

LAWYERS, GUNS AND MONEY:

Content Contextualism and the Cognitive Foundations of Statutory Interpretation¹

Gary Blasi²

Professor of Law, UCLA Law School

LAWYERS, GUNS AND MONEY:.....	1
I. INTRODUCTION.....	3
A. Law, Time and Category	3
B. "Lawyers, Guns And Money," Pornography and Heat.	15
II. ELABORATING THE TEST CASE: ASSAULT WEAPONS IN CALIFORNIA	18
A. A Primer On Guns: The Science And Technology Of Death (and Sport).	20
B. Categorizing Assault Weapons.....	22
III. HARDER CASES AND INTERPRETIVE THEORY.....	26
A. New Textualism.	27
B. Other Interpretive Theories: Intentionalism at the Core	28
C. Reexamining Easy Cases.	29
IV. LEGISLATIVE CATEGORIZATION AS A CONSTRAINT SATISFACTION PROBLEM	30
A. The Costs of Precision: Categorization, Diagnosis and Signal Detection Theory.....	35
B. Constitutional Constraints: Void-for-Vagueness, Separation of Powers and the Non-delegation Doctrine	42
C. Economy	45
D. Openness	48
E. Fitness to the Problem.....	48
V. THE MACHINERY AT HAND: CATEGORIZATION AND HUMAN COGNITION.....	49
A. Classical and Folk Theories of Categorization.....	50
B. Cognitive Scientific Theories of Categorization	52
VI. IMPLICATIONS/CONCLUSIONS	85
A. Better Legislative Categorizations Through Science?.....	86

¹ Draft 3.31 Copyright by author. Not for citation or attribution without express permission

² I am grateful to colleagues at both the UCLA and Stanford Law Schools for encouragement and correction, especially to Paul Brest, George Fisher, Joe Grundfest, Pam Karlan, Dan Lowenstein, Miguel Mendez and Richard Steinberg. This article also reflects the hard work and creative insights of research assistants Michael Gelfond, Sinai Megibow, Amnon Siegel, Nicole Tellem. The staff of the Hugh and Hazel Darling Law Library at the UCLA Law Library provided the usual superb assistance. All remaining errors are mine. The research was supported by generous grants from the UCLA Academic Senate and the Dean of the UCLA Law School.

B. Implications for Theories of Statutory Interpretation.89

ABSTRACT

The field of statutory interpretation is one of central importance to both lawyers and judges, perhaps even more central to their daily work than the analysis of appellate opinions. As a field of academic inquiry, however, the field has become rather stagnant and seems now at a stalemate between contending schools of thought, with most siding against the pure forms of textualism sometimes associated with Justice Scalia and arguing for some form of *contextualism*. What kinds of context should matter is disputed. Thus far, however, scholars have paid remarkably little attention to one crucial contextual factor: What is the statute *about*? What domain of human activity does the law seek to regulate? Justice Scalia urges courts to attend to the plain language of a statute -- any statute -- in order to encourage legislators to clearly say what they mean. This argument is easier to sustain in substantive areas where great precision is obtainable. But should legislatures be barred from acting in substantive areas where precision is very difficult? Legal scholars have acknowledged, then turned away from, this question. This is so in part because scholarship in this area has not thus far taken account of advances in cognitive science and communications theory. In this article I explain how the cognitive science of categorization, along with signal detection theory and complexity theory, allow us to compare substantive domains according to the degree of difficulty in legislating in them, by establishing a metric for the *theorization* of a substantive domain.

The implications of this approach extend well beyond informing the process of drafting legislation. The theoretical foundations of statutory interpretation depend on unspoken, and often incorrect, assumptions about the possibilities of precision in crafting statutes. Once statutory interpretation is understood as an inevitably human process, relying on the tools of human cognition and categorization, the field of statutory interpretation itself might be reconstructed on a more solid, even scientific, foundation.

I. INTRODUCTION.

A. Law, Time and Category

Law exists in time. The very notion of law is *ex ante*. Law says what should, upon certain conditions, happen in the future. Laws are generally applied *ex post*. Certain things having now occurred, what consequence does law require? These notions lie near the core of the concept of law,³ and near the boundary between law and raw power. In law we calculate consequence before identity, the rules before the game.⁴ If we want to write an election code that specifies whether a "dimpled chad" should be recorded as a vote, we recognize that such rules are best crafted before anyone can tell how they might affect a given election. There are powerful philosophical arguments for why such arrangements are morally coherent.⁵ For now, we note that most statutes operate *in futuro*, taking the form, "In the event X should happen, Y should follow."⁶

There is controversy among scholars about what legislatures are doing when they enact statutes; indeed, the entire field of "legisprudence" is devoted to the subject. As an empirical matter, however, legisprudence scholarship is more concerned with *what* legislatures are doing than *how* they are, or should be, doing it.⁷ Rather than engage these questions, I will adopt a view of legislation that is as commonplace in the world in which

³ This is concededly but one, positivist notion of the meaning of *law*. *Law* may also have expressivist and other functions, but in the real world of legislators, citizens, bureaucrats and judges, particular laws also have the more mundane positivist functions I ascribe here.

⁴ There are, of course, also laws that operate *ex post*: for example, reparations and amnesty laws. But a core notion of law itself, as opposed to particular statutes, is that it operates generally and *in futuro*. These principles find expression in the prohibition against *ex post facto* laws and bills of attainder.

⁵ This is the point of Rawls' "veil of ignorance." JOHN RAWLS, A THEORY OF JUSTICE, 136-42 (1971).

⁶ Democratic law thus departs from the democratic ideal in at least two respects. We are all subject to laws made by necessarily imperfect representatives. We are also all subject to laws made by people for whom we had no opportunity to vote for the simple reason that they were dead before we entered the world. Although each new assembly might, in theory, wipe the statute books clean and begin anew, in every real democratic polity, people are governed by "the dead hand of the past." See, e.g., Michael W. McConnell, *Textualism and the Dead Hand of the Past*, 66 GEO. WASH. L. REV. 1127 (1998). Once again, the cognitive limitations of real people constrain the ideal: no legislature comprised of human beings could actually reconsider the existing body of applicable law.

⁷ For a thoughtful, though now somewhat dated survey, see William N. Eskridge, Jr. & Philip P. Frickey, *Legislation Scholarship and Pedagogy in The Post-Legal Process Era*, 48 U. PITT. L. REV. 691(1987). See also, ch. 6 in the Eskridge and Frickey casebook, WILLIAM N. ESKRIDGE, JR. & PHILIP P. FRICKEY, CASES AND MATERIALS ON LEGISLATION: STATUTES AND THE CREATION OF PUBLIC POLICY, (2d ed. 1995).

statutes are written as it is controversial in the world in which they are interpreted and critiqued. I will assume that statutes are communicative acts of legislatures: messages sent forward in time.

Those messages have multiple audiences, including ordinary citizens, administrative agencies and judges. Most scholarship about legislation and interpretation focuses on the judicial audience.⁸ But plainly the fact that law must speak to ordinary citizens introduces an additional set of constraints, some of them deeply rooted in constitutional principle. For example, although a criminal statute may speak with sufficient precision to specialists and judges, it may nonetheless be declared "void for vagueness" as giving insufficient guidance to "ordinary people."⁹ Although I will attend in due course to the constraints introduced by the fact of the multiple audiences for law, I will focus primarily on legislation as communication between legislators and judges. As an empirical matter, it may well be that vagueness or ambiguity in statutory language is the intended result rather than a problem, as Joseph Grundfest and A.C. Pritchard have argued.¹⁰ No doubt this is sometimes the case. For purposes of this article, however, I will assume that on occasion legislatures attempt to say what they mean with as much precision as possible, intending there to be as little variance as possible in how a statute will be interpreted by judges.

One way to think about law-as-communication¹¹ is to invoke as an alternative a new form of political junket: time travel. If, rather than

⁸ A notable exception is Edward L. Rubin, *Law and Legislation in the Administrative State*, 89 COLUM. L. REV. 369 (1989), emphasizing the degree to which modern statutes are intended to speak to administrators rather than ordinary citizens or judges. ,

⁹ *Kolender v. Lawson*, 461 U.S. 352, 357 (1983).

¹⁰ Joseph A. Grundfest and A.C. Pritchard, *Statutes with Multiple Personality Disorders: The Value of Ambiguity in Statutory Design and Interpretation*, 54 STAN. L. REV. 627. Grundfest and Pritchard demonstrate convincingly that Congress achieved almost perfect ambiguity with one statute, as evidenced by the essentially random manner in which the statute was interpreted. Whether intended ambiguity of the sort documented by Grundfest and Pritchard is common is, of course, an empirical question I do not answer here.

¹¹ I hesitate to use this phrase because it has been adopted by scholars working in the traditions of the humanities who use "law" to refer to all talk and writing about legal subjects but who, strangely in my view, virtually ignore the talk and the writing of legislators that matters to both lawyers and citizens: the writing of statutes that carry the force of state power. For example, none of the scholarship collected in *LAW AS COMMUNICATION* (David Nelken ed., 1996) mentions legislators as participants in the communication of, and about, law. I do not discount this enterprise but merely distinguish it. My concerns run to the pragmatic and empirical, in contrast to those highlighted by Nelken in introducing that collection, including: "Can law communicate? Should our definition of legal communication include all communications which refer to law (as Luhmann's does)? Can we communicate with law? Does our thinking about law assume the presence of authority and 'mind'". *Id.* at 15. I have overcome my hesitation in adopting the phrase because I believe my own use of the term coheres rather better with the way

passing statutes at time T1, legislators could simply project themselves, as needed, into the future at time T2, and there decide cases according to their T1 preferences, all problems of both legislation and statutory interpretation would disappear. No communication would be required because legislative preferences would be carried forward in time, in the heads of the time-traveling legislators. No "imaginative reconstruction" of legislative intent of the sort advocated by Judge Learned Hand would be necessary.¹²

In the meantime, legislatures can only pass statutes, which are plagued by problems of at least three kinds: (1) those problems that attend all communication, including ambiguity, vagueness, noise, and risks of accidental misinterpretation; (2) special problems that attend communication in a changing environment about a changing subject matter; and (3) special problems of motivated misinterpretation that exist when those receiving the communication include those with interests that are differentially affected by particular decodings of a given message.

Legal scholarship has attended primarily to the last two problems. Like other forms of communication, statutes encode intention and meaning. Statutory interpretation entails decoding.¹³ Theories of statutory interpretation engage the question of how we, especially the judges among us, ought to decode the meaning of statutes. To be sure, there are objections to law-as-communication. Legislatures are not single-minded entities and can be said to "intend" or "mean" anything only if we loosen the normal usage of those terms, prototypically applied to the utterances of individuals.¹⁴ Nor do many believe that judges are engaged only in a purely technical exercise of decoding the meanings of statutes, acting as the "honest agents" of legislatures. These social facts are decried by some and celebrated by others, who embrace the notion of the judge as "a partner

in which "communication" is most commonly understood, which may or not be its "plain meaning"

¹² [Lehigh Valley Coal Co. v. Yensavage, 218 F. 547, 553 \(2d Cir. 1914\)](#) (L. Hand, J.), cert. denied, [235 U.S. 705 \(1915\)](#). As described by William Eskridge, "Through this imaginative process, the Court seeks to "reconstruct" the answer the enacting Congress would have given if the interpretive issue had been posed directly." William Eskridge, *The New Textualism*, 37 *UCLA L. REV.* 621, 630 (1990).

¹³ Richard Posner, *Legal Formalism, Legal Realism, and the Interpretation of Statutes and the Constitution*, 37 *CASE W. RES. L. REV.* 179 (1987) (analogizing interpretation to a soldier's effort to understand a military order under battlefield conditions when communications have broken down). For a summary of other examples of law-as-communication, see Anthony D'Amato, *Can Legislatures Constrain Judicial Interpretation of Statutes?*, 75 *VA. L. REV.* 561 (1989). See also, Carlos E. González, *Reinterpreting Statutory Interpretation*, 74 *N.C. L. REV.* 585 (1996), for a history of the notion of courts as "honest agents" of legislatures.

¹⁴ This criticism, perhaps by now a commonplace, was cogently made in 1930 by Max Radin, *Statutory Interpretation*, 43 *HARV. L. REV.* 863, 871-72 (1930).

[with the legislative branch] continuing to develop, in what he believes is the best way, the statutory scheme. . .".¹⁵

The objections to the notion of law-as-communication do not extend with equal force to every statute and interpretive problem. In the case of a simple statute about a clearly and commonly understood subject, written in plain language, passed recently by a legislature with an announced and plainly evident intention -- call it an ideal statute -- the interpretations of lawyers, judges and scholars of virtually every interpretive philosophy will converge. In interpreting the ideal statute, we are all originalists, intentionalists, and textualists, drawn to giving expression to the plain meaning of a statute as the drafters originally intended it. Our commitments fade at varying rates with the passage of time and the changing of circumstance, but at the core our normative commitments run to plain meaning.

The real problem, of course, is that many real statutes depart from the ideal: meaning is rarely plain. Leaving aside politics and values (which we exclude to the degree possible in the ideal case), both legislation and interpretation entail solving, with varying success, unavoidable problems of complexity and of communication. With Cardozo, we can imagine an ideal code: ". . . a code at once so flexible and minute, as to supply in advance for every conceivable situation the just and fitting rule," but also recognize, with Cardozo, that ". . . life is too complex to bring the attainment of this ideal within the compass of human powers."¹⁶

Two related kinds of obstacles stand between a legislature and Cardozo's ideal code: problems of understanding and problems of communication. First, in order to "supply in advance for every conceivable situation" a rule, the domain of the statute's operation must be such that the range of "conceivable situations" is understood. A legislature can establish official holidays without difficulty because of the predictable regularities of celestial mechanics that underlie calendars. A judge can determine with some certainty whether tomorrow is Christmas and the courtroom should be dark. We have exceedingly precise means for specifying time, down to the heartbeat of an atom of cesium. In much of life, however, especially in matters of human affairs, we understand far more than we can say in words. We have no equivalently precise measure of the malice in the heart of a scoundrel, but the fact of that malice and its consequences are certainly more real to human beings than the mathematically described resonance of any atom. In acting in domains like

¹⁵ Ronald Dworkin, *LAW'S EMPIRE* 313 (1986)

¹⁶ Benjamin Cardozo, *THE NATURE OF THE JUDICIAL PROCESS* 143 (1921)

these -- the overwhelming majority of areas that concern us -- a legislature has nothing with which to work but words.

Perhaps the most important function of words in legislation is categorization: of objects, events, situations and relationships. There are various ways of categorizing statutes in terms of how statutes themselves categorize. Thus, legal scholars conventionally distinguish between statutes that specify rules and those that describe standards. Other scholars supply supplemental or alternative categorizations, for example: presumptions, factors, guidelines, and principles.¹⁷ The conventional examples are two forms of speed limit: a "rule" specifying a maximum speed limit of 55 miles per hour, or a "standard" prohibiting "excessive speed."¹⁸ One aspect of the distinction is the ". . . extent to which efforts to give content to the law are undertaken before or after individuals act."¹⁹ But as the example makes clear, both rules and standards must be understood as specifying categories, both at the time citizens make decisions and at the time those decisions are assessed by the legal system. A driver seeking to drive slower than 55 miles per hour can rely on a categorical boundary provided by the indicator on her speedometer; the legal system can rely on the instrumentation of the radar or laser. A driver seeking to avoid excessive speed has a much more difficult task in determining the boundary between reasonable and excessive speed -- as does a judge or juror evaluating his behavior later. However one categorizes statutes, it is worth noting that virtually every statute categorizes and fixes boundaries that separate: speeding from lawful driving; murder from manslaughter, and so on.

It could scarcely be otherwise. If law is about things that matter to people, it must necessarily be about categories. Human beings (and other sentient creatures) survive in a chaotic, unpredictable and poorly understood world through categorization. Unmediated by categories, the world confronts us, in William James' memorable description of the infant's world, as "a blooming, buzzing confusion."²⁰ Evolution has enabled us to notice that the large animal of long teeth and tail approaching quietly at dusk resembles the animal that ate our cousin yesterday. While our dog may also categorize the same large cat, only humans (and to some degree, chimpanzees) are able to recognize higher order categories that depend on

¹⁷ Cass Sunstein, *Problems with Rules*, 83 CAL. L. REV. 953 (1995).

¹⁸ Louis Kaplow, *Rules Versus Standards: An Economic Analysis*, 42 DUKE L.J. 557, 560 (1992); Sunstein, *supra* note 14, at 964-65.

¹⁹ Kaplow, *supra* note 15, at 560.

²⁰ William James, *The Principles of Psychology*, 1890, ch. 13, reprinted in 53 Great Books of the Western World 318 (Robert Maynard Hutchins ed., 1952).

relations between other categories -- such as "cousins."²¹ In the absence of law of some kind, relations among the people in large populations or complex societies will also approach the state of James' "blooming, buzzing confusion," even if Hobbes was a bit pessimistic. Felt social norms and instinct will only carry us so far.²² And even norms require some appreciation for category. Neighbors may come to an understanding about how to deal with wandering livestock,²³ but they will need to agree on what constitutes "livestock." And law, whether in the form of imperial decree or modern statute, *must* communicate about categories.

Statutes take the logical form: "If X, Y is the consequence." Statutory interpretation (again, in our idealized world) means determining whether some configuration of events or objects comes within the meaning of X. To cite perhaps the most famous example, is a law prohibiting "vehicles in the park" violated when veterans put an antique tank on display, or a child rides a bicycle through the playground?²⁴ A part of the answer lies in the category label "vehicle" and how one should decode its meaning.

A judge trying to decide whether the tank or bicycle is a vehicle will not want for advice on how to proceed. No fewer than 69 law review articles take note of the problem.²⁵ The general literature on statutory and constitutional interpretation is vast and growing. Battles rage in both appellate opinions and the law journals. Philosophical worlds collide: textualists against purposivists and both against dynamicists; hermeneutics against pragmatism, postmoderns against positivists.²⁶

Linguistics and a small "law and language" scholarly movement have brought some overdue clarity to these conversations. Larry Solan and Peter Tiersma have used linguistic theory effectively to reframe our understanding of interpretation.²⁷ In a series of important articles,

²¹ Recent research indicates that chimpanzees (but not monkeys) can learn categories that depend on relations among other categories. Roger K.R. Thompson & David L. Oden, *Categorical Perception and Conceptual Judgments by Nonhuman Primates: The Paleological Monkey and the Analogical Ape*, 24 COGNITIVE SCI. 363 (2000).

²² And even norms can operate only on the basis of some common understanding of categories. For example, in order for neighboring ranchers to develop norms for dealing with wandering cattle, they must first have a common sense of what constitutes a cow (or heifer, steer or bull). Is the goat eating the laundry on the line subject to understood rules about cattle?

²³ Robert Ellickson, *ORDER WITHOUT LAW: HOW NEIGHBORS SETTLE DISPUTES* (1991).

²⁴ H. L. A. Hart, *Positivism and the Separation of Law and Morals*, 71 HARV. L. REV. 593, 607 (1958).

²⁵ WESTLAW search on JLR file, March 6, 2004.

²⁷ See, e.g., Lawrence M. Solan, *Learning Our Limits: The Decline of Textualism in Statutory Cases*, 1997 WIS. L. REV. 235 (1997) Peter M. Tiersma, *A Message In A Bottle: Text, Autonomy and Statutory Interpretation*, 76 TUL. L. REV. 431 (2001)

Steven Winter has brought to law, especially constitutional law, the cognitive linguistics of George Lakoff.²⁸ A 1995 conference on law and linguistics at Washington University produced an impressive series of articles demonstrating the utility of applying linguistic theory to a wide range of important problems in law in general and statutory interpretation in particular.²⁹

Law and language scholarship has not sought to establish a new interpretivist theory, but rather to bring to jurisprudence what science knows about language and communication. Professor Eskridge is among the rare practitioners of the mainstream jurisprudence of interpretation to make use of the results of scientific linguistics.³⁰ But in my view, no thoughtful person -- of any potential interpretivist persuasion -- can ignore this body of work. So long as the real world of law involves human beings, any useful theory of interpretation must take account the actual use of language.

At the same time, it is worth noting the ways in which law and language scholarship is incomplete. First, like the scholarship on interpretation generally, law-and language scholarship has focused almost exclusively on the problem of extracting meaning from a given legal text, rather than the related but not entirely identical problems of putting meaning *into* text. Generally speaking, as Robert Seidman has written,

The focus of American legal scholarship has failed to follow the shift from appellate decisions to legislation as the principal source

28. Steven L. Winter, *Book Survey: Minding the Law, By Anthony G. Amsterdam and Jerome Bruner*, 99 MICH. L. REV. 1607 (2001); Steven L. Winter, *The "Power" Thing*, 82 VA. L. REV. 721 (1996); Steven L. Winter, *Human Values in a Postmodern World*, 6 YALE J.L. & HUMANITIES 233 (1994); Steven L. Winter, *The Constitution of Conscience*, 72 TEX. L. REV. 1805 (1994); Steven L. Winter, *One Size Fits All*, 72 TEX. L. REV. 1857 (1994); Steven L. Winter, *Confident, But Still Not Positive*, 25 CONN. L. REV. 893 (1993); Steven L. Winter, *Fast Food and False Friends in the Shopping Mall of Ideas*, 64 U. COLO. L. REV. 965 (1992); Steven L. Winter, *The Meaning of "Under Color of" Law*, 91 MICH. L. REV. 323 (1992); Steven L. Winter, *For What It's Worth*, 26 LAW & SOC'Y REV. 789 (1992); Steven L. Winter, *Foreword: On Building Houses*, 69 TEX. L. REV. 1595 (1991); Steven L. Winter, *An Upside/Down View of the Countermajoritarian Difficulty*, 69 TEX. L. REV. 1881 (1991); Steven L. Winter, *Contingency and Community in Normative Practice*, 139 U. PA. L. REV. 963 (1991); Steven L. Winter, *Indeterminacy and Incommensurability in Constitutional Law*, 78 CAL. L. REV. 1441 (1990); Steven L. Winter, *Bull Durham and the Uses of Theory*, 42 STAN. L. REV. 639 (1990); Steven L. Winter, *The Cognitive Dimension of the Agony Between Legal Power and Narrative Meaning*, 87 MICH. L. REV. 2225 (1989); Steven L. Winter, *Transcendental Nonsense, Metaphoric Reasoning, and the Cognitive Stakes for Law*, 137 U. PA. L. REV. 1105 (1988); Steven L. Winter, *The Metaphor of Standing and the Problem of Self-Governance*, 40 STAN. L. REV. 1371 (1988).

29 Conference, *What is Meaning in a Legal Text? Northwestern University/Washington University Law & Linguistics Conference*, 73 WASH. U. L.Q. 769 (1995) et seq.

30 Eskridge and Frickey, *CASES AND MATERIALS ON LEGISLATION*, *supra* note 5, at 642.

of law. We have no theory of legislation to aid in the generation of ideas to guide the legislative process, "no general account of how such statutes should be designed, and what makes them effective or ineffective, desirable or undesirable."³¹

While there are now a handful of counterexamples to Seidman's general proposition, Frederick Bower's *Linguistic Aspects of Legislative Expression* still stands virtually alone in applying even moderately sophisticated theoretical tools (standard linguistics theory) to the problem of the crafting of statutes.

Partly because of emphasis on problems of interpretation, the penetration of the science of linguistics into law has been incomplete. Although linguistics is one of the disciplines that gave rise to cognitive science, linguistics itself has taken limited account of the more recent theoretical and empirical work of cognitive science. Most of the science in law and language scholarship is for that reason rather dated.³² Bower's book relies on conceptions of categorization in linguistics that date to the early 1970's.³³ Steven Winter's work makes effective use of Lakoff's cognitive linguistics, including Lakoff's theory of idealized cognitive models, but Lakoff's theories have not themselves fared very well within cognitive science.³⁴

³¹ Robert B. Seidman, *Justifying Legislation: A Pragmatic, Institutional Approach to the Memorandum of Law, Legislative Theory, and Practical Reason*, 29 HARV. J. ON LEGIS. 1, 3 (1992).

³² Larry Solan, particularly in his use of the mental models theory of Philip Johnson-Laird, is a notable exception. Lawrence M. Solan, *Cognitive Foundations of the Impulse to Blame*, 68 BROOK. L. REV. 1003 (2003); Lawrence M. Solan, Introduction, *A Conference in Celebration of the Publication of Steven L. Winter's Book, A Clearing in the Forest*, 67 BROOK. L. REV. 941 (2002); Lawrence M. Solan, *Causation, Contribution, and Legal Liability*, 64-AUT LAW & CONTEMP. PROBS. 265 (2001); Susan M. Herman & Lawrence M. Solan, *The Jury in the Twenty-First Century*, 66 BROOK. L. REV. 971 (2001); Lawrence M. Solan, *Why Laws Work Pretty Well, But Not Great: Words and Rules in Legal Interpretation*, 26 LAW & SOC. INQUIRY 243 (2001); Lawrence M. Solan, *Convicting the Innocent Beyond a Reasonable Doubt*, 49 CLEV. ST. L. REV. 465 (2001); Lawrence M. Solan, *Refocusing the Burden of Proof in Criminal Cases*, 78 TEX. L. REV. 105 (1999); Lawrence M. Solan, *Law, Language, and Lenity*, 40 WM. & MARY L. REV. 57 (1998).

³³ Frederick Bowers, LINGUISTIC ASPECTS OF LEGISLATIVE EXPRESSION 138-42 (1989)

³⁴ A very rough sense of the relative success of Lakoff's "idealized cognitive models" (ICM) theory can be derived from a quick exercise in bibliometry. In all the psychology journals on the PsycINFO database, only 7 mention the ICM theory in either the title or abstract -- compared to 5289 articles referencing "categorization." (search conducted June 2, 2003) . In comparison, 55 law review articles in the JLR database on WESTLAW mention Lakoff's theory. Of these, 9 were written by the prolific Steven Winter, and many of the remainder cite Winter citing Lakoff. These are not directly comparable numbers in that WESTLAW contains all citations, while PsycINFO is limited to titles and abstracts. Some of the reasons Lakoff's theory of idealized cognitive models has not fared well in psychology and cognitive science are set out in John

In sum, the law journals are virtually devoid of discussion of the problems entailed in the *creation* of the primary texts of law. In comparison to the theoretical riches available to judges and scholars concerned with interpretation, the legislator or lawyer looking for guidance in crafting a statute will find only a handful of articles of comparable sophistication.³⁵ For whatever reasons,³⁶ the vast bulk of sophisticated scholarship about statutory law has been scholarship about interpretation of an existing rather than a potential or intended text.

This essay comes at the problem of law and interpretation the other way around -- from the "sending" or encoding end of law-as-communication. Legislation is, of course, not the only form in which law is communicated. Law is also made by appellate judges, whose opinions also speak to multiple future audiences. I restrict the focus of this essay, however, to the problems of communicating law in statutory form, with

Vervaeke & Christopher D. Green, *Women, Fire, and Dangerous Theories: A Critique of Lakoff's Theory of Categorization*, 12 METAPHOR & SYMBOL 59 (1997). Lakoff's elegant and evocative work on metaphor has found a broader audience, though less so among cognitive scientists. This is not to say that Lakoff is wrong. Indeed, I find his work superior to much of the more conventional cognitive science scholarship in his attention to narrative. He and Jerome Bruner have kept alive a rigorous science-based focus on narrative while much of cognitive science has ignored it.

³⁵ Seidman's essay surely counts as a significant counter example. Other works reflecting significant sophistication in linguistics, philosophy or cognitive theory include foundationally important work by Reed Dickerson, including REED DICKERSON, LEGISLATIVE DRAFTING (1977) and THE FUNDAMENTALS OF LEGAL DRAFTING (1986) (notably, ch. 3: *Drafting and Communication*); aspects of William Eskridge's work on interpretation, especially William N. Eskridge, Jr. & Judith N. Levi, *Regulatory Variables and Statutory Interpretation*, 73 WASH. U. L.Q. 1103 (1995); FREDERICK BOWERS, LINGUISTIC ASPECTS OF LEGISLATIVE EXPRESSION (1989). Works with narrower focus that make a significant general contribution include Julian B. McDonnell, *Definition and Dialogue in Commercial Law*, 89 NW. U. L. REV. 623 (1995) and Steven L. Schwarcz, *A Fundamental Inquiry into the Statutory Rulemaking Process of Private Legislatures*, 29 GA. L. REV. 909 (1995). In using the word "sophisticated," I do not mean to degrade the many works that aim to give practical guidance to the drafters of statutes, guidance in the form of "words to avoid," "how to express time," and so on. See, e.g., G. C. THORNTON, LEGISLATIVE DRAFTING, (4th ed. 1996). I also exclude work by scholars of obvious sophistication in linguistics and other fields whose work has other aims, including improving the language often used by lawyers and judges by encouraging the avoidance of "legalese." See, e.g. PETER TIERSMA, LEGAL LANGUAGE (1999).

³⁶ Professor Seidman offers several hypotheses and suggests research on why scholarship has to so large a degree ignored the problem of legislation. Seidman, *supra* note 28. I would add: Like judges and most lawyers, law professors are generally concerned with the laws that are, rather than those that might be. Many have been clerks to judges, for whom statutory interpretation is part of the job description; relative few have labored in legislative vineyards. To the degree that most law professors envision an audience beyond the academy, they are probably more likely to think of judges than of legislators -- apart from pieces on policy or legal reform most likely to come from the legislative branch. Finally, statutory interpretation has evolved from constitutional interpretation, still regarded as the reputational apex of the academic intellectual enterprise.

particular reference to the problem of specifying categories in legislation. I aim to demonstrate the utility of applying recent work in cognitive science and decision theory to the problems of categorizing in legislation.³⁷

Perhaps the primary benefit of a sustained focus on the problems of encoding of statutory text and a scientific look at the central problem of categorization is the light thereby cast on some current controversies in interpretive theory. The scholarship about what judges should do when they receive these messages is voluminous and varied. Formalists and textualists argue that judges should pay as little attention as possible to anything other than the words of the statute. There are perhaps a dozen species of theory of statutory interpretation that would have judges take account of both text and context. Although these theories differ substantially in the details, all are antiformalist and, in one way or another, contextualist theories. Contextualists differ both about how much attention judges should pay to context and which contexts should matter. Even committed textualists will concede that sometimes context is essential to extract meaning from words, as in deciding whether a statute about "banks" refers to rivers or financial institutions. Others would extend conventional linguistic concepts to bring within the scope of relevant context the particular "interpretive communities" who make sense of statutory language.³⁸ Intentionalists would have judges pay particularly close attention to the circumstances and the legislative process that produced a statute. Others, notably William Eskridge, would have judges interpret statutes "dynamically, in light of their present societal, political and legal context."³⁹ Pragmatists like Richard Posner would have judges attend to the economic and other effects of interpreting statutes in ways that violate "well founded expectations."⁴⁰ Thus, all interpretive theories beyond the purest form of textualism advocate interpretation that takes account of a range of types of contexts, either at the time of statutory enactment, the time of interpretation, or both. Some scholars have abandoned the quest for a universal theory of interpretation, arguing instead for a contextualized choice of interpretive approaches. For

³⁷ As will also become clear, I have not brought to bear all the possible theoretical paradigms originating in the humanities that one might deploy toward explicating the metaphor of law as communication, including especially the work of Habermas. The collection edited by Nelken, *supra* note 8, is a useful introduction to these perspectives.

³⁸ William S. Blatt, *Interpretive Communities: The Missing Element In Statutory Interpretation*, 95 NW. U. L. REV. 629 (2001).

³⁹ WILLIAM ESKRIDGE, DYNAMIC STATUTORY INTERPRETATION, 1479 (1994).

⁴⁰ Richard A. Posner, *Pragmatic Adjudication*, in THE REVIVAL OF PRAGMATISM: NEW ESSAYS ON SOCIAL THOUGHT, LAW AND CULTURE 238, 235-53 (Morris Dickstein ed., 1998).

example, Cass Sunstein and Adrian Vermeule would reframe the question of interpretation as "how should certain institutions, with their distinctive abilities and limitations, interpret certain texts"⁴¹, thus attending to the context of interpretation itself.

Scholars are thus open to considering many kinds of contexts in statutory interpretation. But there are aspects of context rather remarkably absent from this scholarship: First, what is the legislation *about*? That is, what is the domain of human activity the statute seeks to regulate? Second, how well is this substantive domain of legislation understood, either by lay people or by science and experts? Put another way, how *might* a legislature have been more precise had it chosen to do so? In every case the "context" to which contextualist scholars of various species refer is not the substantive domain in which law operates, but the historical, social, and/or legal contexts in which law is created and interpreted.⁴² Although scholars use examples drawn from varied domains to illustrate their points, no conscious awareness of the consequences of the choice of illustration is revealed. Statutes about slum housing, pornography, assault weapons, agricultural subsidies, or tax shelters might be subjected to different interpretive strategies depending on institutional or other contexts, but no scholar has argued that the choice of interpretive strategy should also depend on how well the subject matter of the legislation is understood or theorized, either scientifically or in the general culture, either at the time of enactment or of interpretation.

I make precisely that argument in this article. I propose and explain why one might *simultaneously* share with Justice Scalia an aversion to talk about legislative intent in interpreting statutes about deadlines for filing claims and embrace Professor Eskridge's theory of "dynamic statutory interpretation" when it comes to laws seeking to regulate the production of assault weapons. I argue here that any sound interpretive theory *may* attend to some of the variables of context that have concerned other scholars, but *must* also attend to the state of knowledge about the realm of human affairs in which law operates. As I explain, recent developments in

⁴¹ Cass R. Sunstein & Adrian Vermeule, *Interpretation and Institutions*, 101 MICH. L. REV. 885, 886 (2003).

⁴² For example, Jonathan R. Siegel explains that "Contextualism in administrative law is the interpretation of administrative law statutes in light of background principles of administrative law." *Textualism and Contextualism in Administrative Law*, 78 B.U. L. REV. 1023, 1032 (1998). Others use the term to refer to interpretation that examines the "plain meaning of the statute's language in conjunction with the historic evolution of the statute, along with any other legislative documents that may have accompanied the legislation. . . .According to contextualists, the language of statutes is often ambiguous" and "[a] statute's meaning often depends on its context and purpose," Sande Buhai & Nina Golden, *Adding Insult To Injury: Discriminatory Intent as a Prerequisite to Damages Under the ADA*, 52 RUTGERS L. REV. 1121, 1131 (2000).

cognitive science, decision theory, and other disciplines provide some reasonable means for determining *when* a primarily textualist or intentionalist (or other contextualist) method is more appropriate to the task of interpreting statutory law, consistent with core democratic principles. A key factor in understanding the substantive context of legislation is how well the relevant domains of human knowledge have been *theorized*.

Consider, for example, two social security disability rules specifying the category of persons with what is now called Down syndrome. When my late brother Ricky, who adopted the name "Coach," was born in 1951, the syndrome was understood much as it had been understood when first labeled in 1866. Eight years later, in 1959, science determined that the syndrome was the result of an extra chromosome. Before 1959, the category could only be described by a fairly lengthy set of criteria or rules for specifying the category, or perhaps by reference to exemplar individuals or a prototype. After 1959, the category could be described very precisely: persons with additional chromosome 21, like Coach, have Down syndrome. Of course, a modern legislature might still decide to draft statutes using pre-1959 understandings of Down syndrome. But a modern court should take account of whether the enacting body had other choices, given the state of human knowledge at the time of the legislation.

This focus on the state of human knowledge about the objects of law also provides some insights into fundamental limits of law-as-communication, and thus of law itself. Just as Claude Shannon demonstrated that the flow of any kind of information was limited by the bandwidth of the transmission line,⁴³ I will explore here whether law is also subject to fundamental limits, imposed not by any analogous "bandwidth" of history, but by limitations on capacities for understanding in the fields in which law seeks to act. I will suggest that these limits vary according to how adequately the subjects of legislation have been *theorized* outside of law -- in science or in the general culture. The same approach that makes it possible to speak about theorization in a reasonably rigorous way, also suggests a metric for the intrinsic difficulty of legislating in a given substantive area, and thus the means to think rigorously about not only the pragmatics of creating law, but also of its ultimate limits.

⁴³ Shannon's law describes the theoretical maximum rate at which error-free digits can be transmitted over a bandwidth-limited channel in the presence of noise, usually expressed in the form $C = W \log_2(1 + S/N)$, where C is the channel capacity in bits per second, W is the bandwidth in hertz, and S/N is the signal-to-noise ratio. C.E. Shannon & W. Weaver, THE MATHEMATICAL THEORY OF COMMUNICATION (1963)

Finally, rather more modestly and pragmatically, we might actually learn something useful about the communicative choices of legislatures, and the various costs of pursuing one or another. Clearly, no legislature can actually draft a law in any but the most trivial substantive areas that will obviate the need for interpretation and an interpretive theory. And I realize that sometimes vagueness is the intended consequence; in which case I hope to assist future legislatures in better calibrating their intended vagueness. But I also make a pragmatist's leap of faith, assuming that fairly often legislatures want to say what they mean, and also assuming that drafting statutes is a task of varying difficulty, one that can be done in better and worse ways under given circumstances. I also assume along the way that the progress made in the past two decades in understanding the nature and limits of human communication and cognition more generally may have some relevance to the communication and cognition of legislators and judges.

B. "Lawyers, Guns and Money," Pornography and Heat.

I will use as a continuing example a legislative category of "assault weapons," the specification of which has occupied a good deal of energy of both legislatures and courts, because of an especially volatile mix of "lawyers, guns and money."⁴⁴ The category of assault weapons usefully brings into sharp focus a fundamental problem of legislation: the world is always changing, such that those interpreting our words must apply them to

⁴⁴ "Lawyers, guns and money" are plausibly members of one of two categories. In the original phrase in the late Warren Zevon's 1978 song, the category might be "things to send adventurers in distress":

I was gambling in Havana
I took a little risk
Send lawyers, guns and money
Dad, get me out of this.

[Warren Zevon, Lawyers, Guns, and Money, on Excitable Boy, (Asylum Records 1978)] Readers can listen to the song at <http://morris2k.cti.depaul.edu/zevon/newindex.html> (visited 2/16/04).

A cynic might propose the same as members of the category of "ingredients of bitter and irrational legislative battles".

I am indebted to Professor Peter Tiersma for allowing me to borrow from him the contextual reference to Warren Zevon's song. Professor Tiersma and I found ourselves on the same panel at an academic meeting presenting papers utilizing the assault weapons example to illuminate why it is difficult to write laws. Peter Tiersma, Message in a Bottle: Text, Autonomy, and Statutory Interpretation, 76 TUL. L. REV. 431 (2001) utilizes the same California statute to explore some related issues in statutory interpretation.

a world different than that in which they were uttered. The problem is especially acute in the case of assault weapons because arms manufacturers have actively sought to evade the category, continually modifying weapons to evade the statutory classification. This is a not uncommon phenomenon: attempted restrictions on "tax shelters" produced enormous, categorically evasive creativity by lawyers.⁴⁵ And the more general problem of how to sort millions of objects or situations into a finite number of categories, is universal to legislation. Before turning to "assault weapons" in earnest, it may be useful to calibrate the difficulty of our sample problem by comparing it to two other problems that lie along the spectrum of difficulty.

1. Heat.

Suppose a legislature wants to require that landlords provide adequate heating in apartments within the jurisdiction. A sloppy draftsman -- or one happy to evade the problem -- might simply use words like "adequate heat," leaving the determination of what those words might mean to future landlords, tenants, lawyers, bureaucrats and judges. A more precise definition might reference, for example, "heating facilities capable of maintaining a minimum room temperature of 70 degrees F at a point three feet above the center of the floor in all occupied rooms"⁴⁶ Such a law presents fairly minimal problems in statutory interpretation. To be sure, no thermometer or yardstick is absolutely precise, and there may be odd-shaped rooms in which determining "the center" would be a geometrician's challenge. A bit more trouble is presented by the term "occupied room," because the context begins to move away from physics and toward the complexities of human behavior. In general, however, classifying situations that violate the law from those that do not seems a fairly easy task. A legislature passing such a law in 1920, sending this message toward future judges, can be reasonably confident that it can be interpreted with little difficulty.⁴⁷

⁴⁵ Noel B. Cunningham & James R. Repetti, Textualism and Tax Shelters," 24 VA Tax Rev. (forthcoming, 2004).

⁴⁶ This definition is identical with that in the California Code of Regulations, CAL. CODE REGS. tit. 25, § 34, save for the requirement that the measurement be taken at the "center" of the room.

⁴⁷ Perhaps the most ambitious effort ever undertaken to communicate with precision with future readers was the record placed on the Starship Voyager, in the expectation that beings might try to make sense of it some millions of years in the future. The content of this record is described on the NASA website at <http://vraptor.jpl.nasa.gov/voyager/record.html>.

2. Pornography

Near the other end of the spectrum lies a notoriously difficult problem: how to define "pornography." Like Justice Stewart, we may believe that we "know it when we see it."⁴⁸ But no one has been able to specify this category with very satisfactory precision. The only reliable method of operationalizing Justice Stewart's definition would be to employ Justice Stewart to let us know what he sees, or to find a way to embed Justice Stewart's sensibilities in some form of device or computer program. Yet many people believe that we should not abandon the task of dealing with pornography because of the inevitable interpretive challenges. In the case of pornography, as in many other difficult cases, the courts have explicitly delegated the problem of specifying categorical boundaries to someone else, including future courts, taking account of then "contemporary community standards."⁴⁹ Such a move avoids or delays the problem of categorization, but it does not suggest an answer to the core technical problem: How can a legislature best explain to future Courts, or to future juries applying contemporary community standards, what it *means* by the term "pornography"?

3. Assault weapons, briefly

On first impressions, then, the category of "assault weapons" lies usefully somewhere in the middle of the range of difficulty in problems of categorization. The degree of difficulty is contextual. There are thousands of kinds of weapons. Some ways in which they can be categorized, (for example, as between rifles and shotguns), have been used for hundreds of years and pose few problems.⁵⁰ But the category of "assault weapons" is not one that has been used outside efforts to regulate certain firearms and is one that has no clearly agreed meaning. Moreover, even if a legislature could produce an exhaustive list by manufacturer and model number of banned

⁴⁸ The phrase is most commonly attributed to Justice Potter Stewart's concurrence in *Jacobellis v. Ohio*, 378 U.S. 184, 197 (1964) (Stewart, J., concurring). According to Paul Gewirtz, who has devoted an entire law review article to the origins and use of the phrase, this is "one of the most famous phrases in the entire history of Supreme Court opinions". Paul Gewirtz, *On "I Know It When I See It"*, 105 YALE L. J. 1023 (1996).

⁴⁴ Part of the constitutional test for whether material is pornographic is "whether to the average person, applying contemporary community standards, the dominant theme of the material taken as a whole appeals to prurient interest." *Roth v. United States*, 354 U.S. 476, 489 (1957).

⁵⁰ Rifles fire a single projectile, which spins as it leaves the barrel of the gun because of spiral-cut grooves inside the barrel. Shotguns are normally used to fire many projectiles ("shot") at the same time, and hence have a smooth bore.

weapons, those efforts would be ineffective in dealing with new weapons, some of them created specifically to get around the legal ban. Much as bacteria mutate to avoid particular antigens, gun manufacturers can alter both the names and minor features of their products to keep them legal. For example, one of the weapons used by the killers in the Littleton, Colorado Columbine High School massacre in 1999 was a version of the TEC-9 machine pistol known as the "AB-10", modified specifically to evade federal firearms restrictions. The "AB" in the designation referred to "after ban."⁵¹

At least on first impressions, we ought to have a far easier time specifying "assault weapons" than "pornography." On the other hand, the diversity of weapons and the existence of agents in the market actively seeking to evade categorization complexify the problem. Thus, it appears that the legislative category of "assault weapons" may prove a suitable test case for explicating and evaluating the theoretical points I wish to make.

II. ELABORATING THE TEST CASE: ASSAULT WEAPONS IN CALIFORNIA

The problem of specifying "assault weapons" has probably received the most intense legislative and judicial attention in California. Perhaps the best way of approaching the problem of categorization is to consider the categorizing problem as it was perceived by the California legislature in 1989.

The opinion of the California Supreme Court in *Kasler v. Lockyear* describes the precipitating event, a shooting on the playground of a Stockton, California, elementary school:

While 300 pupils, mostly kindergartners through third graders, were enjoying their lunchtime recess, Patrick Purdy, who had placed plugs in his ears to dull the sounds of what he was about to do, drove up to the rear of the school and stepped out of his car carrying a Chinese-made semiautomatic AK-47. "Impassively, Purdy squeezed the trigger of his rifle, then reloaded, raking the yard with at least 106 bullets. As children screamed in pain and fear, Purdy placed a 9-mm pistol to his head and killed himself. When the four-minute assault was over, five children, ages 6 to 9, were dead. One teacher and 29 pupils were wounded." (Chow,

⁵¹ Paul M. Barrett, Vanessa O'Connell & Robert Tomsho, *Usual Suspect: The Notorious TEC-9 Shows Up Again In High-Profile Killings*, WALL ST. J. Apr. 26, 1999, at A-1.

Slaughter in a School Yard, Time Magazine (Jan. 30, 1989)
p. 29).⁵²

Although perhaps the most shocking, in 1989 the Stockton school shooting was then only the most recent in a series of shooting with assault weapons. Five years earlier, in San Ysidro, California, James Huberty had armed himself with a 9mm semiautomatic pistol, a 12 gauge shotgun and a 9mm UZI semiautomatic rifle and entered a McDonalds restaurant occupied by about 45 people. Huberty fired hundred of rounds, killing 21 and wounding 15.⁵³ And law enforcement officers from across California were reporting that semi-automatic military assault rifles were increasingly the weapons of choice of street gangs.

Not long after the Stockton schoolyard shootings, the California legislature adopted the Roberti-Roos Assault Weapons Control Act⁵⁴ (hereafter, the AWCA), which restricted possession of certain kinds of weapons similar in some respects to the rifle used in the shooting.⁵⁵ The legislature could well have enacted a law with precision equivalent to that of the heating statute; for example, by banning Chinese-made Kalashnikov rifles, or perhaps rifles with a certain sequence of serial numbers. That would not, of course, have accomplished much beyond some market effects in displacing Chinese-made rifles with similar, if not identical, weapons made in Russia or Romania. The legislature might also have adopted a more general and sweeping law, banning all firearms. But that would have gone well beyond the legislature's intention, to restrict particular kinds of weapons associated with these mass killings.

The California legislature's approach to solving this problem entailed both specifying a list of mass-produced weapons, and by empowering a court to determine, after a petition from the California Attorney General and ample notice to the public, that a new weapon should also be included within the "species" of assault weapons. The legislature specified in the process those kinds of changes that would *not* constitute sufficient modification to cause the weapon to fall outside the

⁵² Kasler v. Lockyer, 23 Cal. 4th 472, 483 (2000), 97 Cal. Rptr. 2d 334, 341 (2000). I rely on the generally excellent account of the legislative history of the AWCA in part because my arguments do not rest on the details of this history, but on the general approaches reflected in it.

⁵³ *Id.*

⁵⁴ CAL. PENAL CODE §12275 (1989), Jay Mathews, *California Adopts Ban on Some Assault Guns; Enactment Gives Boost to Other Campaigns*, THE WASH. POST, Apr. 18, 1989; Sherry Bebitch Jeffe, *How the NRA Got Shot Down in California*, L.A. TIMES MAG., July 30, 1989. The circumstances leading to the AWCA are also described at length in Kasler., 23 Cal. 4th at 472.

proscribed category.⁵⁶ The California Supreme Court has upheld both the law and this process of amending the categorization of assault weapons, as against claims that (1) the enumerated list of weapons violated equal protection by being "irrationally under-inclusive," (2) empowering courts to add to the category violated the separation of powers doctrine, and (3) the statute violated due process because the categorization was unduly vague and failed to give fair warning of prohibited conduct.⁵⁷

A. A Primer on Guns: The Science and Technology of Death (and Sport).

The California Legislature had ample reason to act to control the kinds of weapons used in mass shootings. But what "kind" was that? Attention quickly focused on only a subset of the types of weapons involved -- the "assault weapons" -- and *not* on the shotguns and pistols involved. Some of the terms in the preceding sentences will be unclear to some readers, as they may have initially been to some members of the California legislature. A quick primer on firearms is required, both to consider legislation or to understand this essay. Knowledgeable readers can skip forward to the next section.

Guns use small explosive charges in small containers (cartridges) to propel projectiles out of a cylindrical tube (the barrel) at high velocity. Shotguns typically expel some number of pellets; rifles and pistols a single bullet with each discharge. The inner surface of the barrel (the "bore") of a shotgun is smooth; the bore of rifles and pistols has spiral grooves ("rifling") that causes the bullet to spin as it move down the barrel. The effects of bullets on targets, including human bodies, vary according to the laws of physics. The total energy of a bullet is a function of mass and velocity at impact: the higher the energy, the more potential damage. Damage is also affected by whether the bullet deforms as it passes through tissue, and whether the bullet is traveling at a high enough velocity to create shock

⁵⁶ CAL. PENAL CODE §12276.5 empowers the Superior Court to determine that a new weapon not on the proscribed list is nevertheless an assault weapon if it is "identical to one of the assault weapons listed in those subdivisions except for slight modifications or enhancements including, but not limited to: a folding or retractable stock; adjustable sight; case deflector for left-handed shooters; shorter barrel; wooden, plastic or metal stock; larger magazine size; different caliber provided that the caliber exceeds .22 rimfire; or bayonet mount. The court shall strictly construe this paragraph so that a firearm which is merely similar in appearance but not a prototype or copy cannot be found to be within the meaning of this paragraph."

⁵⁷ Kasler v. Lockyer, 23 Cal. 4th at 478.

(sound) waves. Bullets striking tissue at high velocity do considerably more damage than slower moving bullets.⁵⁸

Firearms also vary according to whether a single pull of the trigger can result in firing more than one bullet. "Fully automatic" weapons (the original "machine guns") fire when the trigger is depressed, and continue to fire until either the trigger is released or all cartridges are discharged. The "rate of fire" of a weapon is generally described with reference to the maximum number of rounds that the weapon could fire in a minute. Fully automatic weapons have a higher rate of fire than other weapons, although semiautomatic shotguns can fire a very large number of projectiles in a short time. The potential lethality of a weapon is thus a function of several things: the rate of fire, the weight, velocity and other characteristics of the bullet fired, the number of rounds held in the magazine and the potential rapidity of reloading.

Beyond lethality other features of weapons have attracted legislative interest, sometimes for reasons that are not obviously connected to lethality, including:

- A "thumbhole stock" or "pistol grip that protrudes conspicuously beneath the action of the weapon" -- configurations that may alter how the gun is normally held by the user, or a "folding or telescoping stock" that makes the gun more compact and easier to conceal.
- "Flash suppressors" -- devices on the end of barrels that reduce the visible light emitted when the gun fires
- "bayonet mounts" -- fixtures for attaching a bayonet, a kind of knife or penetrating weapon to the barrel of the firearm.

It is in the context of the characteristics of the range of modern firearms that various legally significant categories have been established. For example, fully automatic weapons of all kinds and short barreled or "sawed off" shotguns have been subject to federal regulation since they were reputed to be the favored weapons of Al Capone and other gangsters in the

⁵⁸ As described by a legislative witness quoted by the California Supreme Court in *Kasler*, "When a high velocity bullet enters the body, Dr. Wintemute explained, "it starts to 'tumble,' as it moves through the tissue . . . greatly increasing the amount of tissue which is damaged by direct contact with the bullet. Moreover, as this high-velocity missile travels through the tissue, it sends out pressure waves: We've all seen pictures of airplanes breaking the sound barrier, and waves moving away from the plane. The same thing happens as these bullets travel through tissue; these pressure waves ... create what is called 'a temporary cavity' behind the path of the bullet, which may be 10 to 15 times--or even greater--the diameter of the bullet itself. As a result of this phenomenon, these high-velocity missiles can damage or destroy organs, break bones--including the femur, possibly the strongest bone in the body--without ever touching them." *Kasler*, 23 Cal. 4th at 484, citing the testimony of Dr. Garen Wintemute of the University of California, Davis, Medical School before the Committee of the Whole. (1 Assem. J. (1989-1990 Reg. Sess.) p. 447.)

1930's.⁵⁹ Specifying these categories of weapons was and is straightforward. Either a weapon can fire multiple times with one trigger pull or it cannot. Subject to disagreements at the microscopic scale, either the barrel of a shotgun is longer than 18 inches or it is not. But other categories are not so easily defined.

The term "assault weapons" has not been used historically apart from the effort to categorize firearms for the purpose of regulating them. The term "assault rifles," on the other hand, comes from the military environment for which these firearms were developed. The most numerically common weapon in the world today is an assault rifle designed by the Russian Mikhail Timofeevich Kalashnikov:⁶⁰ the Abtomat Kalashnikov 47 or "AK-47" similar to the weapon used by the Stockton schoolyard shooter. The most common version of the "AK-47" fires a bullet 7.62 mm in diameter and of intermediate weight and holds 30 cartridges in a removable magazine. If there is a cultural prototype for assault weapons, it is the AK-47. According to an Internet webpage devoted to this one weapon (itself some evidence of the point), there have been between 30 and 50 million such rifles manufactured, making it the most widely used weapon in the world.⁶¹

As is apparent from the preceding paragraph, like automobiles, firearms are generally identified by manufacturer and model. The technologies of mass production insure that all weapons of identical manufacturer and model are substantially identical with regard to function and performance. But, just as there can be many different variations and generations of Volkswagen Beetles, there can be great variations among weapons given the same general name. Assault rifles modeled on the AK-47 were produced in every country allied with the Soviet Union and were produced in many different variations. Among the millions of such rifles produced, there were at least dozens of actual variations, as well as different designations.

B. Categorizing Assault Weapons

Like most citizens, the average legislator is usually unaware of or indifferent to the kinds of details about firearms just discussed, until there is

⁵⁹ 26 U.S.C. § 5861(d).

⁶⁰ A Russian website, <http://kalashnikov.guns.ru/>, provides details on the history of the weapon and its designer.

⁶¹ AK-47.net website: <http://www.ak-47.net/ak47/akru/ak47.html> [March 6, 2004].

reason for concern. But on February 13, 1989, a month after the Stockton schoolyard shooting, the California Legislature met in an extraordinary session, sitting as a Committee of the Whole, to consider what the Speaker described as a matter of "extraordinary importance" and "immediacy": what to do about assault weapons. As proposed legislation was introduced, legislators took different approaches in specifying the category of affected weapons. As introduced in the Assembly, the legislation affected nine specific makes and models of weapons, such as the "AK-47 semiautomatic assault rifle" and the "Uzi semiautomatic assault rifle."⁶² A majority of the Senate preferred a generic description, such as "[a]ll semiautomatic action, centerfire rifles that accept detachable magazines with a capacity of 20 rounds or more."⁶³

As the legislation moved forward in both houses, the Assembly version was amended to drop the generic description, but to expand the list of affected weapons to about 40 specific rifles. Also included for the first time were certain enumerated semiautomatic pistols and shotguns.⁶⁴ After Governor George Deukmejian, a conservative Republican, indicated concerns about the generic description approach, and support for any version of the legislation wavering at the margins in the Assembly, the generic approach was dropped.⁶⁵

The approach to the problem of weapons mutation also evolved over the course of the legislative process. Initially, the Assembly version of the bill would have created an "Assault Weapons Commission" to decide through an administrative process "whether particular firearms are legitimate sports or recreational firearms" and thus not banned.⁶⁶ This provision was quickly dropped⁶⁷, but it reemerged in another form as part of a compromise with the Governor. Under the new provision, which was incorporated into the law as passed and signed, the state Attorney General was empowered to petition a Superior Court for an order including within the affected category weapons that were roughly identical to weapons on the banned list:

Identical . . . except for slight modifications or enhancements including, but not limited to: a folding or retractable stock; adjustable sight; case deflector for left-handed shooters; shorter barrel; wooden, plastic or metal stock; larger magazine size; different caliber provided

⁶² Assembly Bill 357 (introduced January 25, 1989).

⁶³ Senate Bill 292, as amended in the Senate on January 31, 1989.

⁶⁴ *Kasler v. Lockyer*, 23 Cal. 4th 472, 478-79; Assembly Bill 357, as amended in the Assembly, February 27, 1989.

⁶⁵ *Kasler v. Lockyer*, 23 Cal. 4th 472, 486 (2000).

⁶⁶ Assembly Bill 357 (as amended in the Assembly, February 27, 1989).

⁶⁷ Assembly Bill 357 (as amended in the Assembly, March 2, 1989).

that the caliber exceeds .22 rim fire; or bayonet mount. The court shall strictly construe this paragraph so that a firearm which is merely similar in appearance but not a prototype or copy can not be found to be within the meaning of this paragraph.⁶⁸

In enacting the legislation in 1989, the California Legislature declared its intention, and gave additional evidence of both the social problem and the problems of categorization it was attempting to solve:

The Legislature has restricted the assault weapons specified in Section 12276 based upon finding that each firearm has such a high rate of fire and capacity for firepower that its function as a legitimate sports or recreational firearm is substantially outweighed by the danger that it can be used to kill and injure human beings. It is the intent of the Legislature in enacting this chapter to place restrictions on the use of assault weapons and to establish a registration and permit procedure for their lawful sale and possession. It is not, however, the intent of the Legislature by this chapter to place restrictions on the use of those weapons which are primarily designed and intended for hunting, target practice, or other legitimate sports or recreational activities.⁶⁹

The 1989 statute did not end the controversy about how to categorize and regulate "assault weapons." Peter Alan Kasler sued to invalidate the law. The California Court of Appeal agreed with Kasler that: (1) "the 'list' method employed by the Act [to define "assault weapons"] violates equal protection because it does not rationally distinguish between owners of regulated and unregulated guns who are identically situated with respect to the harm sought to be alleviated,"⁷⁰ and that, by requiring "a judge to legislate," the "add-on" provision of the statute violated constitutionally mandated separation of powers.⁷¹

In apparent response to the Court of Appeal opinion, while the matter remained pending before the California Supreme Court, a bill was introduced in 1998 defining assault weapons more generically by reference to a list of features. It was defeated by one vote.⁷² The next year, however, a decade after the first statute was enacted and effective on the first day of the new millennium, the generic description approach the California Senate had preferred in 1989 was resurrected and added to the

⁶⁸ Assembly Bill 357 (Proposed Conference Report No. 1, May 15, 1989).

⁶⁹ CAL. PENAL CODE §12275.5.

⁷⁰ Kasler v. Lungren., 61 Cal. App. 4th 1237, 1243 (1998), 72 Cal. Rptr. 2d 260, 264 (1998), reversed, Kasler v. Lockyer, 23 Cal. 4th 472 (2000).

⁷¹ Kasler v. Lungren, 61 Cal. App. 4th at 1252.

⁷² James P. Sweeney, *New Assault Weapon Ban Shot Down; Proposed Overhaul Falls One Vote Short in Assembly*, THE SAN DIEGO UNION-TRIB., Apr. 14, 1998, at A-1.

statute. Now a weapon can be banned by virtue of being on the existing list of manufacturers and models of weapons, *or* by being added to that list by a Superior Court judge on petition by the Attorney General, *or* by virtue of having one or more particular features. For example, now included within the category "assault weapons" is:

(1) A semiautomatic, centerfire rifle that has the capacity to accept a detachable magazine and any one of the following:

(A) A pistol grip that protrudes conspicuously beneath the action of the weapon.

(B) A thumbhole stock.

(C) A folding or telescoping stock.

(D) A grenade launcher or flare launcher.

(E) A flash suppressor.

(F) A forward pistol grip.⁷³

Notably, in thus reverting in 1999 to the approach taken in the first drafts and in the Senate versions of the legislation a decade earlier, the California legislature declared that:

It was the original intent of the Legislature in enacting Chapter 19 of the Statutes of 1989 to ban all assault weapons, regardless of their name, model number, or manufacture. It is the purpose of this act to effectively achieve the Legislature's intent to prohibit all assault weapons.⁷⁴

Not long after the new statute took effect, the decision of the Court of Appeal invalidating the 1989 law was reversed in the California Supreme Court, which rejected both the equal protection and separation of powers challenges to the Act. As against the equal protection challenge, the Court held that, "[t]he step-by-step approach adopted here--the list plus the add-on provision -- does not violate principles of equal protection.," citing familiar case authority that a legislature may pursue reform "one step at a time, addressing itself to the phase of the problem which seems most acute to the legislative mind."⁷⁵ The Court also rejected the separation of powers argument, finding that the scheme by which the Attorney General presented candidate weapons for inclusion on the banned list for final adjudication by the judiciary was entirely in keeping with the tripartite system of government.⁷⁶

⁷³ CAL. PENAL CODE §12776.1, effective January 1, 2000.

⁷⁴ Section 12 of Stats.1999, c. 129 (Senate Bill 23).

⁷⁵ *Kasler v. Lockyer*, 23 Cal.4th at 488 (citing *Warden v. State Bar*, 21 Cal.4th 628, 645 (1999), (quoting *Williamson v. Lee Optical Co.*, 348 U.S. 483, 489 (1955), 75 S.Ct. 461, 99 L.Ed. 563)).

⁷⁶ *Kasler v. Lockyer*, 23 Cal.4th at 491-498.

III. HARDER CASES AND INTERPRETIVE THEORY.

On first impressions, it is not easy to see how a legislature could have done a better job at specifying the category of assault weapons. As with all efforts at communication, however, the real test comes at the receiving end. In the scheme enacted in California, the meaning of the assault weapon category can be interpreted by courts in two contexts: (1) A court in a prosecution brought under the statute, deciding whether a particular object possessed by the defendant is a banned weapon; or (2) a Superior Court deciding whether a class of objects should be added to the list, on the petition of the Attorney General.

Consider the following possibilities:

1. The personal flamethrower: A defendant is charged with unlawful possession of a weapon he has himself invented and constructed. Rather than bullets, it emits small globules of burning fluid, propelled by highly compressed gas. The defendant admits possession but demurs to the charge on the ground that his weapon is not an assault weapon? A judge must decide.

2. The "Mini M-16": The Attorney General petitions the Superior Court to add to the list of banned weapons a newly modified version of the military's M-16 rifle. In contrast to the 5.56 mm caliber bullet of the original, the modified weapon fires a bullet of the same weight, but one that is longer and smaller in diameter (5.00 mm). A judge must decide whether to grant the petition.

The results in both cases will depend on the interpretive theory, implicit or explicit, of the court. The court in search of an explicit theory of statutory interpretation will have a great many choices from a voluminous and sophisticated literature, including intentionalism,⁷⁷ modified intentionalism,⁷⁸ new textualism,⁷⁹ public justification theory,⁸⁰

⁷⁷ REED DICKERSON, *THE INTERPRETATION AND APPLICATION OF STATUTES* (1975)

⁷⁸ Richard A. Posner, *Statutory Interpretation - In the Classroom and in the Courtroom*, 50 U. CHI. L. REV. 800, 817 (1983).

⁷⁹ Generally ascribed to Justice Scalia, as in ANTONIN SCALIA, *A MATTER OF INTERPRETATION* (1997). See also, William N. Eskridge, Jr., *The New Textualism*, 37 UCLA L. REV. 621 (1990); Lawrence M. Solan, *Learning Our Limits: The Decline of Textualism in Statutory Cases*, 1997 WIS. L. REV. 235 (1997).

dynamic interpretation theory,⁸¹ and public choice theory.⁸² A full comparative accounting of these choices is unnecessary here. A couple of examples, drawn in perhaps exaggerated form, will suffice to illustrate the range of choices.

A. New Textualism.

"New textualists" like Justice Scalia will likely direct the judges faced with our two "hard cases" to look hard at the words of the statute, with dictionary and calculator at hand. These are not, it turns out, hard cases at all. The defendants in both cases are innocent. The personal flamethrower is not an assault weapon because it is not a firearm. The AWCA applies to "firearms" of various kinds and the California Penal Code defines "firearms" as "devices that . . . expel . . . a projectile . . . by the force of . . . any explosion or other form of combustion." Though a burning hunk of fuel might be a "projectile," and releasing compressed gas might be viewed by some as a kind of "explosion", the AWCA clearly means to include only those explosions that are produced by some form of combustion, evidenced by the phrase "explosion or *other* form of combustion" (emphasis supplied). As educated people have known since the time of Lavoisier, *combustion* entails rapid oxidation or other chemical combination and not merely the expansion of a compressed gas. If legislators had wanted to ban compressed gas weapons, they certainly had the capacity to do so. To include this weapon will encourage lack of linguistic precision and general care on the part of the legislature.

The question of whether to add the "Mini-16" to the banned list is even easier to answer, based on the plain words of the statute. It permits extending the assault weapons designation to weapons of altered caliber, "provided the caliber exceeds .22 rimfire." Simple calculations, performed by anyone not trained in fuzzy math, reveal that 5.00 mm is 00.19685 inches, which plainly does not "exceed" .22 caliber in diameter. If the legislature meant "exceed" in some other sense, they certainly could have said so. Next case.

⁸⁰ See Bernard W. Bell, *Legislative History without Legislative Intent: The Public Justification Approach to Statutory Interpretation*, 60 OHIO ST. L.J. 1 (1999).

⁸¹ See ESKRIDGE, *supra* note 35.

⁸² See Daniel A. Farber & Philip P. Frickey, *Legislative Intent and Public Choice*, 74 VA. L. REV. 423 (1988).

B. Other Interpretive Theories: Intentionalism at the Core

Adherents to virtually all other theories of interpretation will find these harder cases, albeit for different reasons. Traditional intentionalists⁸³ will focus on what kinds of weapons the legislature intended to regulate, perhaps through the exercise of "imaginative reconstruction" of what the enacting legislature would have thought of the weapons at issue.⁸⁴ Purposivists like Hart and Sacks⁸⁵ will direct their attention to the apparent purposes of the statute and then try to determine whether those purposes are best served by including, or excluding, the questioned weapons. As William Eskridge has argued persuasively, at least as a general matter, both the intentions of legislators and statutory purposes are frequently indeterminate. Eskridge would counsel against the "archaeological" projects of intentionalism and purposivism and for a "dynamic" approach: a judge should embrace the inevitable and become a willing, pragmatic collaborator in the project of constructing good law, appropriate to the slightly modified problem presented by the new weapons.⁸⁶

One of the aspirations of interpretive theory accounts for one shortcoming of all these theories. Intent on devising a theory of interpretation that is general in application, theorists have been disinclined to take explicit account of the huge variability in the domains in which statutes operate. For example, statutes vary importantly in age, in the rate of change in the realms to which they pertain, in the degree to which precision is possible, and at what costs. Although illustrative examples are often detailed and concrete, interpretive theory attends in the main to legislation that is unspecified and generic, or represented by exemplars with unexamined features. One of my aims here is to suggest principled ways of taking account of the differences in the realms in which the law operates, and the degree to which those realms are predictable and well understood. Similarly, scholars examining the consequences of varying degrees of rule precision have sharpened the analytic tools we have for considering the

⁸³ E.g., DICKERSON, *supra* note 76.

⁸⁴ ESKRIDGE & FRICKEY, CASES AND MATERIALS ON LEGISLATION *supra* note 5 at 526.

⁸⁵ HENRY M. HART, JR. & ALBERT M. SACKS, THE LEGAL PROCESS: BASIC PROBLEMS IN THE MAKING AND APPLICATION OF LAW, 166-167, 1148 79

⁸⁶ ESKRIDGE, *supra* note 35.

costs and benefits.⁸⁷ These tools do not, however, really provide the means to take account of differences in the realms in which law operates.

C. Reexamining Easy Cases.

We can begin by acknowledging one of the useful insights of interpretivist theory, that legal terms typically have "core" and "penumbral" meanings (in the words of H.L.A. Hart) or exhibit "prototype effects" (as most effectively explained by Steven Winter).⁸⁸ Taken together, the various interpretive theories exhibit a similar core-and-periphery schema. All of the various strains of interpretivist theory converge in the easiest cases. If a unanimous Congress yesterday declared March 14 a national holiday to commemorate Einstein's birth, and specified the date with reference to Greenwich Mean Time as indicated on the atomic clocks of the Naval Observatory, then an entire convention of interpretive theorists would quickly agree as to whether a given event occurred on Einstein Day, even if debates about the meaning and intended consequences of the term "holiday" might engage them for days. Interpretivist opinions diverge primarily with regard to what stance judges should adopt when legislative purposes are ambiguous and intent is ancient, unknowable, or unclearly expressed, where intent and text seem to conflict, or where circumstances and/or applicable normative principles have evolved. Since there seems not much to talk about, not much intellectual energy has been expended in the close examination of easy cases.

Against this instinct, I aim to show in the sections that follow that interpretive theory cannot reasonably proceed without an adequate understanding of how legislatures *can* reasonably communicate about categories, even in seemingly relatively easy cases. I will suggest that a close analysis and the application of science to apparently easy cases may offer insights useful in thinking about the hard cases as well. In order to focus on the aspects of theory I want to develop here, I will thus assume away the major problems endemic to what Eskridge calls "statutory archaeology."⁸⁹ I accept Eskridge's point that both intentionalism and purposivism are incomplete, as traditionally articulated, as either a

⁸⁷ Colin S. Diver, *The Optimal Precision of Administrative Rules*, 93 YALE L.J. 65 (1983); Kaplow, *supra* note 15; Sunstein, *supra* note 14; Russell B. Korobkin, *Behavioral Analysis and Legal Form: Rules vs. Standards Revisited*, 79 OR. L. REV. 23 (2000).

⁸⁸ Stephen L. Winter, *Transcendental Nonsense, Metaphoric Reasoning, and the Cognitive Stakes for Law*, 137 U. PA. L. REV. 1105, 1174 (1989).

⁸⁹ For a reasonable summary of the critiques, see ESKRIDGE, *supra* note 35, at 14-47.

descriptive or a normative theory of interpretation.⁹⁰ Nevertheless, all interpretive theories treat text and intent as at least potentially relevant. For simplicity I will assume that the interpretive theory of any future judge will cause him or her to take account of the legislature's words, intent and purposes, at least when there are few difficulties determining any of these. A reasonable judge, I will argue, should also take account of the situation a legislature may have confronted in trying to solve the problem before it.

By the same logic, I agree with Justice Scalia that we should generally attend closely to the statutory text and whether a more careful legislature might have said more clearly what it meant, in order to provide incentives to legislatures to speak as clearly as they reasonably can. Unlike Justice Scalia, however, I believe that we should take explicit account not only of whether the legislature might have spoken more clearly, but also of the degree of difficulty in doing so in the given area of regulation, and the inevitable costs, as well as benefits, of maximal feasible clarity. The work in cognitive science and decision theory described below provides the means of making these assessments.

IV. LEGISLATIVE CATEGORIZATION AS A CONSTRAINT SATISFACTION PROBLEM

What was the problem confronting the California legislature after the shooting in Stockton? Initially, the problem was likely seen as "How do we prevent THAT from happening again." What was meant by THAT was subject to radically different interpretations. The "problem" exemplified by the shooting quickly came to be framed as how to limit the availability of particular kinds of firearms in California. As the gun lobby frequently points out, there is always a person attached to the finger on the trigger -- in this case an obviously deranged man for whom hindsight would prescribe treatment or incarceration. The Stockton shooting that came to epitomize the "assault weapons problem" might as easily have become emblematic of hate crimes or failures in the mental health and parole systems or inadequacies in security in public schools.⁹¹ How the problem came to be constructed as a matter of the availability of assault weapons is beyond my

⁹⁰ *Id.* at 15.

⁹¹ Patrick Purdy, the shooter, was reportedly motivated in part by his hatred of Asians, and had extensive prior contacts with both the mental health and criminal justice systems in Stockton. *Schoolyard Killer's Hatred of Minorities Told in Report*, SAN DIEGO UNION-TRIB., Oct. 6, 1989, at A-4.

scope here.⁹² Once the problem was framed in terms of regulating a category of firearms, it became what is called in some disciplines a *constraint satisfaction problem*.⁹³

Constraint satisfaction problems entail finding a path to a goal while simultaneously satisfying a number of constraints. Life is full of constraint satisfaction problems. How can I find a restaurant that both my wife and my mother will like, within 30 minutes drive of my home that will be reasonably quiet and not terribly expensive? Crossword puzzles are another simple example: a given letter is subject to the two constraints of fitting into two different English words in both the horizontal and the vertical. What are the constraints on legislation? These will vary according to the context.⁹⁴ It may be that some statutes are passed purely for public relations value, with little concern for how they might be interpreted by citizens or judges. But drafting legislation intended to have any significant effect in the real world on behavior, is invariably subject to multiple constraints, including constraints regarding precision, economy, flexibility, and fitness to the problem at hand.

(a) *Precision*. Laws can be imprecise in different ways. To use a familiar example, an ambiguous statute might refer to "banks," but not indicate whether it pertains to financial institutions or the land adjacent to rivers. A vague statute might refer to "property," but not indicate whether it refers to real property, personal property, intellectual property, or all of them. A precise statute will enable both citizens and judges to determine either directly or through some definite procedure whether an object or behavior is within the legislature's intended category.

Philosophers distinguish varying kinds of imprecision and argue about whether vagueness is a property of language, a property of the world, or a consequence of using language to refer to the world.⁹⁵ Although vagueness

⁹² See, for example, Joel Best, *IMAGES OF ISSUES: TYPIFYING CONTEMPORARY SOCIAL PROBLEMS* (1995).

⁹³ "Formally, a constraint satisfaction problem (CSP) is defined over a constraint network, which consists of a finite set of variables, each associated with a domain of values, and a set of constraints. A solution is an assignment of a value to each variable from its domain such that all the constraints are satisfied. Typical constraint satisfaction problems are to determine whether a solution exists, to find one or all solutions, and to find an optimal solution relative to a given cost function."

Rina Dechter, *Constraint Satisfaction*, (paper available at <http://www.ics.uci.edu/~dechter/publications/r68.html>, visited March 6, 2004)

⁹⁴ Steven L. Schwarcz suggests the following constraints at work in the drafting of uniform state laws by private bodies such as the American Law Institute: clarity, simplicity, flexibility, fairness, consistency and completeness. Schwarcz, *supra* note 32, at 924-925.

⁹⁵ A collection entitled, *VAGUENESS: A READER* (Rosana Keefe & Peter Smith eds., 1996), contains a sampling of the major views on the subject. Particularly for philosophers in the analytic tradition, vagueness poses very important problems. An extensive bibliography on the

is, self-referentially, a somewhat vague concept, we can usefully distinguish between ambiguity, generality, and vagueness. Homonyms (like “bank” in the previous example) are examples of ambiguity. We can typically disambiguate them by adding context: once we know the statute is in the Financial Institutions Code, we are reasonably certain what sort of “bank” it is to which a section of that code probably refers.

As one philosopher has written, “vagueness is ambiguity on a grand and systematic scale.”⁹⁶ In contrast to the finite number of disambiguated meanings of an ambiguous term, vague terms have potentially infinite possible meanings. Consider terms like “tall” or “heap.” Is someone 5’11” tall? Do 100 grains of wheat constitute a “heap”? These situations give rise to the Sorites Paradox, after the Greek word for “heap” (*soros*), which went roughly like this:

1 grain of wheat does not make a heap.
If 1 grain of wheat does not make a heap then 2 grains of wheat do not.
If 2 grains of wheat do not make a heap then 3 grains do not.
...
If 9,999 grains of wheat do not make a heap then 10,000 do not.

Therefore, 10,000 grains of wheat do not make a heap.⁹⁷

Philosophers describe vague predicates as having borderline cases, fuzzy boundaries, and being susceptible to the Sorites Paradox.⁹⁸

The lack of precision in categorization leads to interpretive errors of two possible kinds: errors of inclusion and errors of exclusion. Accordingly, legislators face two different kinds of precision constraints. In order to meet the constraint of *inclusion*, the category specification must include as many as possible of the "kinds" of weapons the legislature wants to regulate (e.g., a firearm that has ". . . such a high rate of fire and capacity for firepower that its function as a legitimate sports or recreational firearm is substantially outweighed by the danger that it can be used to kill and injure

subject can be found at a website described as "an evolving resource" for philosophers concerned with vagueness, at <http://www.btinternet.com/~justin.needle/> (visited March 6, 2004).

⁹⁶ Kit Fine, *Vagueness, Truth and Logic*, reprinted in VAGUENESS: A READER, *supra* note 95, at 136.

⁹⁷ The Sorites paradox was one of seven puzzles attributed to the logician Eubulides of Miletus, a contemporary of Aristotle. An analog was the Bald Man (*phalakros*) paradox, in which grains of wheat were replaced with hairs on a head. TIMOTHY WILLIAMSON, VAGUENESS, 8-9 (1994). Williamson's book is itself something of a landmark in philosophical studies of vagueness. This vast and varied literature goes well beyond my purposes here.

⁹⁸ VAGUENESS: A READER, *supra* note 95, at 4.

human beings.")⁹⁹. In order to satisfy the constraint of *exclusion*, the category specification must *not* include the "kinds" of weapons the legislature does not want to regulate (e.g., ". . . weapons which are primarily designed and intended for hunting, target practice, or other legitimate sports or recreational activities.").

These constraints do not merely reflect policy choices for legislators. They can rise to the level of constitutional principle. In order to comport with the "first essential of due process of law," a category with legal significance must be described in language such that "persons of common intelligence" will not "necessarily guess at its meaning."¹⁰⁰ Thus, given a particular firearm, ordinary people must be able to determine with some certainty whether the firearm is, or is not, an "assault weapon." As I will suggest below, there may be means of specifying the category of assault weapons other than detailed lists that avoid the vagueness constraint, but it is likely that the prefatory language in the statute itself, standing alone, would have been declared unconstitutionally vague.

(b) *Economy*. The category must be defined within parameters of time, space, money and information that constrain real legislators, judges and citizens. It might be theoretically possible for a legislature to develop an exhaustive list of all the individual firearms in the world at a particular moment in time, perhaps by serial number, and to indicate as to whether each is an assault weapon. To do so in practice would require an enormous expenditure of time as well as access to information that no legislature can possibly obtain. Indeed, there is apparently no reliable exhaustive list even of the *nuclear* weapons in the arsenal of the former Soviet Union, to say nothing of the 30-50 (!) million Kalashnikov-style rifles in the world. All means of describing a category of objects other than such a list are less perfect.

(c) *Flexibility*. Even a perfect enumeration would not reach weapons not yet produced, either on preexisting or new designs. Unless the Legislature wants to constantly revisit the definition of the category, it must specify a way in which weapons not currently produced can be determined to be in the category. In this case "not currently produced" may extend both to the potential for minor future variations (shortening a barrel by an inch, for example, or replacing a wooden stock with a plastic stock) and entirely new inventions (e.g., weapons that use compressed gases rather than explosives, or burning globules rather than bullets).

⁹⁹ §12275.5 *supra* note 68.

¹⁰⁰ *Connally v. Gen. Constr. Co.*, 269 U.S. 385, 391 (1926).

(d) *Fitness to the Problem.* For our purposes, we assume that the legislation is intended to have an effect on the problem identified by the legislature. In the case of assault weapons, one can make a case that "the problem" was perceived in one or more of the following ways:

(1) *Positivist/Pragmatic Fitness.* The problem is the unregulated availability of weapons that makes it possible for disturbed individuals to kill dozens of people before the police can stop them. This construction of the problem would cause legislators to focus on the actual, technical lethality of weapons.

(2) *Political Fitness, Cynical Version.* The problem may be that other people, especially voters, perceive these weapons to be a problem and identify particular kinds of weapons, especially those used in the most publicized killings, as those that legislators must address. Problems that confront legislators are always socially constructed, and represent real political problems whether the "real" problem exists or not. Cynics will note that some weapons seem to be included based entirely on appearance, especially similarity in appearance to certain military weapons used in highly publicized instances, rather than lethality or function or anything else that might actually matter.

(3) *Political Fitness, Expressivist Version.* As critics of the law-as-communication metaphor have noted, law has social functions beyond the explicit formulation of categories, rules, and consequences. Law has *expressivist* functions as well. Appearances do matter. From this perspective, banning from civilian use weapons that merely appear similar to military weapons may serve an important expressive function, increasing the symbolic separation between the military and civilian spheres of life and signaling disapproval of would-be civilian Rambos. One need only visit the "hobbies" section of any magazine stand to see the evidence of a subculture devoted to guns, especially military-style weapons. Legislation may express a point of view with regard to this subculture. The California Court of Appeal accepted the expressivist interpretation, holding that the Legislature might take account of difference in appearance, banning "meaner-looking" firearms, but not make distinctions between firearms of identical appearance *and* equivalent function.¹⁰¹

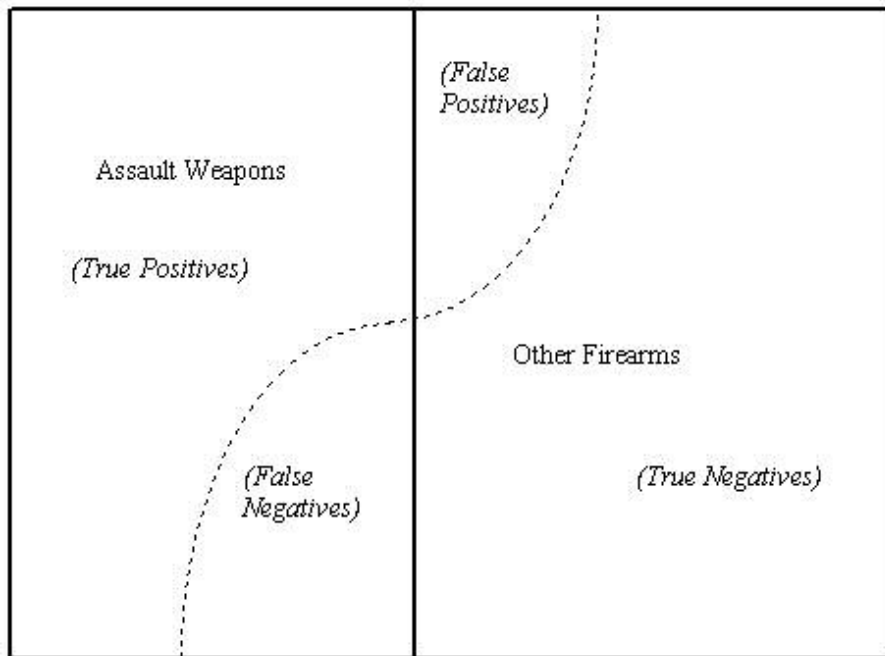
¹⁰¹ *Kasler v. Lungren*, 61 Cal. App. 4th 1237, 1262 (1998), reversed, *Kasler v. Lockyer*, 23 Cal. 4th 472 (2000).

A. The Costs of Precision: Categorization, Diagnosis and Signal Detection Theory

We note initially that some of the constraints on legislation are interrelated. First, because there will always be some uncertainty in classification at the margins, there is an inevitable tradeoff between the constraints of inclusion and exclusion. If the category is specified very broadly, it may be applied to future cases a legislature with perfect foresight would not have included. If specified narrowly, the category will not encompass weapons that were intended to be included. Assume with perfect knowledge and infinite time, a legislature could specify the set of all "assault weapons" perfectly. In other words, the perfect legislature could partition all weapons in the world into two sets: assault weapons and all others. In graphic form, let the vertical line in Figure 1 be the "real" boundary between the set of "assault weapons" and "other weapons" (as determinable with infinite time, resources and foresight). Any categorization in the real world will depart from the ideal. As operationalized by the entire legal system, the result might be as represented by the dotted line in Figure 1, indicating errors of both inclusion and exclusion.

Figure 1

Ambiguity and Two Kinds of Errors



Although the legislature can strive for categorical specifications that will reduce errors of both types, it cannot eliminate them. There are, however, ways of thinking more systematically about the tradeoffs. The two kinds of error (overinclusion and underinclusion) parallel the potential error in any diagnostic regime. Consider, for example, a blood screening test for a particular cancer. The test can only indicate the concentration of a particular chemical compound in the blood. Even leaving aside measurement errors, no such test is absolutely definitive. There will always be both false positives and false negatives, at whatever cutpoint is established. Moreover, there is typically no way to reduce the number of false positives without increasing the number of false negatives.

In diagnostic medicine, the relationships among these variables are described with reference to various statistics computed on the variables in Table 1.

Table 1
Possible Outcomes of a Diagnostic Test

		DISEASE	
		+	-
T E S T	+	True Positive (TP)	False Positive (FP)
	-	False Negative (FN)	True Negative (TN)

The power of the test to detect true positives is called the *sensitivity* of the test, equal to $TP/TP+FN$. The *specificity* of the test, $TN/TN+FP$, is a measure of the power of the test to identify true negatives.¹⁰² "True" and "false" in the case of diagnostic tests is determined with reference to an assumed ideal state of perfect knowledge. In medical diagnosis, the functional equivalent of perfect knowledge is supplied by the passage of time and increasingly more certain tests. Of course, we have no perfect analog in law to the "true diagnosis." Rather, we have postulated as the standard the subjective intentions of the legislature, were they able to be perfectly expressed at one time and perfectly understood at another. With

¹⁰² There are many other measures of diagnostic tests that are commonly computed with regard to diagnostic tests, e.g., the positive likelihood ratio, the negative likelihood ratio, the positive predictive value, and the negative predictive value.

that caveat, by analogy to the medical terms, the *sensitivity* of a categorization indicates how well the categorization will result in the inclusion of cases the legislature would have included, and its *specificity* the degree to which the categorization avoids including instances the legislature would have excluded.

Medical tests, like statutory interpretations, are rarely completely determinate, generating a certain positive/negative decision. But binary decisions must be made: to prescribe a drug or not, to rule for the plaintiff or the defendant. Diagnostic tests typically yield a number in a range, along which a single "cutpoint" must be established in order to yield a binary interpretation. As the cutpoint changes, the probabilities of false positives and false negatives will necessarily vary inversely. Which cutpoint is selected for decisional purposes will depend on how the relative risks of false positives and false negatives are viewed. Risk sensitivity will vary both by interpreter and by situation. In medicine, our willingness to tolerate false negatives varies with the potential seriousness of the disease at issue, and tolerance for false positives with the cost and burdens of a potentially unnecessary course of treatment. There is no perfect analog in statutory interpretation to the concept of "cutpoint," but obviously different statutory interpreters, judicial and otherwise, will have varying interpretations of the same statutory language. In the case of either the diagnostic test or the statutory categorization, the overall risks -- of overinclusion or underinclusion in the case of statutes -- will depend in part on the effectiveness of the legislative categorization.

One measure of the overall effectiveness of a diagnostic test was originally developed to analyze the efficiency of radar systems. Radar operators must decide whether a given "blip" on the screen is an aircraft rather than a flock of geese or electronic noise. The systematic study of such problems is the province of signal detection theory.¹⁰³ Like doctors interpreting tests or judges interpreting statutes, operators must interpret a somewhat ambiguous signal, and like doctors or judges, they may also vary the stringency with which a particular test result or signal is assessed. One way to represent the overall performance of a radar system, a diagnostic test, or a category specification is by plotting a "receiver operating characteristic" curve or "ROC curve."¹⁰⁴ As applied to the specification of categories, for

¹⁰³ A good overview of signal detection theory can be found in Kenneth R. Hammond et al., *Making Better Use of Scientific Knowledge: Separating Truth from Justice*, in JUDGMENT AND DECISION MAKING: AN INTERDISCIPLINARY READER (Terry Connolly et al. eds., 2000).


¹⁰⁴ For a readable introduction to ROC curves and associated concepts, see John A. Swets et al., *Better Decisions through Science*, 283 SCI. AM. 82 (2000). See also, John A. Swets, *Enhancing Diagnostic Decisions*, in JUDGMENT AND DECISION MAKING: AN

every method of specification, there is some probability of both false positives (overinclusion) or false negatives (underinclusion), a probability that varies with the stringency with which the category is interpreted. If we plot the probability of false positives and false negatives as the stringency of interpretation varies, we get the associated ROC curve for the method of specification, as in Figure 2.

INTERDISCIPLINARY READER *supra* note 103. In the former work, Swets indicates that the initials in "ROC curves" stand, somewhat redundantly for "relative operating curves." I adopt the former, given the lower risk of ambiguity.

Figure 2

ROC Curve for a Statute



The diagonal line in Figure 2 represents chance accuracy: a diagnostic test equivalent to flipping a coin or a categorical specification in a language no one can understand. The better the method of categorical specification, the more the ROC curve bends toward the upper left, as the likelihood of true positives and true negatives increases (or, obviously, as the number of false positives and false negatives declines).¹⁰⁵ In Figure 2, the solid curve thus represents a categorical specification that is superior (in terms of eventual errors produced) to that represented by the dotted line. The shape of the ROC curve reflects the relative balance between the sensitivity and specificity of the categorization. The overall effectiveness of a diagnostic test or categorical specification can be reduced to a single number, the area under the ROC curve -- sometimes referred to as the AUC (for "area under curve"). It is worth noting however, that such a single number might obscure differences between categorization methods that vary according to whether they are better at avoiding underinclusion

¹⁰⁵ Because by the probabilities of a false positive and true positive (and true negative and false negative) always sum to unity, by convention we attend only to the curve reflecting true positives and true negatives.

or overinclusion, a fact reflected by the shape of the ROC curve, but not the AUC.

The AUC may nevertheless be a useful single metric of the accuracy of the diagnostic test or, in our case, the degree to which the categorical specification will be interpreted as intended. Put another way, the area under the ROC curve indicates how much of the determination of category membership is left to the discretion of the interpreting judge, agency or citizen. Given the assumptions about democratic theory we have been operating under, the greater the area under the ROC curve, the better the specification. Obviously, a legislature wishing to exercise maximal influence over the future will want to develop methods for specifying legislative categories that will push the ROC curve for the statutory language as far to the upper left as possible, thereby reducing future interpretive errors.

To translate all this back to our continuing problem of how to define "assault weapons," the risk of under-inclusion can be reduced by specifying the category in broader terms, but not without increasing the likelihood of banning weapons the legislature would have not have banned were the question put to it directly. In both the medical and the legal example, no matter how good the diagnostic test or the category specification, there will *always* be some degree of error requiring a decision about the tradeoff between the respective risks of false positives and false negatives, between the risks of overinclusion or underinclusion. The cost of those tradeoffs is entirely contextual and external to the issue of accuracy itself. A falsely positive diagnosis may lead only to unnecessary worry, or it may lead to a dangerous course of unnecessary treatment. Errors in legislative categorization lead to similar considerations: How significant is the unintended banning of a kind of weapon widely used by duck hunters (e.g., semiautomatic shotguns), compared to the risk of leaving untouched a kind of weapon really useful only against crowds of people (e.g., the ominously named Street Sweeper)? Obviously, legislators of different political persuasions and representing differing constituencies will come to different assessments of risks and value.

The final statute adopted by majority votes may obscure these underlying conflicts and concerns. Efforts to create legislative history to guide subsequent interpretations may reveal them, although textualist interpreters may ignore them. Sometimes the legislature will send some interpretive guidance along with the statute, in the form of a statement about presumptions to be applied or statements of legislative intent. These may affect the shape of the statutory ROC curve -- the relative concern with overinclusion or underinclusion -- but they cannot by themselves

increase the area under the curve -- the precision of the statutory categorization.

It remains to be seen whether merely mapping the categorizing decisions of legislators into the language of science has any particular utility. Certainly, legislators operating purely on the basis of naïve theories of categorization and on common sense have long been thinking about risks of overinclusion and underinclusion. What are the consequences of inadvertently banning a "recreational firearm," as against the consequences of failing to ban an especially dangerous "assault weapon"? Clearly, the question can be thought about without ROC curves. But one benefit of the scientific framing of the issue is that it might facilitate being able to consider all the risks of error simultaneously.

The psychological sciences tell us that we may perceive the situation differently according to how we *frame* such questions of risk.¹⁰⁶ Precisely how these phenomena might play out in the instant context is uncertain. For the moment, it is perhaps enough to contrast the categorizing inclinations of two otherwise identical legislators, one thinking most intently on the possibility that a slightly modified assault weapon will escape regulation and be used in another Stockton schoolyard massacre, and the other legislator thinking primarily of the possibility that the innocent hunters in his district will be caught up in inadvertent regulatory excess. Our intuitions and political experience, as well as behavioral science, suggest that both values and sensitivities to different kinds of risk will affect how the two kinds of risk are evaluated.

The weighing of such risks is embedded throughout law. The criminal law must be constructed so as to take account both of the possibility of convicting the innocent and acquitting the innocent. A system that actually seeks to follow Blackstone's ratio -- "better that ten guilty persons escape than that one innocent suffer" -- will look quite different with the acceptable tradeoff set at a different number.¹⁰⁷ Science can help with thinking about how to weigh the risks,¹⁰⁸ and it can help us

¹⁰⁶ Much of this work originated with Daniel Kahneman & Amos Tversky. See, , Daniel Kahneman & Amos Tversky *Prospect Theory: An Analysis of Decision under Risk*, *ECONOMETRICA*, 47 (1979); Daniel Kahneman & Amos Tversky, *On the Psychology of Prediction*, 80 *PSYCHOL. REV.* 237 (1973); Daniel Kahneman & Amos Tversky, *Choices, Values, and Frames*, 39 *AM. PSYCHOLOGIST* 341 (1984); Amos Tversky & Daniel Kahneman, *The Framing of Decisions and the Psychology of Choice*, 211 *SCI.* 453 (1981).

¹⁰⁷ Terry Connolly, *Decision Theory, Reasonable Doubt, and the Utility of Erroneous Acquittals*, in *JUDGMENT AND DECISION MAKING: AN INTERDISCIPLINARY READER*, *supra* note 103 at 229, 232; 4 WILLIAM BLACKSTONE, *COMMENTARIES* 358 (9th ed. 1783); See also, Lawrence M. Solan, *Refocusing the Burden of Proof in Criminal Cases: Some Doubt About Reasonable Doubt*, 78 *TEX. L. REV.* 105, 110 n 22 (1999).

¹⁰⁸ For example, Connolly, *supra* note 107, uses formal Decision Theory, based on notions of

reduce those risks, but it cannot decide the ultimately normative and political questions that inevitably follow the existence of unavoidable uncertainty in applying law to life.

With regard to assault weapons in California, it is clear that the calculation of these risks by the California legislature (that is, different legislatures) changed between 1989 and 1999. In adopting a more general categorical specification in 1999, the California legislature displayed more willingness to risk unintentionally banning some "sporting weapons" in order to achieve the goal of "banning all assault weapons." What the legislature could not do was simultaneously eliminate the risks of both over and under-inclusion. But it might have better satisfied these dual constraints by taking advantage of the science of categorization, as I suggest below.

B. Constitutional Constraints: Void-for-Vagueness, Separation of Powers and the Non-delegation Doctrine

There were other constraints on the California legislature imposed by the Constitution. Here, I refer not to the Second Amendment¹⁰⁹ but to two doctrines deeply embedded in different parts of constitutional law: due process and the separation of powers in a tripartite system of government. The due process clause of the Fourteenth Amendment has long been interpreted as a requirement that the legislature pass penal laws that "define the criminal offense with sufficient definiteness that ordinary people can understand what conduct is prohibited and in a manner that does not encourage arbitrary and discriminatory enforcement."¹¹⁰ Regardless of where the legislature determines to draw the boundary between "assault weapons" and "lawful weapons," if it does so in a sufficiently indeterminate way, it risks a court declaring the statute "void for vagueness". Vagueness is itself a vague concept, which encompasses concepts philosophers would call

expected utility theory, to analyze the consequences of setting different evidentiary thresholds for conviction in criminal cases. Using no formal theory, Colin Diver presented a thoughtful analysis of the tradeoffs embedded in what he called "rule precision," in Diver, *supra* note 87. Notable recent efforts along the same lines are Sunstein, *supra* note 87, and Kaplow, *supra* note 87.

¹⁰⁹ Eugene Volokh, *The Commonplace Second Amendment*, 73 N.Y.U. L. REV. 793 (1998); Eugene Volokh, *The Amazing Vanishing Second Amendment*, 73 N.Y.U. L. REV. 831 (1998).

¹¹⁰ *Kolender v. Lawson*, 461 U.S. at 357 (1983)

ambiguity, vagueness and contestability.¹¹¹ As a constitutional matter, the notion of indeterminacy embodied in "void for vagueness" doctrine encompasses concerns in how statutes may be construed by three audiences: citizens, the police and courts.

Accordingly, a legislature sensitive only to risks of unintended exclusion could not constitutionally ban possession of "deadly weapons" as a means of insuring that "assault weapons" are banned, because "deadly weapon" is a radically indeterminate concept. An astonishing variety of objects and instrumentalities have been used to inflict serious injury or death, and there is no indication that human creativity and/or depravity in this regard has been exhausted.¹¹² In the face of such a law, a citizen would not know whether possessing a brick, a baseball bat or sharp pencil put her at risk. A police officer would be granted enormous discretion in deciding whom to arrest. A court would have similarly unbridled discretion in deciding whether any violation had occurred.

One way the legislature might reduce the uncertainty with which the words of a statute are viewed by the citizenry is to delegate the detailed removal of uncertainty to another branch of government at a future point in time: an administrative agency of the executive branch or, as in the case of the California assault weapons law, to the executive (the Attorney General, who may file a petition to add a weapon) and to the judiciary (a Superior Court judge, who may grant or deny the petition). This may, however, run afoul of other principles deriving from the tripartite structure of government: the separation of powers and the non-delegation doctrine.

In theory, the non-delegation doctrine limits the power of the legislative branch to delegate its legislative function to any non-legislative entity, but in practice, it especially limits administrative agencies and the executive. As Bernard Bell summarizes the argument for using the nondelegation doctrine to limit vagueness in statutes:

. . . [The] remarkable vagueness of contemporary statutes results from one of two causes: (1) elected representatives' efforts to avoid responsibility so as to ensure their own re-election, which critics view as a political pathology or (2) disagreement among a legislative majority, in which case, critics urge, legislation should remain

¹¹¹ Jeremy Waldron, *Vagueness in Law and Language: Some Philosophical Issues*, 82 CAL. L. REV. 509, 512-513 (1994). Waldron's piece is part of a small but very useful Symposium, entitled *Void for Vagueness*.

¹¹² For example, California courts have determined that some dogs and some shoes are "deadly weapons" under the "assault with a deadly weapon" statute (CAL. PENAL CODE 245), *People v. Nealis*, 232 Cal. App. 3d Supp. 1, 283 Cal.Rptr. 376 (1991) (sufficiently large attack dogs); *People v. Aguilar*, 16 Cal. 4th 1023 (1997) (hobnailed or steel-toed boots).

unenacted. Neither avoidance of responsibility nor the passage of legislation despite fundamental disagreements accord with the principles of liberal democracy.¹¹³

Plainly, democratic theory would not countenance a statute empowering a federal bureaucracy or a court to "adopt such rules in furtherance of the common welfare as it sees fit." As with vagueness, the question is not whether, but how much, is left to future determination. The U.S. Supreme Court has taken a rather expansive view, requiring only that Congress "lay down by legislative act an intelligible principle to which the person or body authorized to [act] is directed to conform." *J.W. Hampton, Jr., & Co. v. United States*, 276 U.S. 394, 409 (1928). As Justice Scalia noted,

In the history of the Court we have found the requisite "intelligible principle" lacking in only two statutes, one of which provided literally no guidance for the exercise of discretion, and the other of which conferred authority to regulate the entire economy on the basis of no more precise a standard than stimulating the economy by assuring "fair competition." See [*Panama Refining Co. v. Ryan*, 293 U.S. 388, 55 S.Ct. 241, 79 L.Ed. 446 \(1935\)](#); [*A.L.A. Schechter Poultry Corp. v. United States*, 295 U.S. 495, 55 S.Ct. 837, 79 L.Ed. 1570 \(1935\)](#). *Whitman v. American Trucking Ass'ns, Inc.*, 531 U.S. 457, 475 (2001).

Once a statute passes the test of minimal determinacy, the Court has established in *Chevron*¹¹⁴ a preference that the federal courts defer to the interpretive capacities of administrative agencies. But the vagueness and non-delegation doctrines establish constraints of minimal determinacy for statutes, a constraint that varies with context, including the different audiences to which statutes speak.

Delegating greater specificity to an administrative agency or court may relieve the citizen from the uncertainties of a vague statute, but leave to the agency or court so much discretion that the agency or court is exercising the prototypically legislative function. The risks involved are analogous: will a court or agency be so unconstrained by the language of the statute that its decisions effectively constitute legislating? In the *Kasler* case, the Court of Appeal found that the "add-on" provision of the AWCA "requires

¹¹³ Bernard Bell, *Dead Again: The Nondelegation Doctrine, the Rules/Standards Dilemma and the Line Item Veto*, 44 VILL. L. REV. 189, 197-98 (1999)

¹¹⁴ *Chevron U.S.A., Inc. v. Natural Res. Def. Council, Inc.*, 467 U.S. 837 (1984).

a judge to legislate."¹¹⁵ As noted earlier, the California Supreme Court disagreed, emphasizing that most of the decision-making in the "add-on" provision would rest with the Attorney General, whose decisions were subject to review by the courts, which played only a "very narrow, essentially adjudicatory" role in the determining of whether a new weapon was added to banned list.¹¹⁶ The reasoning of the Court provides some basis for thinking that the California assault weapons statute might serve as a useful model for other statutes seeking to regulate an actively evolving subject matter, while preserving the requirement of fair notice to the citizenry.

C. Economy

Although those closest to the issue might be inclined to think otherwise, dealing with the Stockton shooting was not the only problem facing the California legislature in 1989, nor was defining "assault weapons" the only task involving categorization. No doubt lawmakers might have crafted a more nearly ideal statute, perhaps by employing the same expert resources Kasler employed in challenging the law. As the Court of Appeal opinion revealed, there were clearly gaps in legislative knowledge about firearms. But it is also the case, as the California Supreme Court noted, "[t]he perfect can be the enemy of the good."¹¹⁷ The legislature inevitably must decide how to allocate the various limited resources at its disposal, including the waking hours of legislators and their staffs. An hour spent refining the category of "assault weapons" is an hour unavailable for other purposes. Moreover, every hour spent in the effort likely does not produce equal returns in terms of precision. It is more likely that as the effort to achieve ideal precision is approached, more marginal effort is required to make improvements. Moreover, it is not merely legislators whose time is at issue. As other scholars have noted, "rules" may consume more legislative time than "standards," but demand less judicial or administrative time at the point of application. Louis Kaplow has suggested a framework for taking account of the costs and benefits at all stages of legal regulation of varying degrees of rule precision.¹¹⁸

Put in more general terms, utilizing the framework set out above, with a given set of technological and linguistic resources, increasing the

¹¹⁵ Kasler v. Lungren, 61 Cal. App. 4th at 1253, 72 Cal. Rptr. 2d at 270-71.

¹¹⁶ Kasler v. Lockyer, 97 Cal. Rptr. 2d at 347.

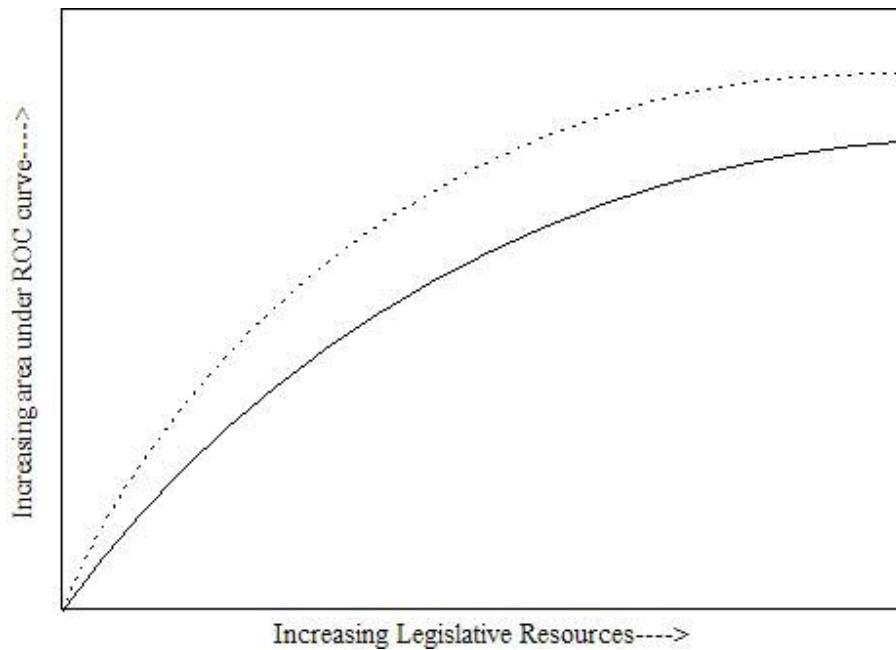
¹¹⁷ *Id.* at 344.

¹¹⁸ Louis Kaplow, *A Model of the Optimal Complexity of Legal Rules*, 11 J.L. ECON. & ORG. 150 (1995).

area under the statutory ROC curve for a single statute requires expending more legislative resources as reflected in Figure 3.

Figure 3

The Costs of Legislative Precision



As is also indicated in Figure 3, the relationship between legislative resources and the precision with which a legislative category can be specified may change if the available level of understanding about the subject of the categorization increases. Consider, for example, the challenges facing legislatures seeking to ban possession of certain plant material before and after the genetic sequence of the plants has been determined. As explained in much greater detail below, the costs of statutory precision are affected both by the means by which a legislature expresses itself, and limited by the level of theorization in the domain law seeks to regulate.

At any particular point in time, however, for a legislature of given collective intelligence and access to a given set of categorization techniques, the constraint on legislative resources translates into a constraint on the total of *all* the areas under the *all* the ROC curves for *all* the statutes it considers in a given time. Put more simply, more time spent

on assault weapons means fewer interpretive errors in the future with regard to assault weapons but more interpretive errors with regard to taxes on agricultural products or other subjects. The decisions with respect to such tradeoffs are generally beyond the view of any judge (or citizen) assessing the legislative work product with regard to a particular legislative category.

D. Openness

The constraint that generated the most controversy in the assault weapons controversies in California was dealing with a changing world in which firearms are constantly changing. Indeed, the reason that the then Attorney General and the California Police Chiefs Association preferred a generic definition rather than a list was the fear that "manufacturers could get around a list merely by renaming a firearm or making simple cosmetic changes in the weapon's design."¹¹⁹ Although especially acute in the face of an intelligent and highly motivated set of actors in the regulated realm, the problem of openness is universal. *Plus ça change*, but not necessarily *plus ça change*. The more exquisitely and precisely a statute is crafted in order to eliminate ambiguity as applied to the current state of affairs, the more quickly it will become irrelevant as the world changes. A United States Constitution written with the precision of the Internal Revenue Code would have quickly become irrelevant, if not entirely meaningless.

A legislative body must weigh the risks entailed in regulating a changing world. Processes of amendment or administrative interpretation can ameliorate but not eliminate the concern with openness. How great those risks are is affected in part by the variability of the domain of regulation. For example, until recently it was relatively easy to identify plants and animals subject to import restrictions, because the domain was variable on a Darwinian time scale. If traditional genus and species designations were not enough, gene science now provides an unimaginably precise way of specifying a species, by reference to a particular genetic sequence. At the same time, however, the same scientific advances now allow the engineering of species in a lab in human rather than evolutionary time.

E. Fitness to the Problem

Finally, and perhaps most centrally, is the question of how nearly a statute attains legislative objectives, what Colin Diver calls

¹¹⁹ *Id.*

"congruence."¹²⁰ The nature of those objectives is, of course, a political matter, external to technical considerations. And objectives are often conflicting. A legislature wanting to limit the availability of weapons that can be used to kill a large number of people in a short time will certainly want to consider restrictions on semiautomatic shotguns, which can fire hundreds of lethal pellets in a matter of seconds. The same legislature, however, operates in a world in which semiautomatic shotguns are used by law abiding citizens to hunt ducks and obliterate clay pigeons. Further, a legislature may reasonably have objectives that extend beyond the purely instrumental regulation of objects to, and including, affective social norms and the culture. As noted previously, there are cultural reasons that weapons of a certain style and appearance are viewed differently by criminals, by some mentally disordered persons, and by ordinary citizens and legislators, entirely apart from technical details of how much harm they can cause to human flesh. Weapons themselves have semantic content as well as operational capacity; they "mean" something different to different people and within different interpretive communities.

V. THE MACHINERY AT HAND: CATEGORIZATION AND HUMAN COGNITION

Unspoken in the discussion thus far is the one overarching, inevitable, universal constraint on legislation: the constraint of human cognition. What matters to legal outcomes is not how categories are specified on paper, but how they are understood in the minds of people -- especially judges but also individual citizens -- who must make decisions based on categories. Obviously, the enterprise of decoding the meaning of categorical specifications is deeply embedded in both general culture¹²¹ and distinct subcultures of groups and professions. But even more obviously, all people have roughly the same basic cognitive apparatus with which to engage the task of making sense of categories. Understanding the possibilities and limits of legislative categorization therefore requires some understanding of categorization at the level of human cognition.

¹²⁰ Diver, *supra* note 87 at 67.

¹²¹ Recent work suggests significant variations in categorization (and other basic psychological processes) in people in different cultures. For an accessible survey (as well as some more controversial speculations), see [RICHARD NISBETT, THE GEOGRAPHY OF THOUGHT: HOW ASIANS AND WESTERNERS THINK DIFFERENTLY...AND WHY](#) (2003).

Legislatures and judges generally proceed on implicit theory of cognition that has remained unchanged since Aristotle. In the past two decades, however, cognitive scientists have made important progress in understanding categorization as a psychological phenomenon. With some significant exceptions, these scientific advances have generally been ignored by law and legal scholarship generally, especially as they pertain to the problem of legislative categories. After presenting a brief summary of the cognitive science of categorization, I will suggest some ways in which legislatures might take account of what we now know about how categories are processed in the human mind. I will also suggest that interpretivist theorists must now take account of the same science.

A. Classical and Folk Theories of Categorization

In most areas of human knowledge, there are both folk and expert theories. "Folk theories" are those theories held by most people in the culture who lack special training. For example, most people without formal training in physics have a "folk theory" of momentum that predicts the course of a stone on the end of a string swung around in a circle and then released. The folk theoretic prediction is that the stone will follow a curved path, as if the stone possesses a kind of dispositional momentum.¹²² In physics, as in many other areas of knowledge, contemporary folk theories bear a remarkable resemblance to Aristotelian theories.¹²³

Untrained people (including legislators and judges) also generally have an Aristotelian theory of categorization.¹²⁴ If we ask the average person to define or describe a "chair," we are likely to get back a list of necessary and sufficient features, the dominant "folk theory" of concepts and categories.¹²⁵ If a dispute arises about the "chairness" of a particular object, we may turn to a dictionary which operationalizes folk theory, definitions typically being lists of defining features. But, as shown most

¹²² In fact, of course, the stone will follow a straight path, the direction of which is determined by the vector sum of the forces operating on the stone at the moment of its release.

¹²³ See, e.g., A. DiSessa, *Unlearning Aristotelian Physics: a Study of Knowledge-Based Learning*, COGNITIVE SCI. 6, 37-75 (1982).

¹²⁴ Aristotle's category theories, along with the rest of his body of work, are available on the World Wide Web at <http://www.mit.edu/classics/Aristotle/categories.1.1.html>.

¹²⁵ "People apparently have a strongly held belief that there are defining attributes for categories, in spite of the failure of psychologists, linguists and philosophers to find any." Gregory L. Murphy and Douglas L. Medin, *The Role of Theories in Conceptual Coherence*, 92 PSYCHOL. REV. 289, 311 (1985).

famously by Wittgenstein,¹²⁶ many concepts people actually use (e.g., "game"), seem to have no defining features.

As philosophers now recognize, it is impossible, at least outside the realm of mathematics, to specify the defining features of a natural concept or category in a way that satisfactorily classifies all cases we might encounter. The philosophical critique is, however, essentially ignored by legislators. Most statutes begin with a set of definitions, a list of features that specify the objects and circumstances to which the statute is intended to apply. Definitions themselves contain words that must be defined. Thus "assault weapons" is defined with reference to "firearms," which must in turn be defined. But for every problem we solve, we create another, because the features and definitions we use are not themselves primitives and must each therefore be defined. "Firearms" are defined with reference to "combustion," which is both somewhat ambiguous and undefined in the statutes. As the list of features and definitions expands, it is by no means clear that uncertainty about the classification of future instances will be reduced. And as a practical matter, the definitional regress must stop, leaving undefined primitives that are themselves subject to varying interpretations.

The length and complexity of the description of the category is not entirely a matter of choice, however. Some categories are easier to describe than others. In mathematics, the "Kolmogorov complexity" of a mathematical object is the length of the shortest computer program required to generate that object.¹²⁷ Objects of apparently infinite complexity, like the Mandelbrot set, may have quite low Kolmogorov complexity and vice versa.¹²⁸ We might think of the "Kolmogorov complexity" of legislative categories as the shortest set of definitions that will serve to reliably classify most situations that actually arise over time, given the same amount of interpretive effort on the "receiving" end of law.

Thus, the Kolmogorov complexity of "official holiday" or "adequate heat" is plainly less than that of "assault weapons." By comparison, the

¹²⁶ LUDWIG WITTGENSTEIN, *PHILOSOPHICAL INVESTIGATIONS*, 1953. More recent and only slightly less definitive work on this topic is SAUL KRIPKE, *NAMING AND NECESSITY* (Harvard University Press, 1972) and Hilary Putnam, *The Meaning of "Meaning"*, in *LANGUAGE, MIND AND KNOWLEDGE: MINNESOTA STUDIES IN THE PHILOSOPHY OF SCIENCE*, vol. 7, (Keith Gunderson ed., 1975).

¹²⁷ For an introduction to Kolmogorov complexity theory as applied to categorization, see Ulrike Hahn & Nick Carter, *Concepts and Similarity*, in *KNOWLEDGE, CONCEPTS AND CATEGORIES* 73, (Koen Lambaste & David Shanks eds., 1997).

¹²⁸ As represented graphically, the Mandelbrot set is the best known of fractal images. Although it appears infinitely complex, and equally complex at all scales of magnification, the set is precisely defined by a short algorithm. The algorithm is described, and the set can be explored interactively, at <http://aleph0.clarku.edu/~djoyce/julia/explorer.html>. [visited March 6, 2004].

Kolmogorov complexity of "pornography" seems truly awesome to contemplate, as we shall see below. Differences in Kolmogorov complexity of categories, I shall argue, depends less on the actual complexity of the realms of life being categorized and more on whether we have adequate theories about the subjects being categorized. As I shall explain, just as algorithms make possible the precise specification of mathematical objects, adequate theories make possible the more precise specification of some kinds of categories.

In most of human affairs, however, and certainly in most of the domains in which law operates, we do not have adequate theories. Why do people kill? How do we assess whether a motorist is driving reasonably? When is a contract unconscionable? Although we may occasionally import theoretical constructs to help answer such questions, in the main our assessments of the relevant categories -- murder/manslaughter, negligence/recklessness, enforceable/unconscionable -- must be made by other means. In most of law as in most of human affairs, we can only utilize those categorization processes that our ancestors evolved long before the first law was spoken or written down. But we now know a great deal more about those processes, as I describe in the next section.

B. Cognitive Scientific Theories of Categorization

Modern cognitive science seeks to replace philosophical speculation about concepts and categories with theories about how people actually engage in categorization and use categories and concepts; theories that generate predictions that can be tested. The past thirty years have produced thousands of articles describing and seeking to explain, not how people *might* develop concepts and categories, but rather how they actually do so.¹²⁹ For example, we know that some categories seem to have a "radial" structure, in which some members are more "central" -- more commonly and more quickly adjudged to belong to the category. People and dictionaries can agree on the meaning of a category like "bird." The Oxford English Dictionary, 2nd ed., offers "[a]ny feathered vertebrate

¹²⁹ This is not an exaggeration. The PsycINFO database now (March 6, 2004) contains 5560 articles with "categorization" as a subject or title term. I do not hesitate to add that I have read but a small fraction of these articles and add the following caveat: One traverses interdisciplinary boundaries only at the risks of both misunderstanding and being misunderstood. As a law professor, I have often read with some pain the efforts of social work professors or psychiatrists to explain legal principles. Part of the pain comes from knowing that my own work in their worlds may seem as ill-informed, or simply wrong. But one does the best one can, trying to make preliminary judgments as between contesting schools of thought that seem as driven by inconsequential difference as the minor sectarian parties of a foreign country.

animal, . . . : distinguished by their warm blood, feathers, and adaptation of the fore limbs as wings, with which most species fly in the air." But people do not assess creatures having all these features as being equivalently "birdlike." If we ask North Americans to list all the birds they can think of, most will list robins or sparrows more often than either eagles or penguins. If we flash pictures of creatures on a screen and ask North Americans to press a key indicating whether the creature is a bird or not, we can determine that it takes subjects less time to classify a robin than an eagle or penguin.¹³⁰ As reflected in human behavior, at least in the context of some kinds of categories, some members of the category are more typical than others, a result that is not accounted for by classical theories.

In the past twenty years, the early work of Kay, Rosch and Rips that demonstrated that people do not have an Aristotelian response to categories has been extended by dozens of further and more refined experiments and more elaborated theories with no single, coherent theory yet achieving dominance. Indeed, the search for a single theory to account for how people categorize has recently been modified by the increasing evidence that categorization is, to some extent at least, a *domain-specific* function. For example, the same theory that accounts for how we categorize birds may not account for how we categorize inanimate objects.¹³¹ Utilizing technologies that reveal the activation patterns in the living human brain (Positron Emission Tomography or PET and functional Magnetic Resonance Imaging for fMRI), in just the past few years scientists have been able to determine that categorization of different kinds appears to be localized in different regions of the brain. These neuroimaging studies have determined, for example, that the categorization of animals and of tools takes place in different regions of the brain.¹³² Similar studies suggest that rule-based categorization and similarity based categorization are associated with quite different regions of the brain.¹³³ Studies like these have cast some doubt on the value of energy spent debating the relative merits of candidates for "the" theory of categorization,

¹³⁰ Barbara C. Malt & Edward E. Smith, *Correlated Properties in Natural Categories*, 23 J. VERBAL LEARNING & BEHAVIOR 250.

¹³¹ For a general survey of this question, see Douglas L. Medin et al., *Are There Kinds of Concepts?*, 51 ANN. REV. PSYCHOL. 121 (2000).

¹³² Michael A. Kraut et al., *Neural Activation during an Explicit Categorization Task: Category- or Feature- Specific Effects*, 13 COG. BRAIN RES. 213 (2002) [summarizing research in this area but also suggesting that differences may be accounted for by activation of the sensorimotor/cognitive systems in the case of tools -- in other words, when we think of tools we also think of using tools. *Id.* at 220]. See also Edward E. Smith et al., *Alternative Strategies of Categorization*, 65 COGNITION 167 (1998).

¹³³ Murray Grossman et al., *The Neural Basis for Categorization in Semantic Memory*, 17 NEUROIMAGE 1549 (2002).

as well as on the possibility of setting out a complete account of those debates in a law review article. For our purposes, however, it may be useful to at least consider what were considered, until recently at least, the three leading contenders for a general theory of categorization: prototype theories, exemplar theories and theory-theories.¹³⁴ My point in briefly describing these theories here is not to argue for the general validity of any of them, but to sketch them in enough detail to suggest their implications for legislative categorization.

1. Prototype Theories

Law is often spawned by prototypes, by the reaction to particular events taken to represent entire classes of possible events. Willie Horton comes to stand for all dangerous parolees. “Willie Hortonizing” comes to stand for the process of attaching a particularly unpleasant prototype to a politician in the public mind. Lucy Williams has written compellingly of how one abusive mother who happened to be on welfare became the “poster child” for draconian welfare “reform” in Massachusetts.¹³⁵ The very notion of “poster child” is itself prototypical, taken from the phenomenon of using particular victims of a disease to represent all victims. When commentators speak of the purpose of The War Powers Act as being intended to prevent “another Vietnam,” nearly everyone understands the reference. That laws are animated by prototypes does not mean, of course, that they are best expressed by reference to prototypes, but the role of prototypicality in the origin of laws bears keeping in mind.

In cognitive science, prototype theories posit that concepts and categories are represented in the mind by reference to an idealized representation of the concept or category. The prototype is a mental

¹³⁴ I do not include Lakoff's "idealized cognitive model" theory primarily because it has not been taken seriously as a theory within cognitive science, as indicated in footnote 31. I do not share with Steven Winter the view that Lakoff's theory is superior to the competition because "it provides more general, overarching abstraction capable of accommodating the many different forms of categorical and conceptual structure identified in the literature" or because "it provides a more supple theoretical tool that better explains the complexities of radial categories." STEPHEN L. WINTER, *A CLEARING IN THE FOREST: LAW, LIFE, AND MIND* 88 (2001) One of the problems with the ICM theory is that it is *too* "supple" and *too much* "a general, overarching abstraction." The same features that make Lakoff's theory appealing in trying to explain categorization to an audience of law professors make it less than appealing as a scientific theory, primarily because it is not sufficiently articulated to generate predictions that can be tested as again, for example, the predictions generated by the more standard scientific theories of human categorization. Indeed, at some point, it becomes difficult to distinguish an ICM from an "idea," except that all ICM's seem to have the feature of "radialness" in the sense of "radial category."

¹³⁵ Lucy A. Williams, *Race, Rat Bites and Unfit Mothers: How Media Discourse Informs Welfare Legislation Debate*, 22 *FORDHAM URB. L.J.* 1159 (1995)

representation similar to Kant's notion of the schema¹³⁶ for, e.g., triangles, which is different from any particular example of a triangle. New cases are evaluated in terms of their similarity to the prototype. Thus, we can ask a sample of people within a particular culture to rate examples of birds or fruit as being "typical" of those categories on a 7-point scale, from least to most typical, as indicated in Table 2. People can do this notwithstanding that it might be extraordinarily difficult for them to articulate what "typicality" means in this instance. At least as forced into this one-dimensional scale, there is a degree of "birdness" or "fruitness" about which most people agree, although there is no perfectly prototypical bird or fruit.

Table 2
Sample Typicality Ratings for "Fruit" And "Bird"¹³⁷

Fruit	Rating		Bird	Rating
Apple	6.25		Robin	6.89
Pear	5.25		Seagull	6.26
Strawberry	5.00		Falcon	5.74
Blueberry	4.56		Mockingbird	5.47
Fig	3.38		Sandpiper	4.47
Pumpkin	2.31		Penguin	2.63

We can also construct a higher dimensional measure of prototypes by asking people to rate the similarity between pairs of members of a category and between members and the category itself, and then mapping the results, as in Figure 4. Thus, in Figure 4, an apple is closer to the category "fruit" than is a blueberry, but the prototype itself does not occupy a space coextensive with any particular example of a fruit. Another way of describing a prototype is that it is "an exemplar with average values on all

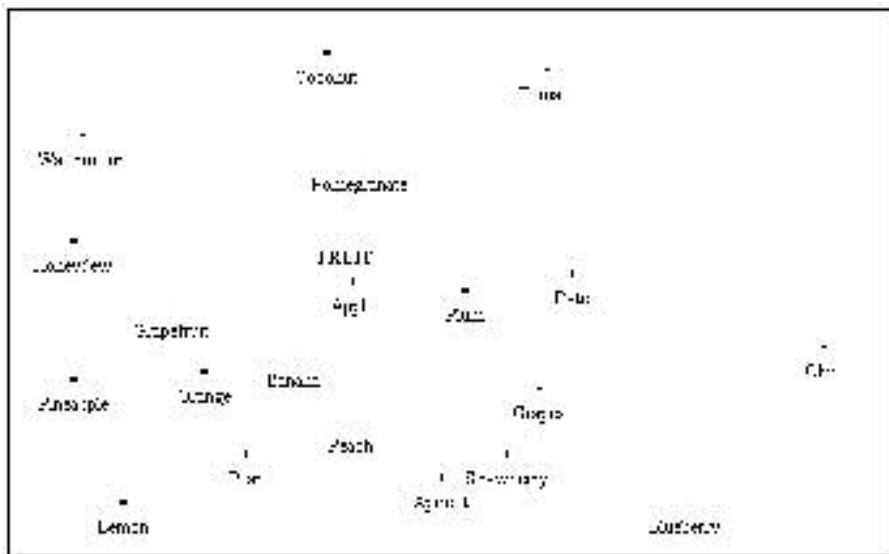
¹³⁶ In THE CRITIQUE OF PURE REASON Kant wrote:

"In truth, it is not images of objects, but schemata, which lie at the foundation of our pure sensuous conceptions. No image could ever be adequate to our conception of triangles in general. For the generalness of the conception it never could attain to, as this includes under itself all triangles, whether right-angled, acute angled, etc., whilst the image would always be limited to a single part of this sphere. The schema of the triangle can exist nowhere else than in thought, and it indicates a rule of the synthesis of the imagination in regard to pure figures in space." [quoted in PHILIP N. JOHNSON-LAIRD, MENTAL MODELS 189-90 (1983).]

¹³⁷ From Malt & Smith, *supra* note 129, reprinted in Edward E. Smith, *Categorization, in THINKING: AN INVITATION TO COGNITIVE SCIENCE* 44 (Daniel N. Osherson & Edward E. Smith eds., 1990).

the dimensions along which the category's exemplars vary".¹³⁸ However described, these examples demonstrate the defining feature, if you will, of prototype theory: what is stored in memory as representing the concept is the abstracted prototype itself, rather than all the exemplars from which it may be derived. In some cases there will exist an exemplar that is the same as the prototype, but this need not be the case.

Figure 4



The Similarity Space for “Fruit”¹³⁹

As is described in more detail below, prototype theory has fallen out of favor among cognitive scientists as a general account of categorization because of its failure to account for several observed aspects of human cognition. For the moment, however, let us examine how prototype theory might inform legislative drafting. We can begin with some of the examples from the experiments described above. Suppose we want to ban the importation of “fruit” from a particular country, in a way that goes beyond listing known species of plants. One way for a legislature to define what it

¹³⁸ Stephen Dopkins & Theresa Gleason, *Comparing Exemplar and Prototype Models of Categorization*, 51 CANADIAN J. EXPERIMENTAL PSYCHOL. 212, 213 (1997).

¹³⁹ From Amos Tversky and J. W. Hutchinson, *Nearest Neighbor Analysis of Psychological Spaces*, 93 PSYCHOL. REV. 3 (1986), reprinted in Smith, *supra* note 136 at 39.

means by “fruit” would be to attach something like Figure 4 to the statute, indicating that “fruit” means any object at least as “fruitlike” as a coconut as measured by some specified distance in the similarity space. Alternatively, the statute might specify the experimental conditions that produce a metric scale like that in Table 2 above, and define a “fruit” as an object with a typicality rating above 3.0. Of course, what is “fruitlike

There is no obvious reason why the same methods could not be used to specify the meaning of a category like “assault weapon.” Experiments and surveys might indicate that the AK-47 assault rifle is closest to the prototype of “assault weapon,” just as the apple is closest to the “fruit” prototype. The California legislature might have defined the “prototypical assault weapon” as the AK-47 assault rifle, and then given the same list of banned weapons, but with an additional feature: an attached set of numbers specifying typicality ratings for each example. These numbers might be determined by surveys or experiments with any set of people: the members of the legislature itself, or a sample chosen by other means. Prototype theory suggests that there need not in fact be any instance of the prototype. The AK-47 may be closest to the prototype of all existing assault weapons, but it need not be coextensive with the prototype.¹⁴⁰

Pornography would appear to pose a much more difficult problem. To begin, we cannot produce a relatively complete set of candidate members of the category -- as with fruits or assault weapons -- because there are an infinite number of possible candidate images. Whether we could get substantial inter-rater reliability in judgments of “pornographicity” -- as among legislators or others -- is an empirical question. I would guess not, even if we used a carefully selected, culturally and sexually homogenous group. But again, this is for now a matter of cultural judgment rather than empirical assessment.

The notion of prototypicality, particularly as expressed in the spatial metaphor of “radial categories,” also turns out to be more complicated than first appears. “Typicality” may mean different things in different contexts. First, we can use “typicality” to describe the extent to which objects are “good examples” of the category. To use an example from Osherson and

¹⁴⁰ For the sake of simplicity and space, I gloss over here the rather important notion that the “typicality” metric cannot be easily specified in a single number. Operationally, we might glean such numbers from a survey, but the answers to the survey might conceal the fact that a typicality metric is an implicit sum of metrics as to many different features: one person may give a Blodgett 47 rifle a .8 typicality rating because of the caliber and barrel length, while another person gives the same rating because of the color of the stock and the name of the gun. Cognitive scientists have also found significant differences in how people learn categories with “linearly separable” vs. “non-linearly separable” (significantly intercorrelated) features.

Smith,¹⁴¹ we may agree that both the biggest selling chair from the IKEA chain and Queen Elizabeth's throne are both chairs, but certainly the IKEA chair is more "typical" a chair in this sense. A second, but distinct aspect of typicality relates to the perceived degree to which a candidate example "belongs" to the category -- the "gradient of membership." Both the Queen's throne and the IKEA chair are, without question, chairs. But what of a three-legged stool, a bus bench, or a hassock? We can experimentally determine that people in the same culture have roughly the same take, expressed roughly thus: the IKEA chair may be a 100% (1.0) chair, but the three-legged stool only a 0.7 chair. The difference we can describe as the category membership function. Surveys and reaction time experiments may produce single numbers for "chairness" of candidate objects, but already we begin to see that the single metric may conceal differences in typicality and category membership.

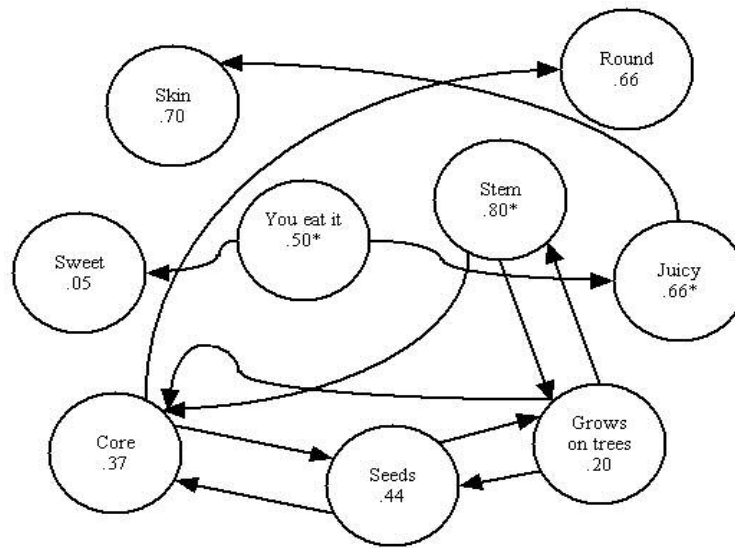
Recent work by cognitive scientists on conceptual coherence suggests there is often an even more fine-grained internal structure to concepts and categories.¹⁴² One method of assessing the internal structure of a concept rests on measuring "mutability," the degree to which people judge that the concept or category remains the same even as the feature in question changes. Mutability can be measured by four tests: surprise, ease-of-imagining, goodness of example, and similarity-to-an-ideal. "Surprise" is measured by asking subjects how surprised they would be to encounter a transformed instance: e.g., an apple that did not grow on trees, an assault weapon made of plastic, or pornography that does not depict the human genital region. "Ease-of-imagining" is measured by asking people to judge how easily they can imagine an instance of the category without the feature. "Goodness of example" is measured by asking how good an example of the category they judge an object without the feature to be: e.g., "How good an example of an apple is an apple that does not grow on trees?". The "similarity to the ideal" is obtained by asking subjects, e.g., "how similar is an apple that does not grow on trees to an ideal apple?" By these means, the dependency relations between features of the concept or category can be described as in Figure 5 below. These experiments suggest that the "central" features of a concept, those that are the least mutable, are those upon which most other features in turn depend. The same experiments suggest the existence of, and means of ascertaining, the parts and aspects of concepts that have their own dependency structures, as in the two

¹⁴¹ Daniel Osherson & Edward E. Smith, *On Typicality and Vagueness*, 64 COGNITION 189, 189 (1997).

¹⁴² Stephen A. Sloman et al., *Feature Centrality and Conceptual Coherence*, 22 COGNITIVE SCI. 189 (1998).

subnetworks of Figure III below that pertain to the reproductive and the food-related features of apples.

Figure 5143
The Conceptual Structure of "Apple"



How might these theories and technologies of assessment be used in drafting legislation? Consider first whether information as represented in Figure 5 about "apples" would be helpful to a judge or administrative

143 *Id.* at 205. Arrows point from a feature to one that it depends on. "Ease of imagining" judgments obtained in experiments are shown beside each category-feature.

agency charged with determining whether a new, genetically engineered fruit is an "apple" for purposes of a statute regulating agriculture. The answer is almost certainly affirmative: There is certainly information in Figure 5, or other representations of the same kinds of assessments of meaning, that is not present in the category label standing alone. We know from the words of assault weapons legislation that features like handgrips and flash suppressors matter, but we do not know how much they matter (or, as we shall see, why). Revealing the internal structure of feature interdependence conveys information on this point (even before we understand why some features matter more than others).

Thus, at least in principle, a legislative analyst charged with reducing legislative action to statutory language might use these methods to get at what legislators mean by "assault weapons," producing a chart like Figure 5 as part of the legislative history of the statute. A judge asked to decide whether a new weapon should be added to the list of "assault weapons" might make use of such information as significant, perhaps even the best, evidence of legislative intentions.

2. Exemplar Theories

Prototype theory has strong intuitive appeal. Unfortunately, as a general theory of how people learn and use categories and concepts, prototype theory now appears to be inadequate in many respects.¹⁴⁴ First, prototype theories treat concepts as context-independent. But experiments demonstrate the opposite. For example, asked to rate the typicality of various "beverages at breaktime," people give a different ordering depending on whether the people taking the break are secretaries or truck drivers.¹⁴⁵

Further, prototype theory suggests certain results when concepts are combined. For example, people think metal spoons are more "spoon-like" than wooden spoons, and smaller spoons "spoonier" than larger spoons.

¹⁴⁴ W.K. Estes observes: "It is interesting to note that prototype theory is by far the most visible variety in the literature, although it can be credited with none of the close quantitative accounts of categorization data that have appeared during the last decade, the majority of which have been achieved by exemplar-similarity models. The popularity of prototype theory appears to derive from a combination of factors -- its intuitive appeal, its long history, and some results of experiments employing categories of objects ((typically irregular polygons or dot patterns)) produced by means of variations on experimenter-defined prototypes. W.K. Estes, *Models of Categorization*, in 29 CATEGORIZATION BY HUMANS AND MACHINES 15 (Glenn V. Nakamura et al. eds., 1993) [citations omitted].

¹⁴⁵ Emilie M. Roth & Edward J. Shoben, *The Effect of Context on the Structure of Categories*, 15 COGNITIVE PSYCHOL. 346 (1983).

But people judge large wood spoons more typical than small wooden spoons. In sum, people seem not to do what prototype theory suggests that they do: retain only the abstracted essence, the central tendency, of the concept, discarding irrelevant detail. It appears that people also take account of the interactions among the features of the objects they categorize, especially when those features are not, in the observed world, independent (or, in the nearly unpronounceable term of cognitive science, "non-linearly separable"). In other words, we take notice of the fact that wooden spoons tend to be large and metal spoons small. The attraction of prototype theory, its parsimoniousness, is surely lost if it predicts that people retain prototypes for every combination of features (small metal spoons, large wooden spoons, and so on). Exemplar theory challenges the very notion of prototypes, while seeking to account for prototype effects

Exemplar theory is something like prototype theory, only less so. That is to say, whereas prototype theory assumes that prototypes are extracted and stored separately from the particular category members or exemplars with which a subject is familiar, exemplar theory posits that subjects directly access traces of multiple exemplars in memory at the (later) point at which the categorization of a new instance is required. Categories thus have no independent representation, but are an emergent property of the collection of exemplars associated with the category. Most of the recent research suggests that exemplar theory accounts for a wider range of phenomena than prototype theory.¹⁴⁶ As one recent account puts it, "the literature has come to favor instead a generalized exemplar principle in categorization."¹⁴⁷

Consider again the California legislature's initial effort at defining "assault weapon" by reference to a list of particular weapons. Exemplar theories of categorization operate in the same way, the category being implicitly defined as "these things and others 'like' them": the category is in the examples themselves. Prototypes, if they exist, are constructed "on the fly" and as needed, rather than stored independently in memory. Exemplar theories of categorization are thus entirely congenial to the ancient

¹⁴⁶ For overviews of the literature as well as particularized experimental studies, see Koen Lamberts, *Exemplar Models and Prototype Effects in Similarity-Based Categorization*, 22 J. EXPERIMENTAL PSYCHOL.: LEARNING, MEMORY AND COGNITION 1503 (1996); Dopkins & Gleason, *supra* note 137; and GREGORY L. MURPHY, *THE BIG BOOK OF CONCEPTS* 73-114 (2002), noting the relatively greater strength of the predictive power of exemplar theories in experiments, but also noting that prototype theories may be more appropriate to natural categories, and that there may be differences in categorization among people. *Id.* at 114.

¹⁴⁷ J. David Smith & John Paul Minda, *Prototypes in the Mist: The Early Epochs of Category Learning*, 24 J. EXPERIMENTAL PSYCHOL.: LEARNING, MEMORY & COGNITION 1411 (1998).

principle of *ejusdem generis* in statutory interpretation. This canon of construction (translated as "of the same kind") provides that when the general description follows a list of more specific examples, the general words are limited to those similar to those enumerated specifically. For example, in the phrase "automobiles, trucks, tractors, motorcycles and other motor-powered vehicles," "vehicles" would not include airplanes, since the list was of land-based transportation.¹⁴⁸

How might the California legislature have used exemplar theory? Rather than assuming that there is a single prototype for a legislative category and trying to probe the internal structure of that concept, the legislature might attend more directly to real exemplars. Providing a rather extensive list of exemplars seems relatively feasible in the case of "assault weapons." There are books that catalog virtually every kind of weapon being mass produced.¹⁴⁹ A legislative committee could go through the catalog and mark those exemplars it would include and those it would not. The annotated catalog could be sent forward with the legislation *as* the category. Moreover, "assault weapons" are not determined only by reference to what they are, but in contrast to what they are *not*: "sporting firearms." The same firearms catalog could be marked to indicate exemplars of sporting firearms, and clear non-membership in the assault weapons category. Indeed, listing the firearms explicitly *not* being regulated is a method that has been used in federal firearms legislation, albeit possibly more out of concerns for political liability than semantic clarity.

How might the exemplar approach be used with legislation concerning "pornographic images"? Here there is no finite list, no catalog, that can be categorized by exemplar/non-exemplar status. A legislature might employ a version of Justice Stewart's method to categorize, say, a sample of 1,000 images taken off the Internet. Future citizens and judges would be left with the prospect of viewing this collection and comparing candidate images with the sorted collection. Not a pleasant prospect, certainly, but it does seem that both citizens and judges would have a better idea of what the modern Justice Stewart sees when he sees and knows pornography, after he has seen it and expressed the state of his knowledge.

In effect, exemplar theory works in the same way that our common law is supposed to work: categories and concepts evolve out of the disposition of particular instances: in situations A, B and C, we have

¹⁴⁸ From <http://dictionary.law.com>.

¹⁴⁹ The Jane's Company produces exhaustive paper and on-line catalogues of virtually every sort of weapon, including assault weapons. See e.g., JANE'S INFANTRY WEAPONS and other publications at <http://catalogue.janes.com/> [visited March 6, 2004].

instances of "gross negligence," whereas in situations X, Y and Z we do not. Although a court may say what it thinks are the principles underlying its decisions, what is supposed to matter is what the court actually does in a particular case. Subsequent cases are to be decided based upon their "similarity" to the decided exemplars in the reported cases.

Of course, critical analysis quickly reveals that "similarity" is not an unproblematic concept. As Nelson Goodman most famously demonstrated, one cannot say with precision whether even two things are similar without specifying the "respects" in which they are similar.¹⁵⁰ Like any two objects, any two cases are "similar" in a potentially infinite number of respects: whether the claim arose on a Tuesday, is prosecuted by a Christian, involves animals, and so on. A legislature cannot merely provide a list of exemplars and be certain that future judges will extract the correct dimensions of similarity from the infinitely many possible dimensions. For example, if all the examples of assault weapons have barrels longer than 6 inches and wooden stocks, future judges may assess these as important (indeed, universal) features of the legislated category, even though legislators may well have banned sawed-off assault weapons with plastic stocks as well, if the possibility had occurred to them. This logic counsels the legislature drafting exemplar-based laws to include as many exemplars as possible, including exemplars that vary as much as possible on the non-significant dimensions, and to include exemplars that are outside the intended category. Deciding which exemplars to use, which features to mention, seems more like an art, however, than the science this article promises. Perhaps we can move a bit more toward the scientific, or at least scientific, by looking at how exemplar-based categories could be represented in connectionist representations of categories or concepts.

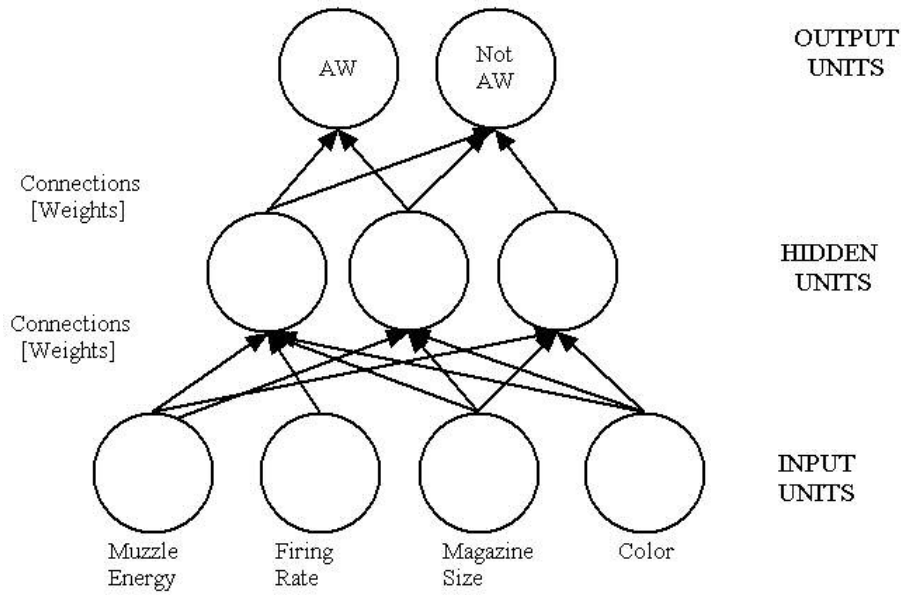
3. Prototypes, Exemplars and Neural Networks

I begin with a brief account of a neural network, the basic architecture of what are called "connectionist" models of cognition. In neural networks, information is representable not in the form of symbols, rules or statements, but in the pattern and strength of the connections between simple processing units modeled on biological neurons. As in the biological brain, each simulated neuron is a processing unit with an input and output side. The strength of the signal on the output side is a function of all of the signals on the input side. The outputs of some neurons are connected to

¹⁵⁰ Nelson Goodman, *Seven Strictures on Similarity*, in PROBLEMS AND PROJECTS (Nelson Goodman ed., 1972).

the inputs of other neurons. The simplest functioning neural network consists of three layers of neurons: an input layer, an output layer, and a hidden layer, organized as indicated in Figure 6.

Figure 6
A Simple Neural Network



Neural networks can be "trained" to categorize as follows: Each input neuron is set according to one feature of the object. Each output neuron is associated with one or another category. The network is then "trained" on a large number of examples -- the "training set." As each example is presented, input neurons are set to match the features of the exemplar. The connection weights are initially set at random and the network initially produces a "guess" at the correct categorization of the exemplar. A "training" supervisor then indicates whether that categorization is correct. With each iteration, the connection weights are adjusted to lessen the difference between the "guess" and the correct output, as supplied by the supervisor. By this seemingly simple process, neural networks possess quite extraordinary capabilities in recognizing complex, subtle patterns, including information in schematic form, and to learn very subtle differences that may distinguish members of one category from members of another. The "knowledge" that permits the neural network to so perform is not contained in any list of descriptions, statements, or propositions. It is entirely contained in the vector of connection weights

between neurons in the network. The source of this "knowledge" is, of course, the supervisor or trainer of the network. The value of the neural network approach is that the trainer need not specify a rule explaining why a particular object belongs in one category or another. The network will "learn" whatever implicit rule the trainer is using, whether the trainer can state it explicitly or not. The trainer need only, with Justice Potter, "know it when he sees it."

Is the neural network model a prototype theory or an exemplar theory? The answer is both. The network will display prototype effects, more confidently classifying objects that approximate a prototype, but the ability of the network to do so is based entirely on training over a set of exemplars. The neural network behaves *as if* it is calculating the difference between a sample object and a stored prototype, and *as if* it was doing so on the basis of information extracted from exemplars.

Applying connectionist implementation of exemplar theory to the problem of legislative categorization might yield something like the following: First, a network is constructed to have only two output neurons, consistent with (a) "assault weapons" and (b) all other objects, as in Figure 6. Second, the "exemplars" are found among the examples of the training set. In order to be useful, the training set must include both members and non-members of the category: single shot pistols as well as AK-47's; guns with pink stocks as well as brown.¹⁵¹ Third, the architect of the network must specify which features of the exemplars are to be associated with particular inputs to the input neurons: barrel length, caliber, bullet weight, speed of reloading and discharge, and so on. Fourth, the network must be "trained" by a supervisor, someone able to provide feedback to the network when it assesses that a particular object is, or is not, within the category. The supervisor need not be aware of any particular rules for determining category membership. Rather, like Justice Stewart, the trainer must know an assault weapon when he sees one. At the end of this process, the network will have an ability to categorize new candidate objects, an ability that approaches asymptotically that of the trainer.

In our thought experiment, both the architect and the training supervisor of the network would be the legislature itself. Having determined which features should be attended to, the legislature would then

¹⁵¹ The training set is crucial. A possibly apocryphal story has an early military target recognition network trained to distinguish Soviet-bloc tanks from those used by NATO forces. The training set consisted of hundreds of photographs. Because of the difference in the difficulty of acquiring pictures of the two categories of tanks, most of the Soviet-bloc tank pictures were of tanks under trees and in the shadows. The neural network learned to associate shadows with things Soviet, such that a US tank under a tree was classified by the network as potentially hostile.

consider a substantial list of exemplars as a training set. Feedback would be presented to the network as to whether each of the exemplars was, or was not, a member of the category. At the end of this training, the specification of the category would be in the network itself, in the connection weights produced by the training. This network structure and vector of connection weights would BE the category definition. Just as the voice recognition and OCR programs on my computer can recognize both the sight and sound of the letter "A" without having any set of rules for distinguishing "A" from "H," a statute banning "assault weapons" might come with an appended neural network, trained by the legislators at the time the legislation was passed. In some senses, this is probably as close as one can get to excavating with precision what legislators actually meant in using a term like "assault weapons."

Such an enterprise might be helpful to the Attorney General and judges called upon to classify new weapons. But, because the "meaning" of the category is embedded in the weights of the network, to ordinary citizens lacking access to the network, the results would be entirely opaque and unconstitutionally vague, at least until a judge has ruled on a particular candidate firearm. Even if this concern were met (by making the network available on the internet, for example), it is unrealistic to suppose that a legislature would engage in the tedious process of training a neural network on a set of possible exemplars of each of the categories the legislature intends to define. In the face of all that modern science can offer, Aristotle begins to seem a more reasonable choice.

But neural network computer programs are not the only means by which a legislature acting on the basis of exemplar theories of categorization might legislate. The legislature might simply include a list of exemplars and near-miss excluded cases. A legislature can then hope that judges will detect the same similarities in the exemplars that they detect, aided by a large and well-crafted list of exemplars. Or legislators can specify the features or "respects" in which the exemplars are similar and reflect the category being specified. The legislature might additionally indicate which of the exemplars are particularly good examples of the intended category. The "goodness of example" measure might be made based on votes as to particular exemplars: exemplars at the categorical margin would be expected to draw closer votes.

4. CATegorization in Context


Thus far, we have taken the categorization task as one that can be performed without much reference to context. At least in the context of

objects, it would seem that we do not need context: we are likely to regard an object as being an assault weapon without regard to the day of the week, the gender of the owner, or any other fact extern to the object itself. In many cases, however, we cannot categorize without context.

To continue the neural network analog, consider a simple neural network designed to "read" marks on paper: an optical character recognition program charged only with classifying handwritten marks as either the letter "A" or the letter "H" -- a task performed readily performed by such programs on Palm Pilots and other "Personal Digital Assistants." The "features" of the marks might be whether the mark touches squares (pixels) on a predefined grid, as in Figure 7, or by characteristics defined by other means, like the angle between the two "legs" of the character.

Figure 7

A Pixel Grid for Letter Recognition



Each character presented to an input device like that reflected in Figure 7 will generate a vector of numbers. The neural network will process those input numbers and generate a "guess" as to whether the character is an A or an H. The trainer of the network will provide feedback to the network indicating whether the guess is correct or incorrect. As the

network is trained on hundreds of examples of different characters, in different fonts, shapes and sizes, its accuracy will improve. Which is not to say that it will ever become perfect. Nor can it learn to do things its trainer cannot do, such as resolving the ambiguity of the middle figure in Figure 7.

As in the case of "river banks" and "banks with money in them," ambiguity can sometimes be resolved by introducing additional constraints from the context of usage. Consider the same ambiguous middle character in Figure 7, in the context of the usage in Figure 8.

Figure 8
Disambiguation through Context



We effortlessly assign the very same object to different categories, depending entirely on the context.

How might the constraints supplied by context help disambiguate an object potentially classifiable as an "assault weapon"? Consider the guns used in some carnival games, in which customers fire pellets at mechanical targets. The guns are typically affixed to the game booth in such a way that they cannot be aimed at anything other than the target area. Many of them are automatic weapons, in the sense that they fire multiple shots with one pull of the trigger. Some of them might fit within some definitions of "assault weapons," particularly if they were detached from the game machinery and able to fire at people in the carnival crowd. Are they therefore "assault weapons"? The context (attached to the game machinery or not) might supply the answer, even if the legislature had never thought to consider whether to include them.

5. Theory-Theories

Some categories seem to depend almost entirely on context. Both prototype theories and exemplar theories have a certain surface plausibility in the context of categories like "fruit," "bird," or "assault weapon." But

consider the following members of an as-yet unidentified category: cats, insurance policies, birds, deeds, jewelry, photographs, dogs, bonds, children, original paintings and passports. No obvious prototype can be extracted from such a list. The list of exemplars provides no clear way to determine whether, for example, "cash" or "cheap reproductions of paintings" belong to the list. However, things become clearer when we are told that these are "things to take out of one's home in the event of a fire". Suppose in such a moment of crisis, as one quickly collects possessions as the smoke thickens, a roommate calls out, "Take the cat." Does she mean the tabby recently liberated from the pound, or the Magritte painting of a tabby inherited from the Dadaist grandfather? The answer might depend on whether the painting is a reproduction, as well as on the presumed values of the roommate.

Or, consider the working title of this piece. What do "lawyers, guns and money" have in common, other than as "things to send to an adventurer in distress"?¹⁵² The category is incoherent absent an explanatory theory in which the items are similar-in-context (things that are highly valued, difficult to replace and more easily moved). Or, to consider a perhaps more legally realistic though remotely analogous example, consider the following list of animals: a gray bat, a black-footed ferret, a giant panda, and a jaguar. We cannot examine a collection of specimens of these animals and induce a particular category to which they belong, in part because they belong to infinitely many: (mammals, not-very-colorful animals, animals that attract fleas, etc.). What gives the category coherence in this case is external to the exemplars themselves: these animals are listed by the Environmental Protection Agency as endangered species.¹⁵³ They are reasonably linked only because these are species that have not fared well in particular ecological settings. These are also animals one would choose to save in the event of a local ecological catastrophe because, like our papers and mementos, they are hard to replace. But note the difficulty we would have in specifying the features of either animals or precious belongings. Describing a prototype for these categories in any systematic way seems equally problematic.¹⁵⁴

The reason is that such categories and concepts make sense only in the context of a purpose, theory, explanation or narrative. Some theorists would extend this notion to all categorizations and concepts, arguing that

¹⁵² Zevon, *supra* note 40.

¹⁵³ Endangered and Threatened Wildlife, 50 CFR §17.11.

¹⁵⁴ There are, however, good reasons for advocates to try. Conservationists will put forward whales or other likeable creatures as "poster species" for preservation; opponents will highlight kangaroo rats or undistinguished little fish.

all categorizations and concepts are embedded in implicit or explicit theories about the world. We know, for example, that the centrality of a feature in determining category membership is determined in large degree by its position within the causal structure of the category.¹⁵⁵ Thus, "roundness" is perceived as a more central feature of basketballs than of cantaloupes. One explanation for this phenomenon is the differential way in which roundness figures into the overall causal relations among features. In the case of basketballs, roundness permits predictable bounces and falling through round hoops. A square cantaloupe might nonetheless be sweet, juicy, gray-green and contain seeds.¹⁵⁶

Whether all categories and concepts are theory-dependent depends on one's theory (or prototype) of theories in general. The notion of "adequate heat," so easily operationalized to numbered marks on the linear scale of a thermometer, does not at first seem to be theory-bound. But, on closer inspection, why do we choose 72 degrees in a statute about "adequate heat" for residential housing? People do not experience as equivalent the differences between 52, 62, 72, and 82 degrees. The linearity of the temperature scale relates to the physics of heat expansion of materials (e.g., mercury), not about anything of direct importance to people. One also suspects that another number might have been chosen at other period of history, when clothing conventions were different. What assumptions underlay the use of the term "occupied room" or the mandate to measure temperature 36 inches above the floor? Here the underlying, invisible assumptions are cultural rather than scientific. As this example demonstrates, even the simplest concept or category rests on a range of unstated assumptions and theories about people and the physical world. We can disregard the underlying theories with some comfort only because the situational assumptions (e.g., clothing conventions, physiological mechanisms) are relatively constant. And we can specify the category with precision because we have an excellent set of theories that enable us to measure heat with great precision.

The category of "assault weapons" provides an interesting contrast. If we are given exemplars of the weapons listed in the California statute, we will notice several common features: All have barrels. All fire projectiles that are propelled by small explosive charges contained in metal cartridges. All have a mechanism that reloads another cartridge automatically when a round is fired. All have a trigger. Most, but not all, have a shoulder stock. Certainly, the similarities are more readily apparent than was the

¹⁵⁵ Woo-kyoung Ahn et al., *Causal Status as a Determinant of Feature Centrality*, 41 COGNITIVE PSYCHOL. 361 (2000).

¹⁵⁶ *Id.* at 365.

case with the list of animals on the endangered list. But is there a “theory” of assault weapons?

Possibly. Most of these weapons have a shoulder stock because of the operation of Newton’s Third Law of Motion¹⁵⁷ and the fact that these weapons fire a moderately heavy bullet at high velocity which results -- by the laws of physics -- in a substantial recoil. Such a firearm that did not transfer that recoil energy to the more massive parts of the human body (i.e., a pistol), might well break the wrist of the shooter. The relative heavy bullet and high velocity are, in turn, functions of the design goal of these weapons: to inflict significant damage on human bodies. The damage a projectile causes on impact is a function of its mass and velocity. In addition, these weapons fire many rounds in rapid succession because of the design goal that they inflict heavy damage in a short period of time, in order to reduce the time the soldier using them is exposed to counterfire. In short, the very reason these weapons are *called* “assault weapons” is that they evolved for optimal use in a range of infantry assault situations, with the overall design goal of enabling the rapid infliction of lethal injury to as many human beings as possible in the shortest possible time. Viewed in this light, the deeper commonality in the exemplars in the California statute extends beyond the more obviously similar features, and might enable the specification of “assault weapons” that share virtually none of the features of contemporaneous exemplars.

A “theory-based” statute banning assault weapons might ban “all weapons capable of firing more than a X projectiles with energies greater than Y within a period of Z seconds.” Such a statute begins to look more like the “heat” statute, and to offer some of the same benefits. First, merely changing the appearance of a weapon no longer removes it from the ban. Second, legislative purposes in reducing access to especially lethal weapons seem better accomplished. Finally, such a statute is rather more easily communicated, understood and applied by citizens, administrative agencies or judges because it contains an operationalized, theory-based test for determining category membership. As with the heat statute, we have reduced the meaning of the category to more basic (arguably the *most* basic) terms: mass, velocity, time.

Note, however, that the “theory” relied upon in this “theory-based” statute is a theory in physics. The political reality (for which we have less satisfying theories) is considerably more complex. One of the challenges of drafting a politically viable assault weapons ban is that there are weapons of equally terrifying lethality that legislators do not wish to ban, such as the

157 Newton's Third Law of Motion: For every action, an equal and opposite reaction.

semiautomatic shotguns favored by duck hunters. There is no constitutional requirement, however, that legislators adhere to a single categorizing device. The legislature might adopt a "theory-based" category, but then explicitly exclude from it what would otherwise be a subcategory: semiautomatic shotguns.

Turning to the most difficult case, "pornography," how might a theory-based categorization of "pornography" be developed? We might begin by looking, as a practical matter, at some of the many reasons legislatures have for seeking to limit pornography. Some of these are matters of esthetic preference about what images ought to be displayed in a society with aspirations of certain kinds of "decency." Others incorporate explicit or inchoate theories of human behavior: that men exposed to certain kinds of images will be more inclined to sexual aggression; that children will be induced to inappropriately early sexual activity. Could a legislature define "pornography highly likely to induce sexual aggression/sexual activity" or "PHLISA"? This is, at bottom, an empirical question. Legislatures can define "adequate heat" in terms of marks on a thermometer because empirically, all thermometers behave in quite similar fashion in the presence of the same temperature. Legislatures can define "assault weapons" in terms of lethality by implicit reference to the underlying physics of how projectiles inflict injury on the human body, safe in the assumption that bullets have roughly the same effect on all human bodies. But one can scarcely say the same of the effect of PHLISA on "men." One could try to operationalize PHLISA in a more refined way, of course. "Men" could be operationalized as "a random sample of men between 18-24" or "a sample of recently paroled, historically heterosexual sex offenders." "Sexual activity" could be operationalized by reference to pupil dilation, galvanic skin response, or changes in blood chemicals, or by reference to performance on some psychological test instrument after exposure to sample stimulus images.

But notice the escalating Kolmogorov complexity of the category specification. The messier and more uncertain the world and the less access we have to reasonable theories, the longer and more detailed must be any purported specification of category. Even then, a law proscribing PHLISA in terms of images that produce a certain blood chemical change in a random sample of adult males will not be adequate to ban images that have effects on those deviant individuals most likely to engage in sexually aggressive acts. Nor would it be possible to define PHLISA by reference to a random sample of deviants, because deviance itself has too many dimensions: some may react in bizarre ways to pictures of Minnie Mouse, to say nothing of Snow White. In short, there is no adequate theory-based

categorization of "pornography" because there is no adequate theory of the human behaviors as to which pornography is relevant. We can and do, of course, utilize arbitrary and unsupported theories. We can also postulate that the specific gravity of witches is greater than 1.0. But many innocent people will drown as the result of our ill-theorized categories.¹⁵⁸

6. Categorizing Scripts and the Function of Narrative.

Thus far our examples entailed the categorization of objects. But what of categories of the kind (!) that occupy much of the categorizing work of law: first degree murder, parole agreements, hearsay evidence, and so on? As I explain in this section, these categories are perhaps best understood in the context of an extension of the "theory theory" of categorization, and a persistent attention to context.

To begin with an easier case, consider the category "waiter." At some food establishments (restaurants in the Sizzler chain, for example), customers pick up food orders from a window, but then a restaurant employee brings drinks and, if ordered, desert. Are these people "waiters?" How do we decide? We may have ideas about the features of a prototypical waiter (by dress, gender, behavior, and so on), but these features don't seem especially helpful here. And clearly each of us has access to numerous exemplars of "waiter." Our concept of "waiter" only makes sense in the term of a stored "script" for "eating at a restaurant." A "script" is a stereotyped sequence of actions that defines a well-known situation.¹⁵⁹ We can think of scripts as schemas for events: the knowledge structures that tell us the sequence of events to expect when we enter a restaurant, as well as what to expect if we order a bowl of fruit. Over time we develop different expectations or scripts. We learn, for example, that in restaurants serving cuisines originating in Asia we are more likely to find chopsticks than silverware. We learn that in restaurants in Kyoto, as in small cafes in rural Oklahoma, one generally pays at the cash register. The Sizzler employee fits the category "waiter" to the degree that he or she matches that role in the prototypical script for restaurant meals. The fit is not perfect (or prototypical). As in the case of objects, context may help: The

¹⁵⁸ There is, in addition, a potential problem of recursive ambiguity: If the categories depend on theories, how do we describe those theories in ways that do not merely move the locus of ambiguity to the expression of those theories? A full answer to this question is beyond the scope of this paper, but this is another area in which the tools of cognitive science, and the notion of parallel constraint satisfaction in connectionist networks, have been usefully brought to bear. See PAUL THAGARD, COMPUTATIONAL PHILOSOPHY OF SCIENCE (1988).

¹⁵⁹ ROGER C. SCHANK & ROBERT P. ABELSON, SCRIPTS, PLANS, GOALS AND UNDERSTANDING: AN INQUIRY INTO HUMAN KNOWLEDGE STRUCTURES (1977).

ambiguously categorized employee might be considered a waiter for purposes of tipping, depending on whether our implicit theory of tipping concerns augmenting the income of low wage workers in such settings, or paying for some kind of *quid pro quo* in service.

The existence in memory of scripts has significant effects for how we interpret and remember events. When given a story about a situation for which we have a script, we are likely to falsely remember having read about events not in the text, but supplied by the script itself.¹⁶⁰ We tend to reorder the sequence of events to fit the sequence supplied by the script for such events.¹⁶¹ Doctors seem to use script-like knowledge to arrive at diagnosis on the basis of the pattern of patient symptoms.¹⁶²

The relationship between scripts and categorization is bidirectional. Consider the classic category of "things to take out of a burning home." The "theory theory" explanation of this category can be understood in terms of specifying those objects that fill one of the slots in a script. In this instance, the script enables and explains the categorization of objects. But we also categorize scripts themselves. The example just given is an exemplar in a category of scripts we might call "situations calling for urgent action to save things of value." Much of the law is chiefly concerned with categories of this kind.

Consider "lying in wait," "provocation," "undue influence" or "insider trading." Imagine trying to convey the import of these doctrines *without* resort to stories that typify and -- together with the scores of similar (!) stories in literature or appellate decisions -- constitute these categories. Some category labels, like "good Samaritan," retain traces of the original story that gave rise to them.

Although it is not clear whether the processes entailed in such categorization are the same as those in the categorization of objects, in part because scripts have dimensions -- temporal sequence and causal relations -- lacking in objects,¹⁶³ it seems likely that the categorization of scripts relies on prototypes or accessing exemplars at the time of categorization. In perhaps the leading theory of script processing, Schank and Abelson

¹⁶⁰ Gordon H. Bower et al., *Scripts in Text Comprehension and Memory*, 11 COGNITIVE PSYCHOL. 177 (1979).

¹⁶¹ Valerie Abbott et al., *The Representation of Scripts in Memory*, 24 J. MEMORY & LANGUAGE 179-199 (1985).

¹⁶² Bernard Charlin et al., *Scripts and Medical Diagnostic Knowledge: Theory and Applications for Clinical Reasoning and Research*, 75 ACAD. MED. 182 (2000).

¹⁶³ Neuroimaging studies suggest that two different kinds of processing in two different regions of the brain are involved in processing scripts: one devoted to the temporal sequence of script events and another devoted to categorizing entire scripts. Arnaud Partiot, *Brain Activation During Script Event Processing*, 7 NEUROREPORT 761, 765 (1996).

propose that people retain in memory "story skeletons" that serve to standardize particular situations,¹⁶⁴ in effect functioning as prototypical scripts.

In general, the categorization of scripts, stories or narratives has been the subject of far less inquiry by experimental psychologists than the categorization of objects,¹⁶⁵ but has drawn considerable attention from researchers on juror decision-making. The leading theory of how jurors make decisions is the "story model" of Pennington and Hastie. In their model, jurors arrive at decisions by (a) organizing evidence into story form (b) learning the attributes of verdict categories and (c) reaching a decision "through classification of the story into the best-fitting verdict category."¹⁶⁶ Determining which story is "best fitting" is, of course, a categorization task.

Several studies suggest that jurors rely on prototypical scripts of crime scenarios, but it is doubtful that there are prototypical scripts for relatively general categories like "first degree murder."¹⁶⁷ Richard Weiner and his colleagues used cluster analysis to determine that the stories generated by subjects asked to imagine "first degree murder" scenarios had no less than 13 different themes (from "battered spouse" to "murder for hire"). They concluded that they had found a "much more complex system of organization than anticipated by schema or prototype theorists."¹⁶⁸ Perhaps more plausibly, their choice of a category at a rather high level of generality ("first degree murder") accounts for their failure to find a single associated prototypical script. It might still be the case that the average juror has prototypical scripts for less general categories, like "murder for hire."

In any event, it is clear that people possess the capacity to categorize scripts, stories and narrative at a very high level of generality, based on

¹⁶⁴ Roger C. Schank & Robert P. Abelson, *Knowledge and Memory: The Real Story*, in 8 ADVANCES SOCIAL COGNITION, 1-85 (Robert S. Wyer ed., 1995). For a more popularized account, see Schank's *Tell Me A Story*, 147-188 (1990).

¹⁶⁵ The field of discourse psychology is chiefly concerned with the processing of narrative discourse, but there has been limited interaction between discourse psychologists and cognitive psychologists working on categorization. See, e.g. Arthur C. Graesser et al., *How Does the Mind Construct and Represent Stories*, in NARRATIVE IMPACT: SOCIAL AND COGNITIVE FOUNDATIONS 229 (Melanie C. Green et al. eds., 2002).

¹⁶⁶ Nancy Pennington & Reid Hastie, *Explaining the Evidence: Tests of the Story Model for Juror Decision Making*, J. PERSONALITY & SOCIAL PSYCHOL. 189, 190 (1992). For a recent application of the story model to juror decision making in a civil context, see Jill E. Huntley and Mark Costanzo, *Sexual Harassment Stories: Testing a Story-Mediated Model of Juror Decision-Making in Civil Litigation*, 27 LAW & HUM. BEHAV. 29 (2003).

¹⁶⁷ Richard L. Weiner et al., *The Psychology of Telling Murder Stories: Do We Think in Scripts, Exemplars or Prototypes?* 20 BEHAV. SCI. LAW 119, 121 (2002).

¹⁶⁸ *Id.* at 135.

perceived similarities. Anthony Amsterdam and Jerome Bruner have produced a compelling exposition of the role of narrative in legal categorization.¹⁶⁹ They come at narrative in the tradition of scholars in the humanities. They adopt an "austere definition" of narrative, as requiring both "a cast of human-like characters, beings capable of willing their own actions, forming intentions, holding beliefs, having feelings" and "a plot with a beginning, a middle and an end, in which particular characters are involved in particular events," a plot that has several distinctive features.¹⁷⁰ Amsterdam and Bruner's "austere definition" is not quite austere enough for my purposes. I would argue, rather, that it describes a useful prototype of the categories of narrative of the kind that concern them, and a more general prototype, perhaps, of "especially interesting and compelling narratives." While Amsterdam and Bruner come at narrative from the humanities, where matters of complexity and nuance are especially salient, for our purposes there may also be some value at coming at narrative from a more elemental level, as a form and extension of schemas, scripts, and theory-theories of categorization. Amsterdam and Bruner interpret narrative as occasioned by the violation of expectations carried in a script. To use their example, "you do not tell about a visit to the restaurant unless something not in the [restaurant] script occurs."¹⁷¹ What this suggests, however, is not that narrative has some different cognitive representation than scripts, but that some scripts are not sufficiently interesting to occasion retrieval and recounting.¹⁷²

Scripts and narratives have a role in the categorization of objects as well. To some extent, the "assault weapons" category falls out of narratives: the Stockton schoolyard shooting and other similar tragedies. To some degree the legislature is saying, "We don't want THAT to happen again," and THAT has by now, sadly, become a script: A deranged individual, often with an acute fascination with guns, acquires an arsenal that includes especially powerful weapons. Note that this script (or, alternatively, narrative) has a slot for "weapons used in mass killings."

Whether an object is or is not an assault weapon, in the United States at a particular point in history, is thus to some extent socially and historically contingent. One way of categorizing objects as assault weapons

¹⁶⁹ AMSTERDAM & BRUNER, *supra* note 20.

¹⁷⁰ *Id.* at 113. Steven Winter offers a more complicated description of the concept of narrative itself as an Idealized Cognitive Model. WINTER, *supra* note 133, at 106-113.

¹⁷¹ AMSTERDAM & BRUNER, *supra* note 20, at 121.

¹⁷² From another perspective, narratives are to scripts as mental models are to schemas: particular instantiations of more general patterns.

emerges from an understanding of the theories in physics discussed above. And, if legislators were functioning purely as scientists, indifferent to the social and political context of their actions, a theory-based approach to categorization might make sense. But, of course, legislators are responding not merely to a cold calculus of the probability of physical injury. In a very real sense, legislators are responding to narrative itself: the sensational stories of a handful of mass killings -- with weapons of military rather than hunting ancestry. On one level -- the level of physics and the calculus of risk -- it seems irrational to ban weapons largely on the basis of appearance while automatic shotguns can be purchased at most WalMart stores. On the level at which legislators -- and law -- operate, however, there are other considerations that can only be understood by reference to the particular narratives giving rise to the legislative action.

Law aims not merely to set up predictable algorithms of punishment and reward, but also to communicate, to *answer* the protagonists in the narratives that animate legislative action in the first place. If these people used pink shotguns or machine guns with graceful curves, then the story -- and the legislation -- would be different. This is not to say that such motivations are somehow irrational. As noted earlier, apart from the criminal and deranged, there is a distinct subculture in which these military style weapons are a salient feature. Legislation banning weapons in this category communicates disapproval of the subculture as well as the weapons.

Whatever the motivation, legislating categories that arise primarily from narratives are an exceedingly challenging enterprise, for two related reasons. First, there are infinitely many sequences of meaningful human action, even given the existence of scripts, because scripts combine in innumerable ways. Second, though we have made some progress, our general theories of human action are in predictive and explanatory power vastly inferior to our theories about the natural world. We are thus left with the more rudimentary means of specifying categories: by providing exemplar stories and names for the similarities we perceive among them. In the real world, this is work better suited to common law judges than legislatures. The interpretations of judges may occasionally be corrected through the legislatively overruling of a judicial interpretation.¹⁷³

In principle, however, a legislature might undertake the same kind of work engaged in by the array of appellate judges in the jurisdiction: considering a range of narratives, historical or hypothetical, and then

¹⁷³ See, e.g., Dan Bussel, *Textualism's Failures: A Study of Overruled Bankruptcy Decisions*, 53 VAND. L. REV. 887 (2000) (study of legislative overrulings of Supreme Court cases interpreting the Bankruptcy Act).

indicating whether or not the particular narrative conforms to the intended legislative category. Of course, the real world, particularly the world containing active agents seeking advantages with respect to statutory categories, is much more variable than the worlds any set of legislators might imagine in advance. But a legislature might, in effect, decide the first instances of statutory interpretation itself. Although not often done in enacting statutes, this is precisely the method often adopted by the Internal Revenue Service in explaining categorical rules in the regulations implementing the tax code.¹⁷⁴

7. A Note on Categories of Categorization Theories

I have presented the three modal types of theories of categorization as if they were contending competitors for approval in cognitive science. For much of the past twenty years, that has been the case. Each theory and particular variants have had strong advocates. Experiments have been designed, then redesigned, to test the predictive power of one theory as against another. As noted earlier, as a general model of categorization, the leading contender of only a few years ago, prototype theory, has lost favor. The most recent scholarship in these areas suggests that a common prototype for scientific progress – weeding out contending theories as against the evidence -- is in this instance wrong. It may well be that there is no single theory of categorization that fully explains what all people do in all cases. For example, as children learn categories, they seem to move from the learning of prototypes toward theory-based categorization. It appears that over the process of category learning, adults make use of prototypes during the early phases of learning and then move toward exemplar-based methods.¹⁷⁵ Experts make categorization decisions on different bases than do novices.¹⁷⁶ There are also cultural and gender differences in categorizations: on average, Western Europeans are more likely to categorize by rule; East Asians by overall similarity.¹⁷⁷ People may rely on one strategy of categorization for easy cases and another for hard cases. As often happens in the case of the workings of evolution's Rube Goldberg contraption that is the human brain, many things are going on at the same time. As noted earlier, it appears that different strategies of categorization may be carried out in completely different parts of the

¹⁷⁴ See, for example, 26 C.F.R. §56.4911-3 regarding expenditures by nonprofits for lobbying and giving 12 detailed examples of the application of the rules to particular scenarios.

¹⁷⁵ Smith & Minda, *supra* note 147 and authorities therein cited.

¹⁷⁶ Pamela T. Hardiman et al., *The Relation Between Problem Categorization and Problem Solving Among Experts and Novices*, 17 MEMORY & COGNITION 627 (1989).

¹⁷⁷ RICHARD E. NISBETT, *THE GEOGRAPHY OF THOUGHT*, 144-146 (2003).

brain.¹⁷⁸ Finally, theory-based categorization can in the end only be as successful as is the underlying theory. This does not prevent us from using folk theories, of course, but these tend by definition to eventually run afoul of how the world actually works.

8. On Domain-Theorization

I suggested above that the difficulty in categorizing was in part a function of how well or poorly the domains of categorization are theorized. It is now time to be as explicit as possible about what I mean, within the context of the present Article. The notion of "theorization" is itself vague. Addressing that vagueness requires a minor detour into the potentially bottomless recursive bog of meta- theory.

But let us begin with simple observations: Two carpenters arguing, first about carpentry and then about love. An argument about the appropriate angle to cut a beam can be settled. In very ancient times, this might have required the cutting of two beams and a test to see which would fit. But at least since the time of Euclid, carpenters have had the ability to settle most such disputes with geometry, without lifting a saw. But the same two carpenters will have more difficulty deciding whether a co-worker's relationship involves real love or mere infatuation. We have little difficulty agreeing that our theories of physical space are more complete than our theories of love.

What it means to have a theory of something is one of the central problems of the philosophy of science, which is chiefly concerned with deciding -- at a meta-theoretical level -- what constitutes an adequate set of explanations: a theory.¹⁷⁹ An extended excursion into the philosophy of science is unnecessary here because our concerns are much narrower -- the communication of category rather than the validity of a given categorization. Consider, again, a law against witchcraft, with "witch" being operationalized as "people who float when placed in water." With or without the ornamentation of terms like "specific gravity," such a law provides clear guidance to the citizenry and judges of the future. The problem of legal categorization is solved, even if many other problems are exacerbated or ignored.

¹⁷⁸ See sources cited *supra* note 131. See also Edward E. Smith and John Jonides, *The Cognitive Neuroscience of Categorization*, in THE NEW COGNITIVE NEUROSCIENCE, available at <http://cognet.mit.edu/Gazzaniga/Entry/chap70.html>, <visited 1/28/01>

For our purposes, then, we need a descriptive rather than a normative or evaluative meta-theory. Two currents within the contemporary philosophy of science offer a useful point of departure. In his recent *Science Without Laws*,¹⁸⁰ Ronald N. Giere argues that the crucial representations in many scientific theories are not the systems of propositions that concerned positivists, but rather the *models* scientists construct of aspects of the real world:

The question for a model is how well it "fits" various real-world systems one is trying to represent. One can admit that no model fits the world perfectly in all respects while insisting that, for specified real-world systems, some models clearly fit better than others. The better fitting models may represent more aspects of the real world or fit some aspects more accurately, or both. In any case, "fit" is not simply a relationship between a model and the world. It requires a specification of which aspects of the world are important to represent and, for those aspects, how close a fit is desirable.¹⁸¹

In Giere's metatheory, theories are like maps; indeed, sometimes theories *are* maps.¹⁸² The paradigm example for Giere is the theory of plate tectonics in geology. Once we have a model -- a three dimensional "map" -- of continental plates "floating" on a viscous substrate, a great many previously puzzling phenomena fall into place: the shapes of the current continents and why continental boundaries seem in some instances to "fit", the spreading of the deep ocean floor, the patterns of magnetism trapped in once molten rocks, and so on. Giere makes a reasonable case that in most of science it is the models that scientists create and test, rather than the equations and propositions that describe those models (and, only indirectly, the "real world"), that constitute the core of scientific theory.

From the perspective of this representationalist meta-theory, how fully a domain is theorized is a matter of how accurate and complete the set of maps and other models we possess with regard to the domain. Because our maps of the familiar physical world can be extremely accurate, a statute referring to a physical boundary can be quite precise. There is little doubt about where Arizona ends and California begins. But imagine a legislative action with respect to mineral rights on a large asteroid of vaguely understood shape and size. Until we have a model or map of the asteroid,

180 RONALD N. GIERE, *SCIENCE WITHOUT LAWS*, (1999)

181 *Id.* at 92-93

182 *Id.* at 24-25

any designation of a portion of the asteroid's surface will necessarily be quite vague.¹⁸³

To continue the analogy, in the various domains of interest to people and legislators we have quite disparate collections of relevant maps and models. Some (perhaps most of the sciences) are more like our modern collections of maps of the earth: fairly precise and generally consistent with each other. Others seem more like the implicit maps in ancient myths, referring only to distant lands and general directions. In particularly ill-theorized domains (perhaps most of the humanities as well as those things that most deeply concern us in our personal and family lives) we lack the equivalent even of a single coherent ancient myth, but have instead have only a buzzing, blooming confusion of inconsistent tales.

We have no certain metric by which to compare the relative completeness and accuracy of models across domains. But, perhaps self-referentially, we can model collections of models. I can tell you that the collections of paper and electronic maps in the UCLA library for the earth's surface far exceeds the collection of maps of the surface of Pluto, and that maps of the "surface" of Jupiter don't seem to exist because the notion of "surface" does not map well to that giant blob of gas.¹⁸⁴ But how do we compare the set of models in geography or space science to the sets of models in psychology or sociology? Reference to the quantity of volumes in the relevant sections of the library is unavailing. Indeed, one of the hallmarks of well-theorized domains is that there is at any one time only one "standard model" (or at most a handful of models contending for that designation) that is regarded as internally coherent and substantially complete. There is always contention at the margins, which may sometimes lead to the undoing of consensus about the "standard model," and the superposition of a new "paradigm".¹⁸⁵ But the existence of a large number of inconsistent contending models or theories generally marks the lack of a "standard model" and a less well-theorized domain.

If the sheer number of models in the domain is an unreliable guide to the degree of theorization, how else might we compare the relative theorization of disparate domains? Another line of work in philosophy and the philosophy of science offers the promise of a metric: the coherence of the models in the domain. To use a seemingly vague term like coherence to

¹⁸³ In principle, we could decide *how* to divide the surface of an asteroid of unknown shape, provided the asteroid is spinning. The axis of rotation defines, potentially, an asteroidal north and south and the means to arrive at an equator. From there it is but a short step to superimpose the notions of longitude and latitude on the asteroidal surface.

¹⁸⁴ For the details and lovely photos, see <http://pds.jpl.nasa.gov/planets/welcome/jupiter.htm> [visited March 6, 2004].

¹⁸⁵ THOMAS S. KUHN, *THE STRUCTURE OF SCIENTIFIC REVOLUTIONS* (1996).

unravel the notion of vagueness may seem a risky endeavor. But recent work by the philosopher Paul Thagard and others has both brought a new rigor to the notion of coherence, and explained the role coherence plays in explanation itself. A full exposition of these ideas is not possible within the constraints of this Article. But Thagard's idea is roughly this: "explanatory coherence" refers to a method of determining the acceptability of a proposition within a set of other propositions by examining the web of relations (consistency, analogy, etc.) among them, which is a kind of constraint satisfaction problem described in Section IV.

Notably, constraint satisfaction models of coherence and of theory itself accord well with the intuitions underlying the "theory-theory" of concepts and categories referenced in Section *, above. For example, as indicated earlier, a standard example of a constraint satisfaction problem is a crossword puzzle, which requires imagining and testing words that meet the constraints that constitute the puzzle. Quite analogously, referring to a category that would include "dogs, insurance policies, birds, deeds, jewelry, photographs, cats, bonds, children, passports" presents a similar puzzle: imagining instances in which such a grouping fit sensibly together.

Thagard and Kirsten Verbeurgt have demonstrated that constraint satisfaction models of coherence have certain formal mathematical properties.¹⁸⁶ Assuming we can reduce the content of a theory or intellectual domain into a set of propositions, it is theoretically possible to compute the overall coherence of that theory or domain by the methods Thagard and Verbeurgt developed. As a practical matter, their formalization of coherence offers at least a consistent way of thinking about the relative coherence of say, quantum physics and theories of historical development as intellectual domains.

If this seems a stretch for our judiciary, it is important to note that courts are not unfamiliar with the problem of assessing the degree of coherence within an intellectual domain. This is essentially the assessment that a court must make in determining whether proffered expert evidence should be admitted at trial. Under the *Daubert*¹⁸⁷ standard, federal judges are tasked with assessing "whether the reasoning or methodology underlying the [proffered expert] testimony is scientifically¹⁸⁸

¹⁸⁶ Paul Thagard & Karsten Verbeurgt, *Coherence as Constraint Satisfaction*, 22 COGNITIVE SCI. 1 (1998).

¹⁸⁷ *Daubert v. Merrell Dow Pharm., Inc.*, 509 U.S. 579 (1993).

¹⁸⁸ *Id.* at 592-93. The *Daubert* factors may be applied to non-scientific expert testimony as well, not merely that offered by scientists. *Kumho Tire Co. v. Carmichael*, 526 U.S. 137 (1999).

valid and whether that reasoning or methodology properly can be applied to the facts in issue.”¹⁸⁹

The Supreme Court indicated that judges must evaluate such factors as (1) “whether a theory or technique . . . can be (and has been) tested,”¹⁹⁰ (2) “whether the theory or technique has been subjected to peer review and publication,”⁽³⁾ “the known or potential rate of error,” and (4) “the degree to which the theory is widely accepted” within the relevant community. The explicit *Daubert* factors do not explicitly take account of the location of the proffered evidence within a coherent set of theories and concepts, but each of the *Daubert* factors is an indirect measure of the degree of coherence in the domain, for the following reasons: (1) *Validity*. Although anti-foundationalist philosophers of science would disagree with the formulation of scientific validity adopted by the court, most would agree that potential theories and propositions should be and are assessed with regard to their coherence with the data and with other provisionally accepted propositions. (2) *Peer review and publication*. The case here is a bit weaker, given that what gets published in the journals of some disciplines appears largely a matter of fashion and the trends of the time. Nevertheless, the converse is generally true: articles that are completely incoherent with the prevailing fashions and paradigms have little chance of publication – except in disciplines where novelty is valued for its own sake. (3) *Error rates*. The only way to arrive at error rates is to compare observed values with some standard and to compute the consistency/correlation/coherence between the two. (4) *Acceptance*. The paradigm prevailing at any time (Kuhn’s “textbook science,” or what is referred to as “the standard model”) is prevailing precisely because of the way in which it coheres with observations and other theories. Indeed, Paul Thagard’s constraint satisfaction models of coherence have been used to model closely the kinds of paradigm shifts that Kuhn described.¹⁹¹

¹⁸⁹ *Daubert*, 509 U.S. at 592-93. As the Court made clear later in *Kumho*, the *Daubert* factors are not exhaustive or meant to be applied mechanically, trial judges retaining considerable discretion in how to determine the question of reliability. Justice Scalia’s concurrence in *Kumho* observes that this “... is not discretion to perform the function inadequately. Rather, it is discretion to choose among reasonable means of excluding expertise that is *fausse* and science that is junky.” *Id.* at 159. Apparently trial judges are equipped to discern the plain meaning of “*fausse*” and “junky.”

¹⁹⁰ *Id.* at 593. Notably, the Supreme Court relied on the philosophy of science of Carl Hempel and Karl Popper, the leading positivists who now probably represent only a significant minority view among philosophers of science. *Ibid.* Thus, the Supreme Court’s own reliance on experts in the philosophy of science fails at the level of metatheory the test it adopts at the level of theory.

¹⁹¹ PAUL THAGARD, *CONCEPTUAL REVOLUTIONS* (1992).

A judge assessing the vagueness or precision of a statute has, of course, a quite different task than a judge assessing the validity of potential expert testimony. First, there is the matter of time. A judge ruling on the expert validity question is concerned with the state of knowledge in the domain at the present time. A judge assessing how a legislature might have been more precise must assess the state of knowledge at the time the statute was adopted. Moreover, the questions concerning expert testimony are generally more narrowly framed, in terms of the proposition to which an expert proposes to testify. Statutes, on the other hand, can seek to regulate quite broadly, in ways that touch on the potential validity of thousands of propositions. These differences are, however, differences in the magnitude of the required judicial enterprise rather than its possibility.

VI. IMPLICATIONS/CONCLUSIONS

How might legislators, judges and legal scholars take account of what science has learned about categorization? First, assuming that it matters what legislators mean to say, then a reasonable pragmatic goal of bringing cognitive science to law is the reduction of interpretive errors, at least when it comes to specifying legislative categories. This is not to say that legislatures may always prefer precision.¹⁹² One of the possible benefits of the kind of analysis sketched above is greater clarity about the costs and benefits of specifying categories to varying degrees of precision. And, where precision is indeed the aim, the cognitive science sketched here suggests some means of better achieving it.

Second, by attending closely to the contexts and domains in which law operates, we may enrich the jurisprudence of legislative interpretation. In particular, whereas much jurisprudential energy has been dissipated in debates about the desirability and consequences of various generalized approaches to interpretation, the science of human understanding and communication of categories suggests that the search for universal principles of interpretation is likely to fail, and that it must generally attend to the specific substantive domains in which law operates. It is one thing for Justice Scalia to insist on a degree of reasonable precision in a fully theorized domain in which such precision can be achieved at low cost. But to insist on a similar level of precision in domains that are less fully understood is merely to deny the legislature the power to effectively act at

all in these areas. In addition, the relatively new cognitive science of categorization may affect old debates about the forms of law, particularly ancient debates about the relative merits of "rules vs. standards." Here again, cognitive science suggests that context and communicative content matter in ways that these debates have long ignored.

A. Better Legislative Categorizations Through Science?

I began this study as a kind of extended thought experiment: would it make a difference if legislators and judges knew some of what cognitive science now knows about categorization. The most reasonable answer is probably, perhaps, a little. Certainly the science of categorization has no direct prescriptions for legislators. Indeed, the fact that several theories still contend for universal acceptance is enough to give us pause. As a pragmatic matter, however, the existence of these contending theories -- each of them with some empirical support -- may suggest how best to take account of them: use them all. As applied to many of the situations we encounter in life, all of the major theories appear to have some validity. Indeed, experimental psychologists spend a good deal of energy trying to devise experimental situations in which one theory will survive and the others will fail. This suggests that pragmatists in law ought to borrow another idea from the methodologists in social science (who adapted it from navigators): triangulation.

Whether in navigation or social science, triangulation is another variant of the constraint satisfaction problem discussed in Section IV: using multiple methods to measure the same feature or phenomenon.¹⁹³ The navigator feels more certain of his location when plots from different landmarks intersect very near the same point. The social scientist is more confident of conclusions supported by cross-sectional surveys, interviews, and the close observation of behavior. We also use triangulation or constraint satisfaction techniques to improve accuracy in communication. The reason that sailors, pilots and others refer in radio communications to letters of the alphabet as Alpha, Baker, Charlie and so on, is that these words carry information that constrains interpretation of their initial letters.

¹⁹³ See, e.g., NIGEL G. FIELDING & JANE L. FIELDING, LINKING DATA (1986); Norman K. Denzin *Strategies of Multiple Triangulation*, in THE RESEARCH ACT: A THEORETICAL INTRODUCTION TO SOCIOLOGICAL METHODS 234-247 (3rd ed. 1989); Todd Jick, *Mixing Qualitative and Quantitative Methods: Triangulation in Practice*, 24 ADMIN. SCI. Q., 602 (1979).

Applying these principles to legislation, a legislature keen on constraining the possible misinterpretations of categorical content might *consider* proceeding as described below. Let me concede at the outset that some of these methods of legal categorization will strike readers as bizarre and unthinkable. For now, consider the following merely as a thought experiment about how a legislative specification of "assault weapons" could be made maximally precise:

1. Utilize a theory, partial or complete, underlying the category, if there is one. For example: "The purpose of this legislation is to limit the availability of weapons of great lethality that are not generally used for hunting purposes. Accordingly, the first component of the category definition of 'assault weapons' is 'firearms capable of firing more than 100 rounds per minute with more than 100 foot pounds of muzzle energy, excepting only shotguns holding 6 or fewer shells in the magazine'".
2. Identify the prototype that best fits the category. For example: "The prototype of the weapons we intend to regulate is the AK-47."
3. Exemplars and Near-Miss Non-Exemplars with "typicality" or "category gradient of membership" judgments of legislators, expressed in numeric form. For example: "By the vote reflected in Column A below, we mean to include the weapon in Column B. By the vote reflected in Column D we mean to exclude the weapon in Column C."

A	B	C	D
98:3	AK-47	Ruger 10-22	4:96
90:10	UZI	Model 12 Win.	12:88
68:32	BAR	Remington 870	33:67

4. A trained neural network or other classifier system. For example, the array of features (caliber, firing rate, barrel length, and so on) of each weapon in an entire current catalog of small arms could be included in the "training set" of a neural network. "Training" would be provided by providing feedback on categorization by means of votes like those reflected above. The nonlinear patterns present in what "assault weapons" appears to mean to legislators would be learned by the network, which would thereafter classify new candidate weapons.
5. Assess the costs of doing each of the foregoing, the possible benefits of both precision and imprecision, and send along an accompanying message that indicates the results of the decisions and preferences for how issues of underinclusion and overinclusion should be resolved. In weighing

those risks, the legislature might want to make use of the "Receiver Operating Characteristic" framework described in Section IV. The legislature may not be able to calculate or specify such a curve, but the very framework focuses attention on the tradeoffs between overinclusion and underinclusion. The resulting statute might merely be accompanied, as part of the legislative history, with an explicit statement on the issue, something like: "Given the limited amount of time available to specifying the category of "assault weapons," we have elected to utilize only methods 1 and 2 for specifying the category of "assault weapons." Given our assessment of the relative costs and benefits to society of not including particularly dangerous weapons, in comparison to the costs and benefits of including weapons that are less dangerous and more utilized for sporting purposes, it is our intention that all reasonable uncertainties be resolved in favor of including firearms within the category."¹⁹⁴ One can even imagine using similar methods of exposition in describing the categories, say, of "acceptable overinclusion" and "unacceptable overinclusion."

There is certainly good reason to suggest that legislators ought to adopt theory-based legislation when they have the means to do so, and that there may be some value as well in adopting the other methods suggested above. But the notion that a real-world legislature will spend time taking votes on individual weapons for the purposes of instructing a neural network seems preposterous. Perhaps. But consider that other legislatures, including the Congress of the United States, have adopted legislation containing extensive lists of weapons. It is doubtless true that few legislators inspected this list before voting on it, but that is an issue of appropriate delegation and procedure for the legislature to determine.

Certainly, one can imagine areas of legislation in which the stakes are sufficiently high that a rational legislature might want to expend enormous effort in categorizing situations in advance of rule-application to particular instances. For example, under a Constitution adopted in a world in which events moved at a less rapid pace, the Congress of the United States effectively possessed the exclusive power to declare war. Yet the strategic forces of the United States are now poised to deliver an unimaginable level of death and destruction in a matter of minutes, upon the rapid classification by the President alone of a set of circumstances as one requiring that response, in accordance with the strategic defense policy of the United States. Congress has effectively delegated this power, surely

¹⁹⁴ Whether inserting such an interpretive rule into a statute is constitutionally permissible is an interesting topic outside the scope of this article. *See, generally*, Nicholas Quinn Rosenkranz, *Federal Rules of Statutory Interpretation*, 115 HARV. L. REV 2085 (2002).

more awesome than that entailed in any previous declaration of war, out of felt necessity. It is not unthinkable, however, that Congress might adopt a different approach, one that would provide more guidance to the military and executive and yet allow for rapid reaction to complex circumstances. Some of the methods suggested here, including the neural network trained on combinations of scenarios by Congress, would permit just such a rapid calculation of likely Congressional intent, not as a replacement for executive decision making, but as an aid to it.¹⁹⁵

B. Implications for Theories of Statutory Interpretation.

Perhaps the principal value of engaging in these thought experiments is not to suggest practical means by which legislatures can enact more precise statutory categories, but rather to suggest to those who interpret statutes a framework for taking account of the context of the substantive domain in which legislation operates. Debates about statutory interpretation tend to be couched in general terms about "the law," illuminated by particular examples chosen to buttress the arguments for the approach to statutory interpretation being advanced. To be sure, there are advocates for "contextualism" of various forms. Of the coherent approaches to interpretation -- recently characterized as the "old war-horses of interpretation: textualism, literalism, plain meaning, original intent, purpose, contextualism, canons of construction, pre- and post-enactment legislative history, imaginative reconstruction, counter-majoritarianism, statutory stare decisis, and dynamic interpretation"¹⁹⁶ -- only the most extreme versions of the first four would deny the relevance of some form of context. Modern textualists recognize that the meaning of a statute, like the meaning of any other text, can only be decoded "according to the common social and linguistic conventions shared by the relevant community".¹⁹⁷ Contextualists vary considerably in what other aspects of context should matter. Those characterized by John Manning as "strong purposivists" would have judges attend to a "statute's overall tenor," including the historical context when the statute was passed, society's values, patterns of policy for related statutes, and statements in the legislative history.¹⁹⁸ For William Eskridge, the relevant contexts are social and political, and context

¹⁹⁶ Shirley S. Abrahamson & Robert L. Hughes, *Shall We Dance? Steps for Legislators and Judges in Statutory Interpretation*, 75 MINN. L. REV. 1045, 1053 (1991)

¹⁹⁷ John Manning, *Textualism and the Equity of the Statute*, 101 COLUM. L. REV. 1, 16 (2001).

¹⁹⁸ *Id.* at 10-11.

itself thus evolves: "[T]he meaning of a statute will change as social context changes, as new interpreters grapple with the statute, and as the political context changes. . . ."199 Pierre Schlag would include a vast "web" of interrelated contexts, even in interpreting so seemingly simple a statute as the classic prohibition of "vehicles in the park":

That is because, as a general word in a legal rule, the term draws its meaning from the interweavings of all manner of webs - webs that are often described as linguistic, cognitive, moral, political, institutional, or cultural. In the rule, the meaning of the term "vehicle" is inscribed in tacit understandings of parks; legal rules; the effects of legal rules; the roles and possibilities of legal rules within the hierarchies of sources of law; the "public" meaning of legal rules for citizens and public officials; and the meaning of legal rules in light of juridical concepts of excuse, justification, prosecutorial discretion, and much more. We are not just talking about parks and vehicles here; we are talking about parks and vehicles in a legal rule in a legal system in a particular culture."200

Understandably, some scholars and judges are troubled by the potential malleability of interpretation seemingly so loosely constrained by "context" so variously described. The reaction, most often associated with Justice Scalia, seeks to eliminate considerations of nonlinguistic context to the maximum extent possible.

As a general matter, we may share Justice Scalia's instinct that strict adherence to the texts of law will encourage legislators over time to be more attentive to the laws they write, with a resulting increase in rule precision. We may share Justice Scalia's distaste for legislation that confers so much discretion on judges that they feel relatively unconstrained in applying their own preferences. But to insist on an arbitrarily high degree of rule precision or categorical specificity in poorly theorized domains, without reference to the possibility of precision, is to deny the *possibility* of law in these domains, with far more inhibiting effects than either due process or separation of powers doctrine requires. Moreover, even where great precision is possible, it is never costless, requiring legislators to determine how to allocate the scarce resources of their time and attention. Democratic theory suggests that making such a determination should not be left, in the first instance, to judges.

199 ESKRIDGE, *supra* note 35, at 199.

200 Pierre Schlag, *No Vehicles in the Park*, 23 SEATTLE U. L. REV. 381, 387(1999).

Like the interpretive stances to which it is a reaction, textualism advances a generalized, universal approach to interpretation. As Cass Sunstein and Adrian Vermeule observe,

[t]ypically, interpretive issues are debated at a high level of abstraction, by asking questions about the nature of interpretation, or by making large claims about democracy, legitimacy, authority, and constitutionalism. But most of the time, large-scale claims of these kinds cannot rule out any reasonable view about interpretation.

Sunstein and Vermeule would replace the misguided search for universal "first best" principles of interpretation with a contextual analysis sensitive to the institutional capacities of legislatures, agencies and courts, as well as to the dynamic effects of particular approaches of both private and public actors. In one sense, this is but another form of contextualism, adding another entry to Schlag's catalog of contexts. But it is more than that, because it springs from the insight that the choice of interpretive rules must itself be sensitive to context.

I share with Sunstein and Vermeule a skepticism about the feasibility of *any* generalized abstract approach to statutory interpretation. But I would add to their insistence on an attention to institutional and dynamic concerns, the requirement that a sound approach to statutory interpretation be sensitive as well to both the possibility and the costs of statutory precision at the time of drafting, recognizing that these will vary dramatically across substantive domains, for all the reasons I have cataloged. The feasibility and costs of precision will vary even within the same institutional arrangements and over time, as the domain of legislation becomes more adequately theorized and the costs of precision therefore decline. These factors do not necessarily lend support to either textualism or any variant of contextualism, but rather add essential dimension to considerations of either.

In other words, there is merit to the argument that the rule of law is well served over time if judges hold legislation to an *appropriately* high standard of precision. What is appropriate will depend on many things; among them are the possibilities and costs of precision, as a consequence of the level of theorization of the relevant domains. In reviewing legislation a court might consider, in addition to the other factors elaborated earlier, the following: First, were there in existence plausible theories that would enable greater rule precision? Second, did the legislature consider the relative costs and benefits of the form of law it selected, given the state of knowledge in the domain? Third, was the legislature's choice of form of

law, given both the range of possibilities and costs and benefits, arbitrary or irrational, or violative of some other standard? Such a contextual approach is preferable to trying to apply a generalized rule to rules, or holding standards to an *a priori* standard. Such an analysis does not end questions of interpretation or fundamental jurisprudence, including how much deference judges should give to the decisions of legislators about *how* to enact law. But the science reviewed in this article provides some means for framing those questions with more precision.

C. Beyond Statutes.

Although this Article has focused on problems of categorizing and communication in the context of statutes, some of the approaches described here are potentially applicable to a far wider range of issues in law and legal scholarship. The problem of categorizing possible objects or situations *ex ante* is pervasive throughout law. Parties to a contract must consider how to describe what will constitute material breach of the contract, and negotiate their intentions into words to be later interpreted not only by themselves, but by a judge or arbitrator. A trustor or testator must consider how to describe those situations that will remove a contingency 50 years hence, as the result of an interpretive act by a trustee or judge. All law, not merely statutory law, exists in time. Much of the same science, and the same technologies for precision I have suggested for consideration in statutory drafting, may also be of some use whenever the most generic problem in law arises: how shall we best say what we mean, when what we mean has consequences, and our words are to be interpreted in the future?