

The Rules of Abstraction

by Douglas Glen Whitman

ABSTRACT: Friedrich Hayek's work on spontaneous order suggests that the emergence of a spontaneous order requires the existence of *abstract* rules of conduct. But how much abstraction is required? Abstraction exists on a gradient, from the highest specificity (pertaining to particular persons and narrowly defined circumstances) to the highest generality (pertaining to all persons in all circumstances). If rules create order by coordinating expectations, either end of the spectrum is undesirable; the most specific and the most abstract rules fail to provide decisionmakers with useful guidance. This article argues that rules that foster coordination must be characterized by an *intermediate* degree of abstraction. This conclusion will be explained and applied to law, language, and etiquette in order to draw out the similar character of rules across various contexts. The article concludes by discussing four properties that rules of intermediate abstraction must also possess to foster spontaneous order.

KEY WORDS: abstraction, spontaneous order, rules, law, language, etiquette
JEL CODE: B53, K0, Z1

Friedrich Hayek's work on spontaneous order, particularly in *The Constitution of Liberty* (1960) and *Law, Legislation and Liberty, Volume 1: Rules and Order* (1973), emphasizes the crucial role of abstract rules of conduct in creating the conditions for such an order. Although Hayek focuses mostly on legal rules that create fertile ground for the emergence of a market economy, he strongly implies that abstract rules are a key ingredient of *any* spontaneous order. He specifically cites language, manners, and morals as areas in which people learn, often unconsciously, an abstract set of rules that guide their behavior (1973, p. 19).

But what sort of abstraction is required? How abstract must a rule be to satisfy Hayek's criterion? Is abstractness a necessary or sufficient condition for rules to foster a spontaneous order? And if abstractness is not sufficient, then what other conditions must the rules satisfy?

In this article, I hope to offer a partial answer to these questions. Hayek's work offers numerous hints about the properties of the abstract rules of a spontaneous order, but the hints are scattered. Moreover, in the passages where Hayek is clearest about the necessary properties of rules, he is usually addressing the specific case of law – and not just any law, but the kind of law that creates a free society.¹ My goal here is to extract from Hayek's work the characteristics of abstract rules that facilitate any kind of spontaneous order. To do so, I will need to draw on orders other than law and markets, especially etiquette and language.

The primary insight that emerges is the notion of an *intermediate degree of abstraction*. It is not enough for rules to be abstract, for they must have *neither too little*

¹ For instance: "There is only one such principle that can preserve a free society: namely, the strict prevention of coercion except in the enforcement of general abstract rules equally applicable to all" (Hayek 1960, p. 284).

abstraction nor too much. In much of his work, Hayek is concerned with countering the modern movement toward command-and-control regulatory systems, which often err in the direction of too little abstraction. “The nature of these abstract rules that we call ‘laws’ in the strict sense,” Hayek says, “is best shown by contrasting them with specific and particular commands” (1960, p. 149). He therefore places great emphasis on the need for high-level abstraction; for instance, he says that law in its ideal form “is directed to unknown people and ... is *abstracted from all particular circumstances of time and place* and refers only to such conditions as may occur *anywhere and at any time*” (1960, p. 150; emphasis added). The possibility of excessive abstraction gets little attention, although in some passages Hayek does seem to recognize its importance by implication.²

The problem of too much abstraction is just as serious as that of too little, I will argue, because it is in many respects *the same problem*: both extremes offer decisionmakers too little guidance about how they should act. If abstractness exists on a spectrum from total specificity to total abstraction, then the ends of the spectrum share similar problems, and the “sweet spot” for abstract rules lies somewhere in the center. In fact, an intermediate degree of abstraction seems to be a characteristic of *most* entities that go by the name of “rules,” whether or not they encourage the growth of a spontaneous order.

The purpose of this article is to explore the relationship between abstraction and the rules that foster spontaneous order. Part 1 of this article will explain the principle of intermediate abstraction in rules. Part 2 will apply that principle to law and language.

² For instance: “Such rules, presumed to have guided expectations in many similar situations in the past, must be abstract in the sense of referring to a limited number of relevant circumstances...” (Hayek 1973, p. 86). That the rules must refer to a *limited* number of circumstances – neither referencing every particular circumstance nor abstracting from all circumstances – seems to indicate that abstraction must be intermediate in character.

Finally, Part 3 will examine four additional features that (intermediately) abstract rules for a spontaneous order must satisfy: purpose-independence, extensibility to novel cases, tolerance for small deviations, and openness to the generation of new rules.

1. Intermediate Degree of Abstraction

1.1. Simple illustrations of the principle

One simple illustration of intermediate abstraction can be found in the rules of etiquette at the dinner table. One rule is to say “Please” when requesting a service from someone, such as handing you the butter; another is to say “Thank you” after a service has been performed; another is to keep your mouth shut while chewing; and so on. Now, imagine if we replaced all these rules with a single, highly abstract directive: “Be polite.” That directive would not provide much useful guidance. Lacking more information, the decisionmaker would have to decide for each and every table interaction what would be the “polite” thing to do.

On the other hand, what if we had a different prescription for every possible table interaction? One directive for when someone passes the butter, another for when someone passes the salt, another for when someone serves food onto your plate, and so on. The decisionmaker’s problem would be very similar to that which he encountered under the abstract “Be polite” rule. Any time he encountered a novel situation – and arguably, every situation is novel in at least some infinitesimal degree³ – he would have

³ O’Driscoll & Rizzo (1996, p. 61-62) argue that even if two situations are identical in all externally visible respects, they will still differ because the passage of time causes the memories and expectations that people bring to those situations to differ. However, this is a subset of the point I’m making here, which is simply that no two situations are exactly identical.

to decide the correct action, without much assurance that it's correct. (Would it be okay for him to use the passing-the-butter response the first time someone passed him I-Can't-Believe-It's-Not-Butter?)

So far we have only discussed dinner table etiquette. But we could imagine the highly abstract rule ("Be polite") being applied to all types of interactions, from proper treatment of the opposite sex to proper behavior in an elevator. Or we could imagine having a different highly-specific rule of etiquette for each and every different activity or interaction, however minutely defined. But we have neither of these. The actual rules of etiquette possess an intermediate degree of abstraction, neither so broad as to include all situations, nor so narrow as to differ for each and every specific situation.⁴ They identify abstractly-defined types or kinds of situation.

The same analysis applies to legal rules. A rule of tort liability that said "Do the right thing" or "Be careful" would not be very helpful in guiding behavior, unless the decisionmaker had some strong notion of what other people think is right in more narrowly defined categories of situation. But if the categories were too narrow, depending too much on the particular characteristics of the particular situation, the decisionmaker's dilemma would be the same. A more useful rule of tort liability tells the decisionmaker how to act in situations with an intermediate degree of abstraction; e.g., "Always yield to cars already on the freeway when merging." This rule picks out some relevant features of a situation (who is merging versus who is already on the freeway), while ignoring countless other features of the situation (relative speeds of the cars, types and models of cars, colors of cars, ages of drivers, number of occupants, ultimate destinations, and so on *ad infinitum*).

⁴ There are exceptions, which will be discussed later.

Actual tort law, and law generally, does rely on rules with an intermediate degree of abstraction. But there is a spectrum from perfect specificity to perfect abstraction, which means that intermediateness lacks clear boundaries. Although the law generally does not incorporate directives at either extreme, it does sometimes rely on modes of decisionmaking that veer in one direction or the other. In the direction of highest specificity, we have case-by-case decisionmaking: the notion that judges ought to consider the particulars of each case, decide based on those features, and refrain from drawing conclusions that stray too far from the particular facts before them. In the direction of highest abstraction, we have the ubiquitous standards of the common law – such as “reasonable man” tests, which direct judges and juries to decide a wide range of cases according to whatever factors would be relevant to a typical, decent person with rational faculties; and “balancing” rules, which ask judges and juries to decide based on an array of relevant factors without necessarily specifying the weight that should attach to each factor.

The virtues and demerits of such abstract decisionmaking guidelines have been discussed extensively in the legal literature on rules-versus-standards, which will be addressed later. For now, it is worth noting that standard-based decisionmaking and case-by-case decisionmaking have much in common; arguably, they are the same thing. In order to make decisions on a case-by-case basis, decisionmakers are likely to appeal, at least implicitly, to broad principles in order to decide what is relevant and what is not. In order to apply standards, decisionmakers end up having to decide each individual case based on its particular merits. Moreover, both forms of decisionmaking lack the certainty of rules with an intermediate degree of abstraction; this impedes the ability of agents

subject to the law to know how to behave in their own situations. To predict the effect of either pure case-by-case or pure standard-based decisionmaking, the agent must possess and process an untenable amount of information and predict the behavior of an untenable number of other agents. The informational burden and cognitive load are simply too great.⁵

Examples aside, I will now attempt to lay out the general problem. Although the word “rules” is sometimes used broadly to refer to any kind of directive for decisionmaking, or any regularity of behavior, it also has a less general connotation. Rules are *those directives that help people make decisions with some degree of certainty about which behaviors are acceptable (or expected) and which are not*. That is, the primary function of rules is to lend predictability to one’s own choices as well as the choices of others (including their responses to one’s own choices). Predictability results from a reduction in the amount of information individuals must collect and process, and it is desirable because it enhances individuals’ ability to make plans and to coordinate with each other.⁶ This does not mean that predictability is the *only* function or value that rules serve. Rather, what distinguishes rules from *other kinds of directives*, such as case-by-case or standard-based decisionmaking, is their effect on predictability.

⁵ Roger Koppl (2002, p. 533) adopts a similar notion of ruleness: “The object [such as an agent] follows a rule if its behavior exhibits invariance across some states but variance across others. It follows a rule if only a subset of variables determining the larger system influence the behavior of the object. The rule-bound object processes less information than it otherwise would. Similarly, the observer needs to process less information than the object unbound by any rule.” In other words, rule-governed behavior responds to some *but not all* aspects of the environment, thereby minimizing information costs.

⁶ There is a close affinity between my approach here and Ronald Heiner’s (1983) model, in which regularity of behavior (and thus predictability) is a means of coping with a gap between the competence of the decisionmaker and the difficulty of the decision problem (the “C-D gap”). In Heiner’s approach, a larger C-D gap corresponds to greater uncertainty. Both case-by-case and standard-based decisionmaking would tend to increase the difficulty of the decision problem.

This predictability – which is, of course, never perfect – is achieved by attaching specific outcomes or behaviors to *classes* of situations. The breadth of these classes corresponds to the degree of abstraction of the rules. If the classes are defined very narrowly, the rules lose some of their rule-like quality because of the difficulty of discerning which class includes one’s own situation. When class definitions depend on a virtually infinite number of characteristics of a situation, the human mind may fail to grasp them all. Some of the characteristics might be difficult to measure, and others might simply be forgotten or not conceived in the first place. As a result, the decisionmaker faces greater uncertainty when assessing the acceptability of his actions and in projecting their likely results. In addition, other agents will face greater uncertainty about how the decisionmakers will act (for instance, potential litigants will find it harder to predict both each other’s actions and judges’ actions).

Virtually the same problem emerges if the classes are defined too broadly; again, the rules start to lose their rule-like quality. This occurs because very broadly defined classes generally correspond to broader sets of possible outcomes, in order to allow a better match between outcomes and the particular features of the situation. (The possibility of very broad classes corresponding to narrow outcome sets will be considered later.) Here, the decisionmaker has little difficulty discerning the class into which his situation falls, but he still faces a knowledge problem: he does not know exactly how the general principles will govern in his case. He does not know which factors will be deemed relevant, and he does not know the precise values to be attached to these factors. In short, the decisionmaker still faces great uncertainty in determining the acceptability or likely outcome of his decisions.

Figure 1, below, depicts the emerging relationship between abstraction and “ruleness,” where ruleness refers to the quality of directives that guide behavior by reducing uncertainty of application. The straight line shows increasing ruleness (from left to right). The circular curve shows increasing abstraction (in a clockwise direction); the endpoints of extreme specificity and extreme abstraction approach each other.⁷ Those directives which are most rule-like will generally display an intermediate degree of abstraction, as shown at the right edge of the abstraction curve.

Figure 1: Abstraction & Ruleness

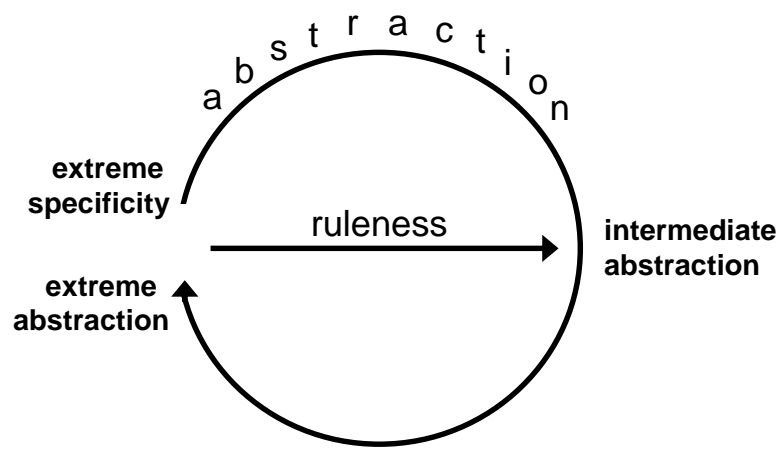
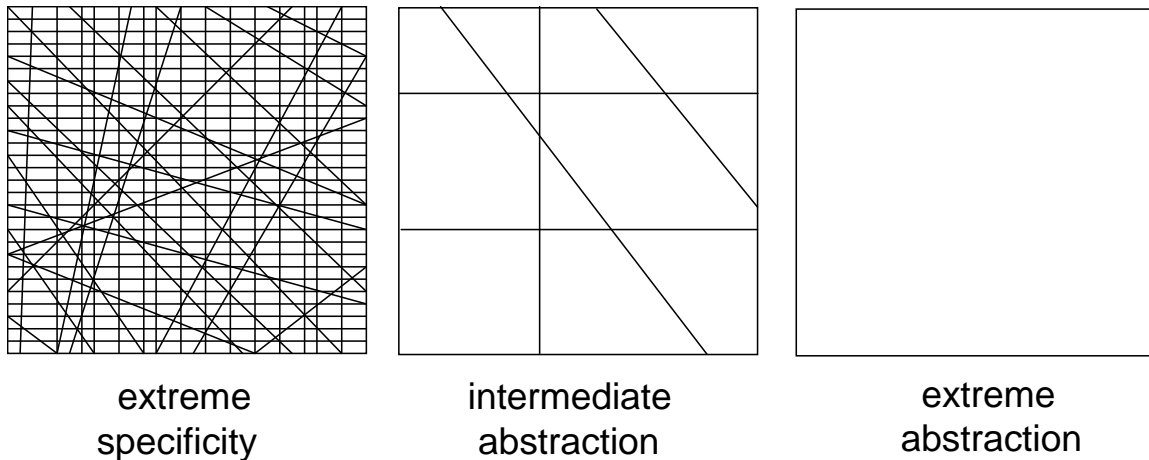


Figure 2 illustrates three different degrees of abstraction. The left panel is characterized by a very low level of abstraction (innumerable distinctions being made among cases), the right panel a very high level of abstraction (no distinctions made among cases), and the middle panel an intermediate level of abstraction (a few

⁷ An alternative figure showing the same relationship would place abstraction on the horizontal axis, ruleness on the vertical axis. The relationship would then take the form of an upside-down U-shaped curve. I have chosen the circular representation to emphasize the similarity of extreme specificity and extreme abstraction.

distinctions made among cases). The spaces between lines can be understood as the classes defined by rules. The uncertainty generated by a proliferation of low-abstraction rules corresponds to the difficulty of determining which space one's own case falls into; there are so many spaces, delimited by so many relevant factors, that it becomes difficult to distinguish one space from another – and therefore difficult to determine the acceptability of one's decision. That source of uncertainty does not exist in the case of extreme abstraction, but it is replaced by something very similar: although you know what class your case falls into (because there is only one), that is not sufficient to determine acceptability of your decision, as cases within this large class are treated differently based on all conceivably relevant factors.

Figure 2: Degrees of Abstraction



1.2. Why we rely on abstraction

The reason why most rules worthy of the name rely on intermediate abstraction is closely related to the reason why we use abstraction in general. Abstraction is a device

that allow our minds to cope with literally limitless amounts of information. As Hayek puts it, “[A]bstract concepts are a means to cope with the complexity of the concrete which our mind is not fully capable of mastering” (1973, p. 29). We deal with that complexity by filtering it, deeming some features relevant while ignoring others. Hayek again: “We never act, and could never act, in full consideration of all the facts of a particular situation, but always by singling out as relevant only some aspects of it...” (1973, p. 30). The more we abstract, the more features we choose to ignore in order to focus attention on the remaining features deemed relevant. But there is a limit to the usefulness of the process, because at some point abstracting further means excluding some features that could be important for some purposes. Indeed, this is true even at the lowest level of abstraction, but the steepness of the trade-off increases as the level of abstraction rises. Statements at the highest level of abstraction will often have a tautological or vacuous character, as they no longer allow us to make useful distinctions.

Consider, for example, the use of abstraction in social-scientific modeling. Koppl and Whitman (2004), drawing on the ideal-type analysis of Alfred Schutz (1932), observe that models of human behavior can be placed on a spectrum of “anonymity.” A model with low anonymity makes very specific assumptions about the characteristics of the people modeled; the least anonymous model would attempt to predict the behavior of just one person – say, Napoleon (Koppl and Whitman 2004, p. 302). A model with very high anonymity makes very generic assumptions about people’s characteristics; the most anonymous model would only make assumptions literally true of every single human being. An example of a highly anonymous ideal type would be Ludwig von Mises’s model of a human being as someone who acts “to substitute a more satisfactory for a less

satisfactory state of affairs” (Mises 1949, p. 13; cited in Koppl and Whitman 2004, p. 304). The analogy to the level of abstraction in rules should be clear: highly specific ideal types are analogous to highly specific directives, highly anonymous ideal types to highly abstract directives. And it turns out that the advantages of rules with an intermediate degree of abstraction correspond to the advantages of ideal types with an intermediate degree of anonymity.

The use of ideal types allows us to make predictions about entire groups or classes of people, economizing on mental effort. But there is a precision-versus-accuracy trade-off involved in moving to higher levels of abstraction; as our models become more abstract, the kind of things we can say about any abstractly-defined group become less precise (Koppl and Whitman 2004, p. 305). In the extreme, we can abstract so much that we can say almost nothing at all. The Napoleon ideal type is not useful for predicting the behavior of anyone but Napoleon. The Misesian ideal type is not useful for predicting much behavior at all, except in the very broadest of terms, without filling in more specific details about the type of people whose behavior we wish to predict (e.g., that they prefer leisure to work, that they discount the future, etc.). The highly specific end of the spectrum is all precision, no accuracy; the highly abstract end is all accuracy, no precision. But these turn out to be very similar! If your point is to be able to say something useful, to make a prediction about behavior, then both have the same problem. The highly abstract end of the spectrum requires us to get more information to make a sufficiently precise statement. The highly specific end also requires us to get more information to determine which model to use – Napoleon or Idi Amin or George W. Bush? So the most useful social scientific models will be those in the intermediate range

(although more and less anonymous models will be useful for some purposes). By settling on models with intermediate levels of anonymity, we can be reasonably (though not perfectly) confident about which ideal type applies in understanding any given social situation.

Rules with intermediate abstraction provide an analogous benefit. Adopting rules that abstract from specifics and focus on a few relevant factors allows us to make predictions for entire groups and classes. For those classes, we can make reasonably confident predictions about “correct” behavior. But if the rules become too abstract, the classes become larger, and we become less confident about our ability to determine “correctness” for any given member of a class.

Why, if rules with an intermediate degree of abstraction have the advantages I have stated, are such rules occasionally displaced by either more abstract or more specific directives? This question is best answered by reference to examples from law, language, and etiquette. But first, we need some additional terminology.

1.3. Input Sets and Output Sets

As noted earlier, we can imagine a very abstract rule whose effects are nonetheless easy to predict. If every case within the entire class will be treated *exactly* the same way, then the breadth of the class poses no problem at all. It is easy to determine what class one’s case falls in, and it is easy to predict the outcome. In writing, an example of this kind of rule would be, “Always end a sentence with a period.” The rule would be clear and easy to apply. Of course, it would also rule out the use of question marks and exclamation points. But there would be little or no problem of

uncertainty. In tort law, an example would be, “The plaintiff shall always win and receive damages in the amount of \$1 million.” This rule might produce all sorts of problems, but it would be highly predictable in its application. So there is no *necessary* connection between high-level abstraction and uncertainty. I will argue, however, that there is a *systematic and empirical* connection.

Borrowing from Pierre Schlag, I will use the word “directive” broadly to refer to any mode of assessing appropriateness or determining outcomes, whether a command, rule, or standard. As Schlag puts it, “The formula for a legal directive is ‘if this, then that.’ A directive thus has two parts: a ‘trigger’ that identifies some phenomenon and a ‘response’ that requires or authorizes a legal consequence when that phenomenon is present” (1985, p. 381).⁸ The same is true of non-legal directives as well.

Any system of directives creates a *mapping* from situations to outcomes. The mapping divides the set of all cases into a number of *input sets*; these sets may be wide or narrow, as illustrated earlier in Figure 1. Thus far, I have implicitly defined the abstraction of rules solely in terms of the breadth of input sets. But the mapping also designates an *output set* corresponding to each input set. The output sets may also be wide or narrow. So a system of directives creates a mapping from input sets to output sets.

Figures 3A, 3B, and 3C illustrate the input sets and output sets of case-by-case, rule-based, and standard-based decisionmaking, respectively. In each figure, the system of directives is illustrated by arrows leading from situations (cases) to outcomes. Within the “situations” column, the divisions separate input sets, so having more divisions means

⁸ According to Schlag, the breakdown of a directive into these two components is conventional in the rules-standards literature (1985, p. 381, n11).


having more (and narrower) input sets. Within the “outcomes” column, the divisions separate different outcomes (for simplicity of presentation, I have assumed only three possible outcomes). The divisions do *not* distinguish output sets. For any given input set, the corresponding output set is *all* outcomes with arrows from that input set.

For a directive system with a low level of abstraction – that is, very narrow input sets as shown in Figure 3A – uncertainty is likely to result regardless of the size of the corresponding output sets. Even if each input set maps to a singleton output set (only one outcome, as in the figure), the decisionmaker has to determine which input set his case falls in. If other “nearby” input sets that differ only in minute details nevertheless map to different outcomes, then the decisionmaker cannot effectively predict the outcome.

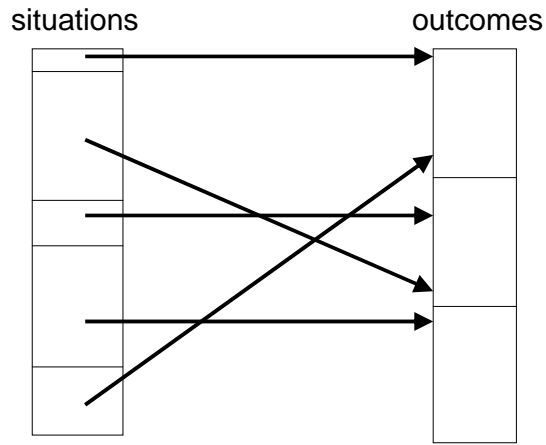
For a directive system with a high level of abstraction, the breadth of output sets matters. It is relatively easy to determine the relevant input set, but uncertainty can arise from a too-broad output set. In Figure 3C, there is only one input set, but it maps to all possible outcomes (a wide output set). We can imagine a version of Figure 3C in which there is just one arrow, leading from the only input set to one possible outcome. But in reality, that sort of directive is rare, and wide output sets are more common. For instance, the standards observed in law – such as balancing and reasonable man tests – are characterized by the vagueness of their recommendations. They specify relevant factors, not particular outcomes. In language, and specifically in writing, style guidelines like “be concise” and “avoid run-on sentences” are similarly vague in their specification of outcomes; they identify some relevant factors to consider in one’s speech and prose, but their application to particular cases is not apparent *a priori*. (For instance, was the previous sentence a run-on?)

Figure 3: Input Sets and Output Sets


3A: Case-by-Case



3B: Rules



3C: Standards



The directives we most often think of as rules have intermediate-sized input sets with relatively small output sets. In Figure 3B, there is only a handful of input sets, thereby minimizing uncertainty about where one's own situation falls; and each input set maps to a single outcome (a singleton output set), thereby minimizing uncertainty about the rule's application.

Although we can imagine directives with very wide input sets and very narrow output sets, such beasts are rarely found in the world. When input sets are large, output sets tend to be large as well. As we shall see in the next section, this fact results from the desire to achieve a desirable match between situations and outcomes.

2. Applications in Law, Language, and Etiquette

2.1. Rules versus standards in the law

The issues raised here have been most extensively considered in the legal literature on rules versus standards. Although the literature is large, legal scholars do not seem to have noted explicitly the importance of intermediate abstraction in rules. On the contrary, rules and standards are typically characterized as the poles of a spectrum. Duncan Kennedy's seminal article on the subject, for instance, says, "At the opposite pole from a formally realizable rule is a standard or principle or policy" (Kennedy 1976, p. 1688). Spencer Overton, in a recent application of the rules-versus-standards debate in the context of *Bush v. Gore*, says, "[T]he form that legal directives may take is actually relative: directives may fall anywhere on a graduated continuum between the extreme poles of rule-like and standard-like" (Overton 2002, p. 73). Vincy Fon and Francesco

Parisi, in developing a model of optimal specificity in the law, say, “Standards and rules can be visualized as two extremes in a one-dimensional space representing the degree of precision of laws” (2007, p. 148). The present analysis indicates that in fact, ruleness does not reside at one end of a spectrum, but somewhere in the center.

Nevertheless, the rules-standards literature does draw attention to the leading reason that directives mapping from very wide input sets to very narrow output sets do not persist – and the reason is closely related to the reason that intermediately-abstract rules sometimes get replaced by more or less abstract directives. While bright-line rules have the advantage of clarity, and thus the salutary effect of reducing uncertainty in a wide range of cases, they also create problems of over- and under-inclusion relative to other values. Kennedy uses age-of-majority laws as an example:

The choice of rules as the mode of intervention involves the sacrifice of precision in the achievement of the objectives lying behind the rules. Suppose that the reason for creating a class of persons who lack capacity is the belief that immature people lack the faculty of free will. Setting the age of majority at 21 years will incapacitate many but *not all* of those who lack this faculty. And it will incapacitate some who actually possess it. (Kennedy 1976, p. 1689)

Because rules are abstract, they will necessarily neglect a variety of details, and those details will sometimes appear important. The rules will therefore be perceived as creating injustice or unfairness in some cases. Although Kennedy is friendly toward the use of general standards, even those who strongly prefer rules recognize this point.

Richard Epstein identifies the demand for “perfect justice” as one of the “enemies of simplicity” in the law: “[Perfect justice] aspires to rooting out error in every individual case. Simple rules do not meet that exacting standard. At best they are only tests; and tests are rules of thumb that work most of the time, but are known and expected to fail some of the time” (Epstein 1995, p. 38).

If the desire to avoid one-size-fits-all (or all-in-the-class) outcomes is the reason for the rejection of rules that map from intermediate input sets to narrow output sets, then it should also be obvious why a mapping from very broad input sets to narrow output sets won't do: it will produce even more errors of over- and under-inclusion than would a mapping with intermediate input sets. I will refer to the desire to avoid such errors as the “matching principle”: ideally, we would like a system of directives to produce a perfect match between situations and outcomes, as determined by some external set of values.

Pursuit of the matching principle creates a tendency to slide toward either case-by-case decisionmaking or standard-based decisionmaking. In the rules-versus-standards literature, the slide is typically characterized as movement away from the pole of rules and toward the more abstract pole of standards. Yet the explanations given for the sliding, interestingly enough, often illustrate movement in the direction of *lower* abstraction. Carol M. Rose, who uses the metaphor of “crystals” and “mud” to mirror the dichotomy between rules and standards, draws attention to the tendency of strictly enforced property rules of the sort that prevailed in the 19th century to generate cases in which sympathetic plaintiffs suffered great losses as a result of seemingly small errors. “It is this booby trap aspect of what seems to be clear, simple rules – the scenario of disproportionate loss by some party – that seems to drive us to muddy up crystal rules with the exceptions and the post hoc discretionary judgments” (Rose 1988, p. 597). Rose describes a process by which bright-line rules are transformed into standards through the accumulation of exceptions. Yet the creation of specific exceptions seems a move in the direction of more *specific* rules, not more abstract ones. Eventually, the mass of exceptions “congeal” into standards that name all the relevant factors without

demarcating clear boundaries. Once again, we see how the ends of the abstraction continuum approach each other.

Kennedy, on the other hand, describes a process by which standards devolve into case-by-case decisionmaking:

Finally, the application of a standard to a particular fact situation will often generate a particular rule much narrower in scope than that standard. One characteristic mode of ordering a subject matter area including a vast number of possible situations is through the combination of a standard with an ever increasing group of particular rules of this kind. The generality of the standard means that there are no gaps: it is possible to find out something about how judges will dispose of cases that have not yet arisen. But no attempt is made to formulate a formally realizable general rule. Rather, *case law gradually fills in the area with rules so closely bound to particular facts that they have little or no precedential value.* (Kennedy 1976, p. 1690; emphasis added)

Kennedy's argument here again supports the thesis that case-by-case decisionmaking and standard-based decisionmaking have much in common; the ends of the spectrum approach each other. Under a standard, the need for some basis of decisionmaking will generate a tendency to create rules – but the rules may end up being highly specific instead of intermediately abstract.

Thus, the legal literature on rules-versus-standards has implicitly recognized the central thesis of this article, although the impulse to cast rules as one end of the spectrum has impeded making the thesis explicit. In Rose's terminology, bright-line rules are crystals and standards are mud; but then what is case-by-case decisionmaking? Take crystals and crush them to make smaller and smaller crystals, and eventually you get sand. As it turns out, mud and sand have much in common.

2.2. Regular versus irregular verbs in language

The issue of levels of abstraction in behavioral directives can also be found in linguistics, perhaps most obviously in the distinction between regular and irregular verbs – a topic that Steven Pinker uses as the leading example throughout his book *Words and Rules* (1999). Most verbs in the English language form their past tense and participle simply through the addition of *-ed*.⁹ Yet there exist 150 to 180 verbs in the English language that form their past tense and past participle in some other way: *go/went/gone*, *run/ran/run*, *sting/stung*, *bear/bore/borne*, and so on (Pinker 1999, p. 16); they are exceptions to the general rule. This phenomenon is not unique to English; other languages also have regular and irregular forms.

Note that the standard rule for forming the past tense possesses an intermediate level of abstraction; it applies only to verbs, and only when you wish to form the past tense or participle. For a wide input set, the rule specifies a singleton output set. Different rules apply to the inflection of nouns – and nouns, too, come in regular and irregular forms. But the irregular forms are loopholes, exceptions made for specific cases. We can tolerate these words that live by their own rules, but their presence in the language does create difficulties for people searching for the right word when writing and speaking – especially those learning the language, like children and non-native speakers. This is a problem of uncertainty, an inability of decisionmakers (speakers) to know what is appropriate behavior (speech) under particular circumstances. Too many irregular words would be intolerable. As Pinker says, “One can imagine a language in which every verb picked its own substitution of vowels and consonants from among the

⁹ This rule is not quite as simple as it seems, since the pronunciation of this suffix depends on the preceding sounds in the verbs to which it attaches, as illustrated by the different sounds at the end of *jumped* and *loaded*.

thousands that are logically possible. But generations of learners have passed down an English language that is very different from that possibility” (p. 91). This imaginary scenario would constitute a case of extreme specificity or too-low abstraction, and it is clearly untenable. Yet the opposite extreme would also be untenable: a language whose only rule for verbs was “form the past tense in a way that is easily pronounced and understood” would fail miserably at coordinating the expectations of speakers, and for essentially the same reasons.

Abstraction is justified in language for the same reason it is justified in general: it allows humans to manage complexity by focusing on some features while ignoring others. Pinker points out the gains in terms of economizing on mental effort and space: “The advantage of a rule is that a vast number of forms are generated by a compact mechanism. In English the savings are significant: The rules for *-ed*, *-s*, and *-ing* (the three regular forms of the verb) cut our mental storage needs to a quarter of what they would be if each form had to be stored separately” (p. 18).

The advantage of clear rules of intermediate abstraction is clear; so why have irregular words at all? One obvious (and true) explanation is history; most irregular forms are simply remnants of older rules of inflection (Pinker 1999, p. 67-68). They do, however, also have some advantages over regular forms, the most salient being ease of pronunciation: “A sequence of sounds that encodes a concept precisely and efficiently may be unresolvable by the ear or unpronounceable by the tongue” (p. 18). Regular inflection can produce words and phrases that are difficult to speak and to hear, thus sometimes requiring a modification. *Ice cream* is now the accepted term for what used to be *iced cream*; *ice tea* does not (yet) have the same acceptance, but both phrases illustrate

how the usual rules can be set aside to ease pronunciation. *Sixths* and *edited* are examples of difficult-to-pronounce words created by the regular inflection rules; “Monstrosities like these are never found among the irregulars, which all have standard Anglo-Saxon word sounds such as *grew* and *strode* and *clung*, which please the ear and roll off the tongue” (p. 19). That irregular forms are typically euphonious points to the role of evolution in selecting rules, and exceptions, to fit the needs of speakers.

As in the case of law, then, we can see that rules can create bothersome outcomes that apparently conflict with underlying values – whether fairness in the law or ease of pronunciation in language. As Frederick Schauer argues, any rule is likely to create errors of both over- and under-inclusion relative to the real or perceived functions of the system: “Just as the factual predicate [of a rule] may sometimes indicate the presence of the justification [for the rule] in cases in which it is absent, so too can the factual predicate occasionally fail to indicate the justification in cases in which it is present” (1991, p. 32). As a result, rules are necessarily suboptimal when measured against the yardstick of ideal decision-making (Schauer 1991, p. 100). Similarly, Ronald Heiner considers suboptimality a defining feature of rule-governed behavior: “If we use the jargon of standard economics, rule-governed behavior means that an agent must ignore actions which are actually preferred under certain conditions” (Heiner 1983, p. 568). The matching principle urges us to create exceptions or to allow greater discretion in order to reduce the divergence between our values and the output of rules. Seeking a perfect match in every case, however, can reduce clarity and, paradoxically, undermine the pursuit of those values.

2.3. Generalities versus minutiae in etiquette

The issue of levels of abstraction arises in etiquette as well, though the literature on etiquette is not as well developed as that of law and linguistics. Two popular commentators on etiquette, Amy Vanderbilt and Judith Martin (a.k.a. “Miss Manners”) both affirm the predictability-generating character of the rules of etiquette. Vanderbilt says that “The rules of etiquette give us something to lean on in a world that grows increasingly large and complex” (1968, p. 98). Notably, Vanderbilt does *not* call for the rules to become more complex as society grows more complex. The rules are valuable because they help manage complexity, not because they reflect it. Similarly, Martin says, “In its symbolic function, etiquette provides a system of symbols whose semantic content provides for predictability in social relations, especially among strangers” (1993, p. 354).

Both authors recognize the hostility that some people have toward etiquette. Vanderbilt attributes the hostility to their having effectively given up; in essence, they say, “I don’t know anything about it. I’m not going to try. I might make a mistake” (1968, p. 98). Presumably, the fear of making a mistake derives from a sense that etiquette is dominated by *minutiae* – that is, directives that apply to very specific situations, such as differences in how to hold a tea cup versus a demitasse. Directives like these err on the side of too much specificity, and thus leave people uncertain about whether their actions are appropriate. The alternative to the etiquette of minutiae might seem to be a general statement of purpose, such as the aforementioned “Be polite.” But this, too, would be problematic. The notion that people should simply be nice to each other leaves us, as Martin puts it, with “the Jean-Jacques Rousseau School of Etiquette, with its charmingly naïve directive that we should all behave like (noble) savages” (1993,

p. 350). Whether presented with directives that are too specific or too general, people still don't know how to behave.

The more useful rules of etiquette, through their intermediate abstraction, serve to divide up social interactions into broad classes. Martin explains:

Once learned and correctly interpreted, the symbols of etiquette permit one person to recognize such essential attributes of other persons as their intentions, status, friendliness or hostility, and thus to deal appropriately with a wide range of social situations and relationships. (1993, p. 354)

Notice that the “symbols” of etiquette simultaneously provide useful information while nevertheless allowing the user to deal with a broad range of circumstances; this is the key feature of intermediate abstraction. Martin gives the instructive example of someone reaching out to shake your hand, which indicates that this person will likely treat you better than someone who spits at your feet (1993, p. 354).

What, then, drives a system of etiquette away from intermediate abstraction? Again, the answer is the matching principle: a fixed rule will generate outcomes that do not perfectly match the goals of people using the system. Martin says the functions of etiquette include “communal harmony, dignity of the person, a need for cultural coherence, and an aesthetic sense” (1993, p. 351). Perhaps more importantly, etiquette can function to distinguish insiders from outsiders. A secret handshake allows people within a select group to recognize each other. This is a special case of the broader function of distinguishing classes of people who will treat you differently. But because rules are both under- and over-inclusive, the symbols of an etiquette system may fail to send the precise signals that people want to send. A handshake, for instance, only distinguishes the tolerant from the actively hostile. More complicated (and specific) practices can develop to show gradations of friendship. In some social situations, it is

now necessary to choose whether to perform a standard handshake, three consecutive hand-clasps, or a fist-bump. The expansion of options allows a more finely tuned message, but it also can create uncertainty and discoordination – such as when someone meets an attempted hug with an attempted handshake. While etiquette generally is supposed “to soften personal antagonisms and thus to avert conflicts” (Martin 1993, p. 352), crossed signals like these can actually foment conflict.

As with law and language, we see that the desire to achieve a perfect match between goals and actions creates a tendency to make exceptions to rules, or to replace them with more open-ended standards. But if too many exceptions accumulate, or if the standards become too open-ended, the resulting uncertainty can undermine the achievement of those same goals.¹⁰

3. Other Rules Conducive to Spontaneous Order

Are rules of intermediate abstraction a *necessary* condition for the emergence of a spontaneous order? Apparently not, given the existence of both standards and case-by-case decisionmaking in the law, and the existence of irregular forms in language. Somehow we are still able to coordinate our economic behavior and to communicate effectively. However, more and less abstract modes of decisionmaking like these seem to exist in the context of an overall system that is mostly populated by rules of intermediate abstraction. As long as such a background exists, it seems that a spontaneous order can

¹⁰ See Heiner (1983, p. 564): “[A]llowing flexibility to react to information or to select [more] actions will not necessarily improve performance if there is uncertainty about how to use that information or about when to select particular actions. Thus, an agent’s overall performance may actually be improved by restricting flexibility to use information or to choose particular actions.”

tolerate some degree of exception-making (in the form of more specific directives) and vagueness (in the form of broader standards).

Intermediate abstraction of rules also does not guarantee the emergence of a spontaneous order. As Hayek repeatedly observes, not just any kind of abstract rules will do.

[I]t is evident that in society some perfectly regular behaviour of the individuals could only produce disorder: if the rule were that any individual should try to kill any other he encountered, or flee as soon as he saw another, the result would clearly be the complete impossibility of an order in which the activities of the individuals were based on collaboration with others. (1973, p. 44)

To understand spontaneous order, we need to identify the other characteristics that abstract rules must have. “The question which is of central importance as much for social theory as for social policy,” says Hayek, “is thus what properties the rules must possess so that the separate actions of the individuals will produce an overall order” (1973, p. 45). In this section, I will attempt a partial listing of those properties.

3.1. Purpose-independence

The characteristic most clearly identified by Hayek – other than abstraction – is that rules need to be *purpose-independent*, meaning they apply without reference to the specific goals of specific people. The rule governing a situation does not change depending on who is doing an activity or why he is doing it; as Hayek puts it, “[T]he rules governing a spontaneous order must be independent of purpose and be the same, if not necessarily for all members, at least for whole classes of members not individually designated by name” (1973, p. 50). Purpose-independence is important for two reasons.

First, purpose-independence allows the order to persist even when a particular goal has been accomplished or is no longer relevant. Because the rules can accommodate a wide (though not unlimited) range of goals, new members of the community will find it in their interest to continue using them to coordinate behavior. A language persists, for instance, because people always have communicative needs – but the needs can change dramatically from person to person and from time to time. The abstract rules of language can assist in any activity from the writing of a contract to the planning of a terrorist attack. Similarly, the market economy persists in large part because it responds to the ever-changing needs of economic actors. The rules of just conduct that foster the market process need not make reference to specific resources, products, services, and technologies; this is one of the advantages of (intermediate) abstraction. Instead, the rules generally refer to broad classes and types of situations, from transfer of control over resources (movable and immovable) to accidents between strangers. Because these classes of interaction occur regardless of the particular people and goals involved, the system provides a means that can serve innumerable different ends.

Second, purpose-independence allows for the continual incorporation of new knowledge by allowing individuals to respond to a changing environment. A changing environment guarantees that *no* system can satisfy all expectations; it is therefore a given that some expectations will be dashed. “The protection against disappointment of expectation which the law can give in an ever changing society will always be only the protection of some expectations but not all” (Hayek 1973, p. 102). If the system aimed, in the face of an ever-changing environment, at the satisfaction of all expectations of specific individuals, it would necessarily have to obstruct the discovery and use of new

ideas and information by other individuals. The inability to respond to change would lead to ever greater disappointment of expectations of other people. Purpose-independence substitutes satisfaction of certain kinds of expectations across persons for the satisfaction of all expectations of a few.

3.2. Extensibility to novel cases

A closely related characteristic of rules of intermediate abstraction is their *extensibility to novel cases* – that is, their capacity to be extended to cover new and unforeseen activities and interactions. This provides a degree of certainty to people in new and unfamiliar situations. As long as those situations possess some of the features identified as relevant by existing rules, it is possible to make some predictions about how the situations will be treated.

In language, the existence of regular rules of inflection facilitates the entry of new words into the language, often in the presence of new activities people wish to communicate about. As Pinker observes, when the new verbs *fax*, *mosh*, and *spam* entered the lexicon, there was no doubt about how to form their past tenses: *faxed*, *moshed*, and *spammed* (1999, p. 13). The same is true when a word is repurposed into a different part of speech; if we decided to *verb* a noun, we would know the noun had been *verbed*. Interestingly, even irregular forms can exert some control over the formation of new words. When words are formed by adding a prefix to an existing irregular word, such as *understand* from *stand*, we know that the past tense will be formed in the same manner as the root word: *understood*, not *understanded*. If we invented new words

based on irregular stems, such as *withset* or *overhold*, we would automatically know their past tenses as well (*withset* and *overheld*) (Pinker 1999, p. 45).

The fact that even irregular verbs can plant the seeds of patterns illustrates two related points. First, when exceptions are made to rules, often the exceptions themselves will take the form of rules. As long as there are not too many exceptions, the intermediate abstraction of the system of rules can be maintained and its beneficial properties – like extensibility to novel cases – preserved. Second, logical consistency is *not* a necessary component of a system of rules. As Mario Rizzo (1999) argues in the context of legal rules, *logical consistency* and *praxeological consistency* are not the same thing. Praxeological consistency refers to whether the rules allow individuals to coordinate their actions with others by making mutually compatible plans (Rizzo 1999, p. 502). If two classes A and B are treated by two different rules, even though logically A and B would seem to be the same, individuals can still rely on those rules to guide their behavior so long as they can tell whether a case falls in A or B. The distinction between the two classes need not be logically relevant, so long as it is visible. For instance, if English units of measure are used for most everyday purposes, but metric units are used for medical and scientific purposes, the system is not logically coherent. But if people can nevertheless coordinate their measuring activities (because they can easily distinguish everyday from medical/scientific contexts), then the system is praxeologically coherent. The system can tolerate at least some degree of logical inconsistency.

3.3. Tolerance for small deviations

Any system of abstract rules will produce occasional failures in satisfying the customary functions of the order, such as communicating ideas or coordinating economic behavior. This fact creates an endogenous impulse to make exceptions or appeal to vague standards – the matching principle at work. To resist extreme applications of the matching principle, a system of rules must have a means of accommodating it to some degree.

Commonly, the process involves turning exceptions into general-purpose rules. In the field of etiquette, an example is provided by queuing norms. In a study of ticket queues for football matches in Australia, Leon Mann (1969) observed that queue members followed a general rule of first-come-first-served; however, the rule had exceptions that lessened the burden of waiting by allowing people to leave the queue from time to time – to eat, use the bathroom, and even play games. The two most common exceptions were (a) that one person could hold up to four spots for people in his group, and (b) that an object (such as a cardboard box) could hold a spot for a single person for two to three hours (Mann 1969, p. 344-5). Notably, these exceptions took the form of rules with relatively well-defined boundaries. They allowed better satisfaction of the matching principle without creating great uncertainty (although some uncertainty remained, as we will see below).

The same process is evident in the common law, where judges often refuse to make exceptions unless they can find some principle that will tell them when to make similar exceptions in the future. Of course, there is some danger in this approach. If too many exceptions accumulate, the list of relevant factors that distinguish them can morph

into a vague standard such a balancing test, which asks judges simply to weigh all the relevant factors in specific cases. Alternatively, a proliferation of narrow exceptions can move the system toward the pole of extreme specificity. Either way, the intermediate abstraction of rules gets eroded. Fortunately, as observed earlier, a spontaneous order can tolerate some degree of such erosion in the rules.

3.4. Openness to generation of new rules

A dynamic system governed by abstract rules will continually generate situations that expose vagueness in the rules (even when the rules appear very clear) or conflict among different rules (even when no conflict was apparent initially). Hayek, in explaining the role of common-law judges, observes that the exposure of residual vagueness may require the creation of new rules to supplement the old:

Since new situations in which the established rules are not adequate will constantly arise, the task of preventing conflict and enhancing the compatibility of actions by appropriately delimiting the range of permitted actions is of necessity a never-ending one, requiring not only the application of already established rules but also the formulation of new rules necessary for the preservation of the order of actions. (Hayek 1973, p. 119)

The emergence of cases that expose conflict or vagueness may well result from the intermediate degree of abstraction. Sufficiently specific directives are likely to have “cracks” between them, because they are not abstract enough to cover every new kind of case. Sufficiently abstract directives, on the other hands, are unlikely to have cracks because they are broad enough to encompass everything – but their corresponding vagueness in terms of output (that is, the breadth of their output sets) will create the need for more specific rules. Intermediately abstract directives will have more cracks than highly abstract ones, but fewer than highly specific ones.

In language, the existence of conflicting rules can put words “in play” with regard to their proper inflection. Consider, for example, the verb *know/knew*. The formation of the past tense by means of this specific vowel change applies to a whole class of irregular verbs, including *grow/grew* and *throw/threw*, all of which involve stems that end in a vowel sound and start with a consonant cluster. How did *know/knew* end up in this group, even though it begins with a single consonant sound? Originally, the *k* was pronounced, but then English began to lose the pronunciation of *kn* and *gn* at the beginning of words. So the emergence of a new rule (no more *kn* and *gn* consonant clusters) created conflict with an old one (inflection for this group of words). As it turned out, the old rule won out, with *know* keeping its old past tense *knew*. (Pinker 1999, p. 71-72.)

In etiquette, Mann’s football queues provide a nice example of how vagueness gets exposed. The rule that an object can hold a spot has an uncertain time limitation. In one football queue in 1966, “irate latecomers, who noticed that many people in the middle of the queue had not made an appearance for most of the day, spontaneously seized their boxes and burnt them,” and the result was a melee (Mann 1969, p. 346). This is a notable example of how uncertainty can undermine the basic function of the system – in this case, avoiding interpersonal conflict. To avoid such a conflict, it would be necessary to create new rules or clarify the old ones.

One commonly followed rule of etiquette says, “Always hold a door open for a woman”; another common rule says “Always allow a woman to enter first.” These two rules generate a conflict in cases where a door opens *into* the place to be entered, because there is a leverage problem in trying to hold open such a door without actually walking

in. In practice, a different rule had to evolve to govern this case (“enter just enough to open the door”). But what is to be done in the case of a revolving door? A conflict emerges here as well; one writer on etiquette suggests that “man first” governs the revolving-door situation (Post, n.d.), but custom may still be evolving on this point. (The emergence of a clear rule in this case may have been slowed by weakening of the old norm of deference toward the “weaker sex.”) For an example of where etiquette had to develop a new rule for a novel situation, consider the emergence of “ATM distance,” the appropriate distance to stand behind someone at an automatic teller machine. Prior rules of social distance addressed distance for regular conversation, flirting, standing in a subway, etc., but not for standing in line at an ATM.

The need to generate new rules for cases of vagueness and conflict among rules highlights a significant role that high-level abstraction can play in the maintenance of a spontaneous order. Although I have argued throughout this article that a system cannot be based entirely on directives at very high levels of abstraction, such directives can be very helpful in *supplementing* rules of intermediate abstraction. The intermediately abstract rules cannot, by their nature, provide the solution or answer to their own contradictions and omissions. Rules do not explain themselves; they do not include their own justifications. Standards, however, can provide a list of the relevant factors that could enter into the creation of new rules or resolution of conflicts among existing rules. This matters for two reasons. First, it means standards are useful in allowing the system of rules to grow and change – as it must to accommodate a changing environment. And second, standards can lend a degree of certainty when agents have to make decisions in situations where the rules are unclear, even *before* the new rules congeal. The certainty

will not be as great as when the rules for one's case already exist, but it will be certainly be greater than when there is neither a rule nor a standard to rely on.

4. Conclusions: On Structure and Content

The rules that encourage the growth of a spontaneous order will, taken as a whole, tend to display an intermediate degree of abstraction. A system governed by rules that are too specific or too abstract will fail to coordinate expectations of people working within the order that emerges, whereas rules of intermediate abstraction economize on mental space while minimizing vagueness. Intermediate abstraction is not enough, however; the rule set also needs to be purpose-independent, extensible to novel cases, robust to small deviations, and open to emergence of new rules.

It is useful to distinguish between the rules that create the conditions for a spontaneous order and the spontaneous order itself. A market economy is a spontaneous order that emerges from certain regularities of human economic behavior, some of which result from the rules of just conduct embodied in the law. That does not mean the law itself is a spontaneous order, although it might be. Hayek points out that “[I]t is possible that an order which would still have to be described as spontaneous rests on rules which are entirely the result of deliberate design” (1973, p. 46). The rules of language set the stage for a spontaneous order of human communication; in principle, the rules of the language could have been designed, like Esperanto or Klingon. The rules of hypertext markup language (HTML) provide a partial basis for the spontaneous order of the Internet, yet HTML itself was deliberately designed.

Nevertheless, it turns out in practice that the rules that facilitate spontaneous order are often themselves the product of spontaneous order (or some process of evolution without design). The rules of law, language, and etiquette emerged spontaneously, albeit with numerous attempts at central direction by governments, grammarians, and guardians of manners. The discussion in Part III of the properties that abstract rules need to satisfy helps to explain why. In order for the system of rules to be extensible to novel cases, robust to small deviations, and open to creation of new rules, the rules themselves must be capable of some amount of evolution in response to changing needs and new circumstances.

Imagine a designer of rule systems who wishes to create a spontaneous order. He might start by creating rules characterized by an intermediate degree of abstraction, and then set the system in motion. But there is no way the designer can foresee all possible novel cases, conflicts, and changes in circumstances *a priori*, so the rule set will assuredly have to change. So the designer faces a choice: either he can maintain constant watch over the system and make every rule change deliberately, or he can find a way to allow the system to evolve on its own. The latter will, of course, require less effort on the designer's part; moreover, the changes that occur in the rules will have to result from changes that occur in the system. In short, an ongoing spontaneous order seems to require an ongoing evolution of the rules that help generate it. To the extent that the evolutionary process responds to the human need for predictability, we should expect the resulting rules to possess an intermediate degree of abstraction.

References

- Epstein, Richard A. 1995. *Simple Rules for a Complex World*. Cambridge: Harvard University Press.
- Fon, Vincy and Parisi, Francesco. 2007. "On the Optimal Specificity of Legal Rules." 3(2) *Journal of Institutional Economics* 147-164.
- Hayek, Friedrich A. 1960. *The Constitution of Liberty*. Chicago: The University of Chicago Press.
- Hayek, Friedrich A. 1973. *Law, Legislation and Liberty, Vol. 1: Rules and Order*. Chicago: The University of Chicago Press.
- Heiner, Ronald A. 1983. "The Origin of Predictable Behavior." 73 *American Economic Review* 560-595.
- Kennedy, Duncan. 1976. "Form and Substance in Private Law Adjudication." 89 *Harvard Law Review* 1685-1778.
- Koppl, Roger. 2002. "Custom and Rules." 61 *American Journal of Economics and Sociology* 531-537.
- Koppl, Roger and Whitman, Douglas G. 2004. "Rational-Choice Hermeneutics." 55 *Journal of Economic Behavior & Organization* 295-317.
- Mann, Leon. 1969. "Queue Culture: The Waiting Line as a Social System." 75 *American Journal of Sociology* 340-354.
- Martin, Judith. 1993. "A Philosophy of Etiquette." 137 *Proceedings of the American Philosophical Society* 350-356.
- O'Driscoll, Gerald P., and Rizzo, Mario J. 1996 (1985). *The Economics of Time and Ignorance*. London: Routledge.
- Overton, Spencer. 2002. "Rules, Standards, and *Bush v. Gore*: Form and the Law of Democracy." 37 *Harvard Civil Rights-Civil Liberties Law Review* 65-102.
- Pinker, Steven. 1999. *Words and Rules*. New York: Basic Books.
- Post, Peggy. Not dated. "Revolving Doors." GoodHousekeeping.com, accessed Feb. 17, 2008, at http://magazines.ivillage.com/goodhousekeeping/etiquette/peggy/qas/0,,284571_431459,00.html.

Rizzo, Mario J. 1999. "Which Kind of Legal Order? Logical Coherence and Praxeological Coherence." 9(4) *Journal des Economistes et des Etudes Humaines* 497-510.

Rose, Carol M. 1988. "Crystals and Mud in Property Law." 40 *Stanford Law Review* 577-610.

Schauer, Frederick. 1991. *Playing by the Rules: A Philosophical Examination of Rule-Based Decision-Making in Law and in Life*. Oxford: Clarendon Press.

Schlag, Pierre. 1985. "Rules and Standards." 33 *UCLA Law Review* 379-430.

Schutz, Alfred. 1932. *The Phenomenology of the Social World*. Translation by G. Walsh and F. Lehnert. Evanston, Ill.: Northwestern University Press.

Vanderbilt, Amy. 1968. "Bad Manners in America." 378 *Annals of the American Academy of Political and Social Science: The Changing American People: Are We Deteriorating or Improving?*, 90-98.