Artificial Meaning

Lawrence B. Solum

ARTIFICIAL MEANING[*]

Lawrence B. Solum[**]

INTRODUCTION: FROM ARTIFICIAL INTELLIGENCE TO ARTIFICIAL MEANING

As artificial intelligences (AI) become more powerful and pervasive, communication by, with, and among AIs has become a common feature of everyday life. Early in the history of AI, there was ELIZA—a simple program that utilized simple pattern-matching algorithms to simulate a psychotherapist interacting with the user of the program.[1] Human communication with AIs has been depicted in film and fiction, from the iconic confrontation of humans with HAL in *2001: A Space Odyssey*[2] to the very human Theodore who falls in love with an artificially intelligent operating system named Samantha in *Her*.[3] But complex communication with AIs is part of everyday life in modern technological societies, including Apple's Siri, automated telephonic service systems for airlines, online ordering systems for merchants like Amazon.com, and characters in video games who converse with human players. All of these contexts involve *artificial meaning*, which we can contrast with communications by natural persons (human beings)—which we can call *natural meaning*.

This Essay investigates the concept of artificial meaning, meanings produced by entities other than individual natural persons. That investigation begins in Part I with a preliminary inquiry in the meaning

---

1. Joseph Weizenbaum, *ELIZA—A Computer Program For the Study of Natural Language Communication Between Man And Machine*, COMM. ACM, Jan. 1966, at 36, 36.

2. ARTHUR C. CLARKE, 2001: A SPACE ODYSSEY 155–56 (1968); 2001: A SPACE ODYSSEY (MGM 1968).

3. HER (Annapurna Pictures 2013).

of "meaning," in which the concept of meaning is disambiguated. The relevant sense of "meaning" for the purpose of this inquiry is captured by the idea of communicative content, although the phrase "linguistic meaning" is also a rough equivalent. Part II presents a thought experiment, *The Chinese Intersection*, which investigates the creation of artificial meaning produced by an AI that creates legal rules for the regulation of a hyper-complex conflux of transportation systems. The implications of the thought experiment are explored in Part III, which sketches a theory of the production of communicative content by AI. Part IV returns to *The Chinese Intersection*, but *Version 2.0* involves a twist—after a technological collapse, the AI is replaced by humans engaged in massive collaboration to duplicate the functions of the complex processes that had formerly governed the flow of automotive, bicycle, light-rail, and pedestrian traffic. The second thought experiment leads in Part V to an investigation of the production of artificial meaning by group agents—artificial persons constituted by rules that govern the interaction of natural persons. The payoff of the investigation is presented in Part VI. The communicative content created by group agents like constitutional conventions, legislatures, and teams of lawyers that draft complex transactional documents is *artificial meaning*, which can be contrasted with *natural meaning*—the communicative content of those exceptional legal texts that are produced by a single individual. This insight is key to any theory of the interpretation and construction of legal texts. A conclusion provides a speculative meditation on the implications of the new theory of artificial meaning for some of the great debates in legal theory.

I.   MEANING: WHAT AND HOW?

First the *what*, then the *how*. We begin with a general inquiry into the meaning of "meaning."

A.   *What Is Meaning?*

This Essay is about meaning, but what does "meaning" mean? The word is ambiguous, as illustrated by the following examples:

- Do you think that those clouds mean it is going to rain?
- I didn't mean to hurt your feelings.
- Do you know how much you mean to me?
- I got the letter translated! Now, I know what it means.

The ambiguity of "meaning" afflicts the law as well. When we ask about the meaning of a constitutional provision, statute, rule, regulation, or judicial opinion, we could be asking about the kind of meaning that

translators provide—roughly linguistic meaning, but more precisely communicative content. But we might be asking about the implications of a text for legal doctrine or for a particular case. Or we might be asking about the purpose or function of the text. These are all distinct senses of "meaning," and before we proceed further, we need to sort them out. There are at least four senses of "meaning" relevant to legal texts:

Communicative Meaning. Suppose someone asks a question about what the Framers meant by using the phrase "arms" in the Second Amendment: were they referring to weapons or the upper limbs of the human body? This sense of meaning refers to communicative content— the content conveyed by the legal text given the context in which it was authored. In this Essay, I will use the word "meaning" to refer to this sense of meaning, and for greater clarity I will sometimes use the phrase "communicative meaning."

Legal Meaning. The communicative content of legal text is not the same thing as the legal content associated with the text. For example, the communicative content of the phrase "freedom of speech" is quite sparse, but the legal content associated with the phrase is very rich indeed: it encompasses a complex set of doctrines, addressing topics like prior restraint, defamation of public officials, and even the regulation of billboards. We can call meaning in this sense "legal meaning."

Application Meaning. Sometimes when we ask about the "meaning" of a legal text, we are asking about the implications that it will have, usually in a particular context. For example, we might ask, what does the First Amendment freedom of speech mean for my defamation suit? Does it provide me a defense? When "meaning" is used in this sense in the context of a legal text, we are concerned with the application of the text to a particular case or to some set of cases. Let us call this "application meaning."

Purposive Meaning. Meaning is also used to refer to the purpose or motive that produced a particular legal text. For example, we might refer to the aim of a constitutional provision by saying, "What did the drafters mean to accomplish through the Privileges or Immunities Clause of the Fourteenth Amendment?" We can call this "purposive meaning."

So, there are four senses of "meaning" in legal contexts. Our focus will be on communicative meaning, but before we begin to explore artificial meaning we need to say something about how communicative content is produced.

B. How Is Meaning Produced?

Theories of meaning are the terrain of the philosophy of language and theoretical linguistics, but in a short essay, even a short and dirty tour is

impossible. Instead, we will examine one important approach to theorizing about theories of meaning. That approach is associated with philosopher Paul Grice, but this is a simplified version that differs from his account in several ways.[4] We can begin with Grice's idea of speaker's meaning,[5] although we will primarily be concerned with legal texts, so our version of the idea could more precisely be called "drafter's meaning."

We can begin with a definition:

> *Speaker's Meaning*: The speaker's meaning (or utterer's meaning) of an utterance is the illocutionary uptake that the speaker intended to produce in the audience on the basis of the audience's recognition of the speaker's intention.

Quite a mouthful! Grice's idea is quite general; it can be illustrated with an example in which there is neither speech nor text, but communication occurs nonetheless:

> Imagine that you have stopped at night at an intersection. The driver of another car flashes her lights at you, and you make the inference the reason for her doing this is that she wants to cause you to believe that your lights are not on. And based on this inference, you now do, in fact, realize that your lights are not on.[6]

Communication succeeds because the other driver recognizes your communicative intention. In the language of game theory, you and the other driver have common knowledge of this communicative intention. Of course, drivers recognize many communicative intentions when they are on public roads and highways. In the United States, octagonal red "stop" signs communicate a legal command, roughly "halt before proceeding."

Here is another example, from the California Driver Handbook:[7]

---

4. *See* PAUL GRICE, STUDIES IN THE WAY OF WORDS 3–143 (1989).

5. *See generally* H.P. Grice, *Utterer's Meaning, Sentence-Meaning, and Word-Meaning*, 4 FOUND. LANGUAGE 225 (1968).

6. This example is a slightly altered version of the thought experiment originally offered by Richard Warner. *See* Richard Warner, *Introduction: Grice on Reasons and Rationality*, *in* PAUL GRICE, ASPECTS OF REASON ix (2001).

7. CAL. DEP'T OF MOTOR VEHICLES, CALIFORNIA DRIVER HANDBOOK 29 (2014), *available at* https://apps.dmv.ca.gov/pubs/dl600.pdf.

A *five-sided sign* indicates that you are near a school. Stop if children are in the crosswalk.

Traffic signs are a particular kind of legal communication. They convey communicative content, and some of them (stop signs, for example) have associated legal content. If you run a stop sign, you can get a ticket for an offense punished by a fine.

In conversation, speaker's meaning can rely on a rich communicative context. If Ben and Alice are having a conversation at a coffee house, there are multiple channels for producing *uptake*—the recognition of communicative intent. Imagine the following dialog:

Alice: Another?

Ben: Black this time.

Alice: Sugar? [rising inflection]

Ben [shaking his head from left to right]: Uhhh.

Alice: OK.

The communicative content is far richer than the literal meaning of Alice and Ben's utterances. "Another" meant "would you like another cup of coffee." "Black this time" meant "I would like the coffee without milk or cream." "Uhhh" plus the shake meant "I would like the coffee without sugar." If we take away our knowledge of context, the same words might mean something totally different. Try your hand at inventing different contexts for the same words and you will see what I mean.

Some legal communication occurs in information-rich contexts. Face-to-face negotiation of an oral contract is a good example. But other legal texts are produced in communicative situations that filter information about the drafter's communicative intentions. For example, the original United States Constitution was drafted by an exclusive group in a secret convention held in Philadelphia,[8] and much of the detailed drafting work was done by the Committee on Style, a subgroup that did its work away from the main convention.[9] How does this kind of communication

---

8. Vasan Kesavan & Michael Stokes Paulsen, *The Interpretive Force of the Constitution's Secret Drafting History*, 91 GEO. L.J. 1113, 1115 (2003).

9. *Id.* at 1206–07.

work? How can groups form communicative intentions? And how do the readers of such texts grasp these intentions? These questions are so hard that it is difficult to see how we can approach them. This Essay explores an unorthodox route to an understanding of the production of communicative content in legal texts drafted and promulgated by complexly structured groups. That exploration begins with a thought experiment.

II.   THE CHINESE INTERSECTION, VERSION 1.0

Imagine that you are in the not-so-distant future, looking back at the history of traffic management in Shanghai, an urban megalopolis that we shall imagine became even larger and more densely populated. Somewhere in the heart of this imaginary version of Shanghai there was a supremely complicated intersection, where ten major roads, three highways, six surface light rail lines, and twenty-three pedestrian walkways intersect. Traffic engineers and the city officials tried all the usual tools to manage the intersection. Stoplights, left-turn-only lanes, pedestrian walk/don't-walk signs, speed limits, don't-block-the-box signs with big fines, and so forth. None of it did any good. The "Chinese Intersection" as we shall call it was the source of traffic jams of epic proportions—get stuck at the wrong place at the wrong time, and your five-hour commute lasted for five days. Literally.

At some point, the officials realized that human beings just weren't up to the task of managing traffic in the Chinese Intersection. For every move the human regulators made, frustrated drivers had a countermove—always attempting to circumvent the rules in order to shave a few minutes off the journey through the Chinese Intersection. Economists wrote articles that explained that driver behavior can be modeled using the techniques of game theory. The Chinese Intersection, the economists said, involved multiple Prisoner's Dilemmas. If every driver obeyed the rules, then each individual driver would benefit, but no driver could trust the others, so almost everyone cheated, and the result was chaos. Other economists calculated the effects of the Chinese Intersection on productivity, transportation costs, and real estate values. Everyone knew that the number would be a big one, but when the results of the study were announced, it precipitated a political crisis. The Chinese Intersection imposed annual costs of close to 1.6 trillion Yuan (or approximately 100 billion dollars).

The good news was that the budget for solving the problem became functionally unlimited. Tasks forces met. Entrepreneurs proposed solutions. Politicians called for immediate action. The obvious solution was to eliminate the Chinese Intersection altogether—perhaps a nice

park could replace it. But then a simulation was run, and it turned out that this would make the traffic problems in Shanghai far worse. Eventually, a solution was proposed. Traffic in the Chinese Intersection (and eventually all of Shanghai) would be managed by an AI that would operate a system of traffic controls. The system would operate a massive system of programmable lane markers, signs, signals, gates, bumps, broadcasts, robot vehicles, and vehicle removal cranes. The goal was to create a system that would write the traffic laws, inform drivers of their contents, and then enforce the laws by removing noncomplying vehicles and collecting on-the-spot fines from the drivers.

And so a team of programmers and traffic engineers began the process of developing the Shanghai Artificially Intelligent Traffic Authority (SAITA). Early on in the process, the developers realized that SAITA would need the capacity to adapt itself to changes in driver behavior and traffic flow. The programmers produced a system that was capable of changing the traffic code, speed limits, signage, and lane configurations—just about every element of traffic regulation in the city. The system was designed to introduce random variations and run controlled experiments to evaluate the effects of various combinations on traffic patterns. Violations would be detected by an elaborate system of electronic surveillance. Offenders were identified and immediately would be removed from traffic by a system of cranes located at key intersections. An automated adjudication system imposed fines for minor offenses and jail terms for more serious violations.

The rollout of the new system was not great. There were bugs and for several days, traffic actually became much worse. SAITA produced a variety of legal texts, signs, and symbols. For example, SAITA periodically revised the punishments associated with traffic code violations—optimizing the level of punishment by modeling the costs and benefits in light of the massive amounts of data it collected on driver behavior. SAITA operated the complex system of traffic signs. In the early days, these signs were simply electronic versions of the familiar signs:

But as the system introduced random mutations, the signs began to

morph as the system learned new techniques for communicating with drivers. Running a static stop sign resulted in a fine of ¥2000, but if the stop sign was flashing, the fine increased to ¥5000 and if the stop sign was flashing, spinning, and cycling from red to orange and purple, then the fine increased to ¥10,000. The system was even capable of introducing its own vocabulary. When the AI detected a new violation type that had a particularly harmful effect on traffic flow, it could invent a word (represented by a logogram or character) to name the type and add a provision to the traffic code. SAITA could even initiate a public information campaign with entertaining YouTube videos and giant billboards defining the new word and instructing the public about the new law.

SAITA's ability to create new words was put to use when the system noticed that particularly aggressive drivers would pull off the street and onto the sidewalk to circumvent particularly bad traffic. SAITA experimented with several different words and logograms to communicate a new traffic code provision that imposed a mandatory minimum sentence of ninety days in jail for this behavior—finally settling on a neologism that would literally translate into English as "dogging." Why "dogging"? SAITA had posted an animated video that portrayed a car driving on a sidewalk as a snarling, aggressive dog. Programmed to monitor the comments to its video, SAITA observed that "dog" was the most common word and then experimented with several variations. "Dogging" was most effective, and once the data revealed that "No Dogging" signs worked, SAITA added a "No-Dogging" provision to the traffic code.

Traffic improved noticeably, even in the first few weeks, but as the system began to learn from experience, the effects became dramatic. Lanes were adjusted for traffic flow, the laws were optimized, and the system did a much better job of communicating with drivers than the old human-run system had ever done. Within two years, the traffic flowed smoothly through the Chinese Intersection. Commute times went down, happiness levels went up, and there were even calls to extend SAITA's domain to regulation of things other than traffic.

What accounted for the success of SAITA? An obvious factor was that SAITA used sophisticated algorithms to adjust traffic patterns—two-way streets could become one-way thoroughfares, a through-lane of traffic could become a dedicated left-turn lane, speed limits could be adjusted—and these adjustments were made dynamically in real time in response to the number of vehicles, their destinations, and so forth. But there was another factor: SAITA was supremely good at communicating traffic rules to drivers. Why? Because SAITA was able to determine

whether its efforts to communicate were working. If SAITA changed the law, but driver behavior did not respond, SAITA would try another form of communication, again and again, until the message was getting across.

* * *

A Digression on the Role of Thought Experiments in Legal Theory

Legal theorists live and die by the hypothetical—the particularly legal version of the thought experiment. We use thought experiments as intuition pumps—the hypothetical is designed to provoke (or "pump") an intuition or reaction. A well-constructed legal hypothetical elicits intuitions about what should be done, and these intuitions can then be used to critique existing rules or to show that rules are justified. Thought experiments can be used to pump other kinds of intuitions as well. I am about to use the Chinese Intersection to pump some intuitions about meaning—the way it works and how it comes to be. I am going to attempt to show you that some of the things you may believe about the production of meaningful legal texts are not quite right. But you should be on your guard. Accidental features of the setup might produce the intuitions pumped by the Chinese Intersection thought experiment. Although the Chinese Intersection tells a story that seems plausible—we can imagine something like SAITA might become possible in the not-too-distant future—skepticism is appropriate. Is this science fiction doing illegitimate work?

* * *

III. THE PRODUCTION OF COMMUNICATIVE CONTENT BY ARTIFICIAL INTELLIGENCES

SAITA would be an AI and the meanings it would produce would be artificial meanings. But is this even possible? Isn't meaning a function of intentions? And aren't intentions mental states? And don't mental states require a mind? SAITA would just be a machine. Doesn't it follow that SAITA couldn't have a mind? And lacking a mind, doesn't it follow that SAITA could not produce meanings? So if there are meanings in the thought experiment, they must be produced by human individuals—perhaps the programmers or the working groups that gave the programmers their instructions. "Artificial meanings" must be a fiction, "transcendental nonsense," as it were—the real meanings must always be natural meanings, produced by natural persons.

The attitude conveyed by the rhetorical questions in the immediately prior paragraph is skepticism about artificial meaning. In this Part of the Essay, we will investigate that skepticism. In the end I hope to convince

you that such skepticism is unwarranted and that artificial meanings are both conceivable and possible. Let's begin with what I think is the root of the skepticism—a view that we shall call the "folk theory of meaning."

Here is a picture of meaning—communicative meaning that is. Meaning begins in our minds—in our heads, so to speak. We have a thought or idea. We then encode that thought in some medium of communication—usually a natural language like English, but creative beings that we are, we can use other techniques—like flashing our headlights or making a gesture. Someone else, a listener or reader, then receives the encoded message. They hear what we say or read what we wrote. Then they decode the message. If they do that successfully, then they will have a thought or idea that is identical to the one we had.

The folk theory of meaning seems quite plausible on the surface. We sometimes talk in ways that seem to presuppose this theory. "You get my idea, don't you?" Or perhaps, "we are on the same wavelength." "You really understand me. It's like you are reading my mind." If the folk theory were true, it would create serious problems for the Chinese Intersection thought experiment. SAITA would not have conscious experiences. SAITA has no head and so there can be no thoughts in the head that SAITA lacks. So SAITA would not be encoding its thoughts, and if it couldn't encode its thoughts, then we couldn't decode them! If the folk theory of meaning were true, then whatever SAITA would be doing, it would not be communicating. The legal code that SAITA wrote would be meaningless. The traffic signs and lane markings it produced would be gibberish. No mind, no meaning.

So if we accept that the folk theory of meaning is correct, then SAITA could not produce communicative content. But the Chinese Intersection experiment seems to elicit strong intuitions to the contrary! So either the folk theory of meaning is false, or its defenders must produce some explanation that accounts for meaning in the Chinese Intersection thought experiment. The most obvious line of reply is to try to show that what appears to be artificial meaning is actually natural meaning (meaning produced by natural persons) in disguise.

How could we argue that artificial meaning reduces to natural meaning? The obvious move is to locate the meaning in the heads of the programmers or those who gave the programmers their instructions. When SAITA uses the word "stop," the meaning comes from the programmer who added the word "stop" to the traffic vocabulary module in SAITA's vocabulary. Consider ELIZA, the primitive AI that simulated a psychotherapist. In ELIZA's case, it seems plausible that all of the sentences produced by ELIZA have meanings that are reducible to

meanings produced by natural persons. Some of the meaning was produced by Joseph Weizenbaum, who wrote the program. Likewise, some of the meaning was produced by the users of ELIZA, because things they said could be incorporated in ELIZA's responses. But all of the meanings produced by ELIZA can be reduced to meanings produced by natural persons.

Of course, SAITA is more complex than ELIZA—by several orders of magnitude. There were many programmers and many others, individuals and committees of individuals, giving instructions to the programmers. But in the end, the argument would go, all of the meaning produced by SAITA could be traced back to meanings produced by natural persons. SAITA's traffic code is a complex concatenation of meanings that began as thoughts or ideas in the heads of thousands of individuals. The reason that SAITA can communicate is that drivers are able to decode the logograms and symbols that SAITA's algorithms used to encode the mental states of the programmers who created SAITA.

But the folk theory of meaning cannot account for SAITA's behavior or the reactions of drivers and pedestrians to the things SAITA does. This is because SAITA goes beyond repeating and recombining expressions provided by its programmers. In the dogging example, SAITA actually recognizes a new type or kind of behavior and then coins a new word to represent the new concept. None of the programmers who wrote SAITA's code knew anything about dogging (the concept) and none of them knew the meaning of "dogging" (the word). "Dogging" has artificial meaning—meaning created by an AI that cannot be reduced to meanings created by natural persons.

"Dogging" was built into the thought experiment to make it clear that artificial meaning cannot always be reduced to natural meaning. The thought experiment reveals the possibility that artificial meanings could be brought into the world by something other than a natural person. It is true that SAITA itself was created by humans. The important feature of the thought experiment is *not* that SAITA would be independent of natural human persons—it would depend on humans for its very existence. The important thing *is* that SAITA's meanings are not reducible to meanings produced by natural human persons. Dependence is one thing; reduction is quite another.

This point does not rely on the "dogging" example. It would still be the case that SAITA could create artificial meanings, not reducible to natural meanings, even if SAITA's vocabulary was limited to words and phrases whose invention was traceable to natural persons. Even if SAITA's vocabulary were entirely composed of words that were already present in its program, SAITA would create artificial meanings if it

could combine these words in new ways. In fact, this is also how natural persons create new expressions. Ordinarily, we do not coin new words when we communicate. We combine words we borrow from others in new ways—just as I am doing now in writing this article. The fact that my words are borrowed from others does not entail the conclusion that other people produced the meaning of this article. That conclusion is obviously false. Just as natural persons create new meanings through novel combinations of words, so too can we imagine that SAITA might do so.

The folk theory of meaning is based on the premise that meanings are necessarily "in the heads" of natural persons. The assumption is that meaning is a mental event that occurs in the mind of the speaker or author. That meaning is then encoded into language—something said orally or written down in a text. If that were true, then SAITA couldn't produce meaning, because SAITA doesn't have a mind of its own—it is just a computer program. SAITA (by hypothesis) is not conscious; it has no self-awareness. So if SAITA does produce meanings, that is, the various directives that enable the Chinese Intersection to function smoothly, then it has to be the case that we can have meanings that are produced artificially, by entities that are not conscious natural persons.

IV.   AFTER THE APOCALYPSE: THE CHINESE INTERSECTION, VERSION 2.0

Here is another thought experiment. Suppose that the SAITA were destroyed or disabled by some calamity. Perhaps a particularly destructive computer virus infects SAITA and the technical support staff is unable to repair the program. Of course, the immediate effect of this calamity is chaos at the Chinese Intersection; once again, automobile traffic in Shanghai is paralyzed. Suppose that the traffic officials respond to this calamity by using human beings to replace SAITA.

In the immediate aftermath of SAITA's demise, the traffic officials decide to appoint a traffic czar. One official is given responsibility for making all the decisions about the various traffic signals. The traffic czar is tasked with the job of rewriting the traffic regulations in response to changing conditions. But it turns out that the traffic czar is simply unable to keep up with the complexity of the problem. The czar cannot keep up with the massive flow of information from the human traffic spotters who take the place of SAITA's electronic monitors. The czar attempts to revise the traffic code in response to new problems, but that job is too big for one individual. Either the czar falls behind, unable to produce well drafted regulations at the required pace, or, even worse, the czar does keep up, but writes regulations that are confusing, ambiguous, and

have unintended consequences. Under the czar, the Chinese Intersection is a complete mess—worse than the "bad old days" before SAITA.

So the czar is replaced. City officials assemble a team of programmers and managers who analyze the code that produced SAITA's operation, and give teams of human beings responsibility for performing the roles that various modules of the SAITA program performed. Of course, the new system is not as efficient as SAITA. Electronic monitoring of the Chinese Intersection is replaced by human observers. The system of electronic signs is replaced by hand-painted signs. SAITA's almost instantaneous response to traffic problems is not matched, but the humans who implement the routines and algorithms in SAITA's programs are able to do a pretty good job. Like SAITA, they implement new traffic rules in response to changing conditions.

Experience with the traffic czar had demonstrated that a single individual could not duplicate SAITA's operation. Instead, the work of monitoring traffic, producing traffic signs, and creating new traffic rules is distributed among several large teams of human beings. Spotters produce reports on the behavior of the various vehicles, trains, and pedestrians. Another group analyzes the reports and routes them to teams that control the traffic signals and lane gates. These teams issue commands to runners who post new signs and open and close the various gates.

Like SAITA, the new system was capable of producing new traffic rules. Instead of assigning the work of drafting to a czar, the job of writing the rules was broken down into subtasks. One group worked on developing descriptions of prohibited behavior. Another group developed a system of sanctions. Yet another group monitored the effects of the sanctions and implemented an algorithm for adjusting the penalties to optimize traffic flow.

No one individual drafted the traffic code. Indeed, no one individual reviews the whole text of new regulations as they are produced. Instead, different sections are produced by different teams and assembled by other groups. The assembled regulations are then considered by another group (the Regulation Committee), which operates by majority vote, either approving the regulation, rejecting it, or sending it back for additional work by the drafting and assembly teams. The Regulation Committee consists of senior-level managers who have multiple responsibilities. It turns out that it is quite rare for each and every individual member of the team to read the whole text of a new regulation. Indeed, many regulations are approved on the basis of executive summaries that are read by the assistants to the team members. It is not uncommon for the final text to be approved without any one

individual having read the whole text. The creation of the meaning of new traffic regulations is a group effort.

This new system for generating traffic regulations is the Group Agent for Shanghai Traffic (GAST). The regulations that GAST produces have communicative content. When a new regulation is promulgated by GAST, the rules are posted on large billboards and enforced by human traffic police officers. The rules are applied by human traffic judges, although the human judges are not quite as accurate as SAITA had been.

SAITA worked because it was a very good communicator. It was able to formulate rules and regulations that were understood by drivers and pedestrians. GAST may not have been quite as good at communicating as SAITA, but it was good enough. The humans who implemented SAITA's algorithms were not as fast as SAITA, but they were able to produce a workable system of traffic regulation. Traffic flowed at a reasonable pace through the Chinese Intersection—slower and with more frequent interruptions because of accidents, but still much better than the "bad old days" and much, much better than the brief regime of the traffic czar.

V.   THE PRODUCTION OF COMMUNICATIVE CONTENT BY GROUPS

What are the lessons of the Chinese Intersection, Version 2.0, for our understanding of legal communication? Recall that the folk theory of meaning is based on the assumption that meanings are a function of the mental states of individual human minds. That theory was unable to accommodate successful communication by an AI. SAITA produced meanings without mental states. Similarly, the folk theory of meaning has difficulty with meanings that are produced by groups rather than individuals. If meanings just are the encoded and decoded mental states of individuals, then it follows that the notion of "group meaning" is incoherent. Groups do not have mental states or minds, and therefore the very idea of a group meaning would be incoherent. Instead, the folk theory must postulate that *apparent* group meanings are actually meanings produced by individual natural persons. So folk theory requires that the meanings produced by GAST must actually be assemblages of meanings by the individual members of the teams that make up GAST.

But it turns out that the folk theory of meaning is simply unable to account for the way that GAST actually works. GAST's procedures for the production of legal texts were modeled on the algorithms used by SAITA. Although fragments of legal texts correspond to the mental states of team members, the meaning of the text as a whole does not

correspond to the mental states of any one individual. Just as SAITA could produce meanings without mental states, so too could GAST produce a meaningful traffic regulation with content that never was "in the head" of any single individual.

Importantly, the communicative content of the various sections of the traffic code did not correspond to the mental states of the members of GAST's Regulation Committee. Some Committee Members voted for regulations on the advice of staffers. Others read summaries. And some members even read some of the sections of the regulations. But it was very rare indeed for a Regulation Committee member to read the whole text of a new regulation. Given these circumstances, no one could think that the meaning of the regulations could be understood as corresponding to the mental states of the members. Those mental states simply did not exist. Rather, GAST produced meaning in much the same way that SAITA did—through a complex algorithmic process. Of course, SAITA functioned without human parts, whereas GAST operated through the actions of individual human agents. But so far as meaning was concerned, GAST and SAITA were essentially identical— they operated on the basis of complex rules that were produced by testing the effect of various regulations on traffic flow.

VI.   NATURAL MEANING AND ARTIFICIAL MEANING

Consider again the distinction between natural meaning and artificial meaning. We are using these two phrases in stipulated senses. "Natural meaning" is meaning produced by a natural person. The folk theory of meaning seems capable of accounting for natural meaning, because natural persons have mental states and their use of language can be explained by the role that their mental states play in the production of their speech and writing.

"Artificial meaning" is meaning produced by entities that are not natural persons. One example of artificial meaning is provided by communications that are generated by AIs. Today, such meanings are produced by the complex algorithms that produce the various texts generated by Amazon.com or Siri. Even today, these AIs produce meanings that cannot be reduced to the meaning generated by the programmers (and others) who have input into the texts produced by Amazon.com or the oral communications generated by Siri. These AIs say new things because their algorithms combine words into sentences that have never been uttered before.

Artificial meanings can be produced by AIs, but there is another common source of non-natural meaning. Artificial meanings can be produced by groups. In Version 2.0 of the Chinese Intersection thought

experiment, GAST was modeled on the algorithmic processes of an AI. In the actual world, the groups that produce artificial meanings are frequently what the law calls "artificial persons." Such artificial persons include business corporations, municipal corporations, school districts, state governments, and the governments of nation states.

These artificial persons speak and write. Sometimes, the communications produced by an artificial person are identical to the communications of a single natural person who acts on behalf of the group. In an absolute monarchy, the king or queen might speak on behalf of the nation state. But in complex modern societies, it is frequently the case that the communications of artificial persons are produced by group agents. Thus, the text of a contract entered into by a corporation may be drafted by a committee of employees and approved by another group of managers. The text of a municipal ordinance may be drafted by a committee of staffers, modified by a subcommittee of the city council, and approved by a majority vote of the council as a whole. The text of a federal statute may be drafted by an industry group, modified by the staff of a congressional committee, modified again in a committee markup session, and then approved by a complex process of committee votes, votes of the House and Senate, and finally signature by the President.

VII. LEGAL MEANING AS ARTIFICIAL MEANING

And this brings us to the law. In the usual cases, legal meaning is produced by group agents. Constitutions, statutes, regulations, ordinances, rules of procedure, and even judicial opinions are usually group efforts. There are exceptions. The meaning of a single judge's ruling from the bench is a natural meaning—produced by a single natural person through the operation of that natural person's mental states. And some judges write their own opinions or simply sign off on opinions that are the work product of a single law clerk. But these are the exceptions that prove the rule. Most legal communications and almost every constitution and statute are produced by group agents. Teams of lawyers write contracts. Teams of legislative staffers (or lobbyists) write statutes ratified by a complex system of voting rules. Constitutions are written by special conventions or parliamentary committees. The law is full of artificial meaning.

Legal interpretation is (usually) the parsing of artificial meanings. Grasping these meanings is not a matter of inferring the mental states of a particular individual or group of individuals. When it comes to group agents, mental states play a role in the production of artificial meanings, but the meanings themselves cannot be reduced to those mental states.

The fact of irreducibility is demonstrated by the Chinese Intersection thought experiment. Version 1.0 shows that we can have meanings without mental states. Version 2.0 shows that a group agent can produce meanings that do not correspond to the mental states of the individuals who make up the group.

The Chinese Intersection Versions 1.0 and 2.0 are just thought experiments. SAITA and GAST are fiction, not fact. But group agents are fact and not fiction. Complex artificial persons like Congress or Microsoft produce artificial meanings in the real world. They are able to do this because their actions are structured by complex rules that enable many individuals to produce texts to communicate effectively, even if there is no single individual who intended that communicative effect.

Of course, there is no guarantee that artificial meanings will be good meanings. The complex group processes that produce a statute might produce a well-functioning regulatory regime, but they might also produce a text with obscure, ambiguous, and seemingly contradictory provisions. Some statutes work; others fail. But this is also true of natural meanings. Natural persons are just human beings. Some of us write clearly, others obscurely. Given the complexity of modern law, reliance on artificial persons to produce legal meaning seems inevitable. The Internal Revenue Code and the associated regulations operate as a complex structured whole, but no human being could possibly produce that text in a single lifetime—and even if that were theoretically possible, no one thinks that wholesale reform of the Code would best be accomplished by a single author.

CONCLUSION: THE ARTIFICIAL MEANING OF THE LAW

Today, the artificial meaning of the law is produced by group agents, and these groups are composed of natural persons. This is a familiar state of affairs, although it gives rise to a well-known set of problems in the theory of legal interpretation. We debate about the original intentions of the framers. Some legal theorists argue that the framers' intentions cannot provide the meaning of the constitutional text, because different framers had different mental states. Similar arguments are made about the intent of Congress and the meaning of federal statutes. The problems are well-known, but they are not insurmountable. Group agents communicate successfully. Our puzzlement about group intentions does not translate into an inability to grasp the artificial meaning of the texts produced by group agents.

Perhaps not tomorrow, but someday, we may face similar problems with the interpretation of legal texts that are produced by AIs rather than group agents. We cannot yet imagine a plausible scenario where the

Constitution of the United States is replaced by a text written by an AI. But we can imagine the drafting of contracts, wills, and trust instruments by AIs. If the AIs are well designed, the legal texts they write will have clear but artificial meanings. Indeed, we can imagine a day when contracts written by AIs are generally regarded as better than those written by natural persons. The AIs might do a better job than humans of avoiding unintended consequences and accidental ambiguities. When we read these texts, we will inevitably rely on familiar tools. We will assume that words are intended in their ordinary senses (conventional semantic meanings) and that standard forms of grammar and syntax are employed. We will not be tempted to search for idiosyncratic meanings that reflect the mental state of the drafters—because we will know that these mental states simply do not exist.

Indeed, the interpretation of artificial legal meanings produced by AIs is already occurring in small but significant ways. Online orders may create binding contracts, and some of the terms of some of these contracts are generated by AIs. Thus, the pricing of airline fares by algorithms (rather than humans) already provides part of a legally binding agreement. We have no difficulty understanding these terms— even though they do not reflect the mental states of a human being.

As time goes on, it seems likely that the proportion of legal content provided by AIs will grow in a fairly organic and gradual way. Indeed, the first time a human signs a contract that was generated in its entirety by an AI, the event might even escape our notice. It seems quite likely that our parsing of artificial meanings generated by AIs will simply be taken for granted. This will be no accident. Today, our social world is permeated by artificial legal meanings. Indeed, we can already begin to imagine a world in which the notion of a legal text authored by a single natural person begins to seem strange or antiquated.

Our world is already inhabited by AIs. Our law is already composed of artificial meanings. The twain shall meet.