, solicitude for the reasonable expectations of contracting parties, the insight of individuals into their own position and needs in the market, and so on.⁵¹ If anything, the law’s enforcement of agreements entered into voluntarily by human beings is overdetermined.⁵² Every one of these justifications, however, makes reference to the intentions and mental states of persons. Not one of them applies to LLMs, which do not have intentions or mental states. LLMs don’t have autonomy, they don’t have ongoing legal lives towards which the law might be justifiably solicitous, they don’t have relationships, and they don’t have reasonable expectations of one

⁴⁹ See, e.g., James Toomey, *Language to Law* (under review).

⁵⁰ See *id.*

⁵¹ See, e.g., FRIED, *supra* note [FRIED2015]; Curtis Bridgeman, *Twenty-First Century Contract Law is the Law of Agreements, Not Debts: A Response to Lewinsohn*, 129 Yale L. J. F. 535 (2020); Dori Kimmel, *Promise, Contract, Personal Autonomy, and the Freedom to Change One’s Mind*, in GREGORY KLASS, GEORGE LETSAS & PRINCE SAPRAI, EDS., *PHILOSOPHICAL FOUNDATIONS OF CONTRACT LAW* (2013); Thomas Scanlon, *Promises and practices*, 19 PHIL. & PUB. AFFS. 199 (1990); Erik Encarnacion, *Contract as Commodified Promise*, 71 VAND. L. REV. 61, 63 (2018); Aditi Bagchi, *Separating Contract and Promise*, 38 FLA. ST. U. L. REV. 707, 708 (2011).

⁵² See *id.*

another, they don't have insight into their own market needs, because they don't have any.

Moreover, there are even reasons why the law should not directly enforce the genuinely agreed-upon content of human commitments but rather *imagine* their exchange from the mouths of reasonable persons—this Solan's point, that, in the absence of better evidence, that *probably is* what they genuinely agreed upon, perhaps even more likely it is what a party who isn't lying understood, and there is a risk of fraud otherwise.⁵³ (It is plausibly worth it, in other words, for the law to create a set of rules for imagining fictional speakers *in contract cases involving human parties*—doing so might *better*, on average, approximate the actual intents of the parties. And even, if it does not, it may be more efficient for courts, and close enough, compared to directly attempting to litigate mental states).

But these, too, are normative considerations that make irreducible reference to the intents and mental states of the parties. They don't apply to LLMs untethered from human intents and mental states either—treating LLMs as imagined reasonable speakers cannot be justified as a way to approximate their intents, because they don't have them. Doing so can't even be justified as an efficient alternative to litigating the actual mental states of LLMs, because there would be nothing there to litigate. A legal rule enforcing the “agreements” entered into by LLMs, totally disconnected from anyone's mental states—whether literally, or by imagining them as certain kinds of fictional speakers, which appear to be the only alternatives—would have to stand on its own normative two-feet. And I, personally, can think of no reason why the law ought to be in the business of doing this.

Finally, while you're coming up with some reason, you might bear in mind that the issues on which I have focused here (having to do with the foundational normative purposes of contract law and the philosophy of language) are, in fact, only the tip of the iceberg—in the hypothetical exchange with which we opened the Part, one LLM purported to sell widgets to another. But LLMs cannot, at present, own property—we'd have to change that, and recognize LLMs as property owners if enforcing their contracts would be coherent at all. But doesn't our law's recognition of private property ownership have something to do with the mental states of owners?⁵⁴ And in cases where it doesn't look like a complete contract was reached, is there any reason to rely on promissory estoppel on behalf of entities that don't have, not just expectation interests, but interests of any kind?⁵⁵ Is it possible for LLMs to commit the tort of fraudulent misrepresentation without intent to

⁵³ See *supra* ____.

⁵⁴ See, e.g., Margaret Jane Radin, *Property and Personhood*, 34 STAN. L. REV. 957 (1982).

⁵⁵ See, e.g., Restatement (Second) of Contracts § 92.

defraud, intertwined as these torts are with contract and contracting?⁵⁶

From this perspective, enforcing the “agreements” of LLMs wholly disconnected from their relation to human intent might be possible, but would demand *systemic* change in how private law generally operates. Contract law, indeed, interacts seamlessly with the rest of (at least) private law—and enforcing these sorts of contracts would demand, ultimately, taking LLMs to be generalist legal persons for private law purposes. To the extent that they *are* philosophical zombies, and behave in ways indistinguishable from human beings, we *could* do this. But we need some normative reason to do so, and it is simply not clear what that would be (*unlike* human beings, the legal personhood of whom we have *ample* reason to recognize—once again, overdetermined if anything).⁵⁷

In short, understanding the “objective” theory of contracts to entail enforcing apparent “agreements” between LLMs like that proposed above is an exercise in the absurd—there are reasons to believe it is not meaningfully possible, and no reason to believe we *ought* to do what we *could* in that regard. At bottom (and as far as I can tell) the only possible justifications for contract law are subjective—an effort to enforce the actual intentional contents of mental states held by persons.

III. IMPLICATIONS

I’ve suggested that the possibility that the “objective theory” of contracts could demand enforcement of arbitrary “contracts” apparently “agreed” to by LLMs probabilistically spouting words at one another ought to serve as a *reductio* of the doctrine—this would be absurd, and any doctrine that would enforce the “agreements” of philosophical zombies must be rejected; or, at least, cannot be read to mean what it says. This understanding of the “objective theory” of contracts has several implications—I’ll highlight two.

First, it is worth emphasizing that, although the “objective theory” cannot make sense of the *normative foundations* of contract law, it may be (as Solan has argued) a perfectly reasonable evidentiary rule for a system of contract committed to enforcing subjective agreements. Second, nothing in this analysis suggests that LLMs cannot play an important role in contemporary contracting in the ways that they currently do—“entering agreements” ultimately related in some way to human mental states.

⁵⁶ See, e.g., Restatement (Second) of Torts § 525.

⁵⁷ Shawn Bayern, for example, argues in favor of liberally granting legal personhood to algorithms—but this is because he thinks doing so would help human beings facilitate their own purposes. See, e.g., BAYERN, *supra* note [BAYERN2022], at 5. What we’re looking for throughout this Part is some normative reason to extend these legal powers to LLMs *in their own right*. Bayern hasn’t provided one; I’m aware of no one who has.

A. Re-Situating the “Objective Theory”

For all its hysterical phrasing in the tomes of American legal pedagogy, it has never been clear that the “objective theory” of contracts actually captures the normative foundations of contract law. Courts *do* seek the genuinely shared meaning of contractual language, and rely on such proxies as the particular usage of the parties and the trade in which they are involved.⁵⁸ More or less all of the affirmative defenses contract law recognizes are proxies for whether the parties subjectively consented—if I told you last week that I’ve hired a sniper to follow you around, and that you must meet me in the town square to sell me your car, a reasonable third-party beholding the transaction might see nothing objectively amiss, but such a sale is obviously void.⁵⁹ Add to these examples the possibility that the “objective” theory might enforce the “agreements” of LLMs as philosophical zombies—contract law is not, from this perspective (nor could it plausibly be) foundationally committed to enforcing entities going through the motions of agreement.

But, of course, a legal regime committed to enforcing genuine agreements is vulnerable to the possibility of fraud—the possibility that one party might intentionally mislead the other by saying something other than what they mean; or might later claim in court that they meant something other than what they said. This is a challenge, of course, because individuals have privileged epistemic access to their own mental states, and it is very hard to prove the mental state of another person.⁶⁰ A legal regime committed to enforcing genuine agreements that simply deferred to self-serving post-facto testimony about their contents would fail routinely at its own aspirations.

So there *are* good reasons for such a legal regime to adopt something like the objective theory as an evidentiary rule—to *hold people to what they said* as their interlocutor would reasonably understand them, because that probably is what they meant, as against the risk they might later claim something else and we wouldn’t be able to prove otherwise. This is all Solan’s re-casting of the “objective theory”—as an evidentiary rule in the end justifiable in a legal regime that aspires to enforce genuine, subjective agreements, and tailored to do so.

But even if something like the “objective theory” may indeed be

⁵⁸ See, e.g., Restatement (Second) of Contracts §§ 201-202.

⁵⁹ See, e.g., Restatement (Second) of Contracts § 174.

⁶⁰ See, e.g., Alice Ristroph, *State Intentions and the Law of Punishment*, 98 J. CRIM. L. & CRIMINOLOGY 1353–65 (2008); Michael J. Kaufman & John M. Wunderlich, *Messy Mental Markers: Inferring Scienier From Core Operations in Securities Fraud Litigation*, 73 OHIO ST. L. J. 507, 507.

justifiable in a contract regime ultimately founded on the aspiration of enforcing genuine agreements, it is important to keep this role in mind and not get carried away by the fanciful behaviorist language in which the rule is often framed. This is because if it is an evidentiary rule justified by its tendency to shield against fraud and enforce genuine, subjective agreements, it only makes sense to the extent that it *does that*—and casting the rule in this light lets us assess it against its alternatives in the same ways we might assess any evidentiary rule or presumption; with reference to whether it is causing comparatively more false positives or false negatives than other possibilities, and which we are more concerned to avoid.⁶¹

Any legal rule, of course, will cause some errors—the question is whether the “objective theory” will end up failing to enforce more instances where there was no genuine, subjective agreement, or vice versa, than alternatives (and at what cost, and with what administrative complexity).⁶² It might, from this perspective, turn out that the civilian rule which relies on courts to more directly seek the subjective understandings of parties is preferable—or not; it might be the courts are not good at this, or that doing so is too expensive to be worth the number of errors it corrects.⁶³ Perhaps there ought to be more exceptions to the “objective theory’s” focus on external actions, where evidence of what the parties actually intended is very strong, or perhaps fewer.

But we can only engage in this kind of analysis if we are clear about the normative point of the doctrine—if we think that the objective theory articulates the normative foundations of contract law, it is meaningless to ask whether it effectively accomplishes its normative goals alongside alternatives; it just is its own normative goals. Understanding its proper role as an evidentiary rule *designed to* enforce genuine agreements over the possibility of fraud, however, lets us think about whether it’s working.

B. The Role of LLMs in Contracting

For those steeped in the practice of high-finance, my claim that courts enforcing agreements concluded entirely by LLMs would be literally absurd may seem peculiarly Luddite—indeed, this may well already be happening; at least, I understand, in the world of high-speed trading, transactions are routinely “agreed” by algorithms “talking” to one another in ways that

⁶¹ See, e.g., SITKOFF & DUKEMINIER, *supra* note [S&D], at 144.

⁶² See, e.g., Michael S. Pardo, *Second-Order Proof Rules*, 61 FLA. L. REV. 1083, 1107 (2009); Margaret Jane Radin, *Market-Inalienability*, 100 HARV. L. REV. 1849, 1910-11 (1987).

⁶³ See, e.g., William Tetley, *Mixed Jurisdictions: Common Law v. Civil Law (Codified and Uncodified)*, 60 LA. L. REV. 677, 712 (2000).

everyone assumes to be enforceable.⁶⁴ But *these* uses of LLMs—to conclude contracts that in fact the people who developed or deployed the algorithms *would in fact want to agree to* in a more efficient way—differ from the cases I take to be absurd discussed above in precisely the relevant respects. Nothing here is inconsistent with the assumption that these sorts of contracts are and ought to be enforceable.

Human Being A and Human Being B can agree to hold themselves to “contractual” output of an LLM, or of two designated LLMs talking to one another, if they’d like.⁶⁵ So can Group of Human Beings A (“Corporation A”) and Group of Human Beings B (“Corporation B”).⁶⁶ There’s nothing particularly novel about this—A and B could agree to hold themselves to content put out by, say, the American Federation of Musicians;⁶⁷ or to the meaning they take to be contained in a particular cloud formation (perhaps as interpreted by a specified oracle, to provide the definiteness required for enforcement); or anything at all (whether any of this is a good idea is, of course, a separate question—one contract law *leaves to A and B*).⁶⁸ There is a real agreement in these cases between entities with mental states—namely, A and B (or, at least, collectivities composed of entities with mental states, corporations; whose decision-making is the product of the agreed-upon procedures of the entities that compose them).⁶⁹

And that—the agreement between A and B—is *the essential agreement in all of these cases*, without which there would be no contractual content. The contents of this agreement would be the target of a subjectivist conception of contract, even if its contents direct courts to facts about the world independent of the parties themselves—whether the output of specified LLMs, form contracts, the background assumptions of a particular industry, or cloud formations and accompanying oracular proclamations. In all these cases, it is a kind of incorporation by reference of some content not emerging between the parties into the agreement between the parties—and that incorporation takes place *by agreement* of the parties.

Indeed, note that none of these things—cloud formations, oracular proclamations, the form contracts of a particular industry, or the outputs of

⁶⁴ See, e.g., BAYERN, *supra* note [BAYERN2022], at 158; Gina-Gail S. Fletcher, *Deterring Algorithmic Manipulation*, 74 VAND. L. REV. 259, 259 (2021); Yesha Yadav, *How Algorithmic Trading Undermines Efficiency in Capital Markets*, 68 VAND. L. REV. 1607, 1607 (2015).

⁶⁵ See, e.g., BAYERN, *supra* note [BAYERN2022], at 34-35.

⁶⁶ See, e.g., Victoria F. Nourse, *Elementary Statutory Interpretation: Rethinking Legislative Intent and History*, 55 B.C. L. REV. 1613, 1615 (2014); Peter A. French, *The Corporation as a Moral Person*, 16 AM. PHIL. Q. 207, 207 (1979).

⁶⁷ GRAHAM V. SCISSOR-TAIL, INC., 623 P.2d 165, 812 (1981).

⁶⁸ See, e.g., BAYERN, *supra* note [BAYERN2022], at 34-35.

⁶⁹ See *supra* note ____.

LLMs—have legal effects on their own. It is only by virtue of the agreement of A and B that they *might* play any role in ascertaining contract obligations. And whatever that role might be is governed and constrained by the agreement between A and B—who might well elect to govern their obligations by the output of an LLM “as interpreted by a reasonable person, imagining the LLM as a reasonable person” or “interpreted literally,” or “the meaning of the clouds as interpreted by three independent oracles,” and so on.

Compare, from this perspective, use of LLMs in this way to form contracts—another staple of the modern economy sometimes thought to challenge subjectivist justificatory views of contract.⁷⁰ When I purchase a product from Amazon pursuant to a lengthy set of terms and conditions I haven’t read, there is a genuine agreement between me and Amazon—we’ve agreed that they’ll ship me the product, and I’ll pay the listed price.⁷¹ To the extent the written content of the form contract reflects that genuine agreement, more power to it.⁷² To some extent, surely, our genuine agreement contemplates incorporating content from Amazon’s form contract to precisify our obligations on such matters as return policy, about which I have general expectations but am happy to defer to their specified choices.⁷³

But it is *the genuine agreement* between me and Amazon that gives this language any legal effect at all, and constrains the extent to which it does—any language Amazon has inserted in there that falls outside of my “reasonable expectations” (that is, outside the scope of our genuine agreement) is unenforceable.⁷⁴ To the extent that some courts have been led astray by the insistent language of the objective theory of contract to enforce language in form contracts and end-user license agreements outside of the scope of the parties’ genuine agreement, they are doing something non-contractual and confused (another reason to be clear about the purpose of the objective theory—to ensure that courts don’t reify and apply it beyond its proper scope in a variety of ways, as apparently some have).⁷⁵

⁷⁰ See, e.g., Michael I. Meyerson, *The Reunification of Contract Law: The Objective Theory of consumer Form Contracts*, 47 U. MIAMI L. REV. 1263, 1265 (1993).

⁷¹ See, e.g., Robin Bradley Kar & Margaret Jane Radin, *Pseudo-Contract and Shared Meaning Analysis*, 132 HARV. L. REV. 1135, 1165 (2019). Obviously, “Amazon” is not an entity with mental states, but it is composed of entities with mental states who, through their own agreed-upon processes, have agreed to a standard offer to ship people glue for a specified price when those people complete certain processes.

⁷² See, e.g., Kar & Radin, *supra* note [K&R2019], at 1173.

⁷³ See *id.* at 1208.

⁷⁴ See, e.g., *C&J FERTILIZER, INC. v. ALLIED MUT. INS. CO.*, 227 N.W. 2d 169, 176 (Iowa 1975); Robert E. Keeton, *Insurance Law Rights at Variance with Policy Provisions*, 83 HARV. L. REV. 961, 967 (1970).

⁷⁵ See, e.g., Kar & Radin, *supra* note [K&R2019], at 1139; see *supra* ____.

And contrast all this with the kinds of “agreements” under analysis here—two LLMs “talking” to one another disconnected from any agreement between entities with mental states or indeed disconnected from any mental states at all. The critical distinction, of course, is that in these cases—like the example that opens Part II—*there is no genuine agreement* between entities with mental states at *any* point. It is not that two entities with mental states have genuinely agreed to defer their obligations to the output of the model, nor incorporate the output of the model, in any sense, into their agreement, nor take the output of the model to approximate and precisify the content of their genuine agreement (conditional on its at least roughly doing so). There *is* no genuine agreement *at all*. And it is in precisely that sense that enforcing these sorts of apparent agreements is absurd—like enforcing the contents of a “form contract” that, say, just appeared on the internet and bears no resemblance to anything anyone actually agreed to, or enforcing the contents of purported messages contained in the clouds *sua sponte*.

From this perspective, the account offered here can hardly be read to purge LLMs from an important role in contracting in the contemporary economy—indeed, it *supports* it. If our financiers think that deferring to algorithms for the precise timing and content of high-speed transactions is a good idea, more power to them, in the same way that it is impossible to imagine the efficiency of the modern economy without some role for form contracts. Courts will, and should, consistent with the normative aspirations of basic contract law, enforce these agreements *because* the financiers agreed to them, *as* the financiers agreed to them.

But all this is a far cry from the absurd entailment of the objective theory of contract, taken literally, sketched above—that it might lead courts to enforce the outputs of two LLMs generating words at one another *unrelated to* any actual agreement between human beings. This *is* absurd, and getting clear about what this absurdity tells us about the normative foundations of contract law can, moreover, help clarify how the law ought to approach algorithmically drafted contracts in the context in which they are actually used.

CONCLUSION

“I suspect,” says Captain Renault, “that under that cynical shell you’re at heart a sentimentalist.”⁷⁶ This is what I’ve suggested about contract law. For more than a century, courts and commentators have insisted that contract law does not care about subjective mental states, understandings and aspirations of the human beings that are party to contracts—what matters, we’ve been

⁷⁶ Casablanca (Warner Bros. – 1942).

told, is the *appearance* of agreement, as an objective third-party would see it, not what any actual human beings ever thought or wanted.

Applied literally to the arbitrary outputs of large language modules to one another, the results are absurd—the objective theory would have courts enforcing the “agreements” of philosophical zombies; entities that act indistinguishably from humans, but lack mental states. There is an alternative—to understand the objective theory of contracts as “a simple—and sensible enough—rule of evidence . . . designed to prevent litigating parties from relying upon testimony of an earlier state of mind without evidence.”⁷⁷ Seeing this clearly can help assess whether the posture is any good as an evidentiary role, and clarify how courts ought to think about the entirely reasonable role of LLMs in contemporary contracting.

⁷⁷ Solan, *supra* note [SOLAN2007], at 379.