

Three Strategies of Interpretation

Adrian Vermeule*

We may distinguish three styles or strategies of decisionmaking. Under a maximizing approach, the decisionmaker chooses the action whose consequences are best for the case at hand (defining “best” according to some value the decisionmaker holds). Where decisionmakers choose the action that is best relative to constraints, accounting for the direct costs and opportunity costs of decisionmaking, we may call the approach optimizing rather than maximizing. Whereas the maximizer focuses only on the case at hand, the optimizer acts so as to maximize value over an array of cases. In contrast to both approaches, satisficing permits any decision whose results in the case at hand are good enough—although we will see that satisficing, like optimizing, may itself represent an indirect strategy of maximization.

In what follows, I will suggest that these distinctions illuminate legal interpretation. Interpretation is just another type of decisionmaking, so interpreters must use some decision-procedure or other. Many approaches to the interpretation of statutes and the Constitution are maximizing approaches that attempt to produce as much as possible of some value the interpreter holds—for example, fidelity to legislative intent or original understandings. Optimizing approaches to interpretation condemn maximizing interpretation as a simpleminded approach that neglects the costs of decisionmaking and the costs of interpretive error. An alternative to both maximizing and optimizing approaches is a satisficing style of interpretation, in which interpreters eschew the search for the very best interpretation (even within constraints), instead selecting an interpretation that is good enough, in light of whatever value theory the interpreter holds. The choice among decisionmaking strategies is utterly agnostic about the underlying value theory; whatever such theory the interpreter holds, there is always a separate question about which decision-procedures are best suited to promote the decisionmaker’s aims.

I will criticize the maximizing style of interpretation and praise its two competitors. Both the optimizing and satisficing perspectives, I suggest, help in different ways to justify some controversial principles of statutory and constitutional interpretation, such as the rule barring resort to legislative history where statutes have a plain meaning, and clause-bound (as opposed to broadly holistic or “intratextualist”) interpretation of statutes and the Constitution. Although maximizing interpretation is untenable, neither the optimizing approach nor the satisficing approach is globally best; each is an attractive decision-procedure in some contexts. Where the interpretive stakes are either very low or very high, satisficing is reasonable (whether or not rational in some stronger sense), while optimizing is best suited to medium-stakes decisions.

I begin, in Section I, by clarifying the conceptual distinctions between and among maximizing, optimizing and satisficing. Section II identifies interpretive styles that draw upon these decisionmaking strategies, criticizes some prominent examples of maximizing interpretation, and examines some contested interpretive principles that the optimizing and satisficing perspectives might justify. Section III offers some considerations that bear upon the choice between optimizing

*The author wishes to thank Eric Posner, Fred Schauer, and Cass Sunstein for comments, and Justin Ruben for research assistance. This version is published without footnotes.

and satisficing as interpretive strategies. Section IV considers the wide appeal of maximizing styles of interpretation, and suggests some mechanisms that cause maximizing interpretation to appear

more attractive than it should. This is not an occasion for ponderous work, so my aim throughout is simply to sketch ideas; I make no effort at a comprehensive and detailed treatment.

I. Maximizing, optimizing and satisficing

The standard model of rational decisionmaking defines rational choice as the choice that maximizes some value. In a common interpretation, the value to be maximized is welfare, in turn defined (contentiously) as the satisfaction of subjective preferences. Nothing inherent in the model, however, requires this; the idea of rational choice deployed in decision theory is strictly formal. The decisionmaker simply ranks the outcomes of possible actions according to some scale of value and chooses the maximizing action—the action that produces the highest value. The model can be extended to cover situations of risk, in which the outcomes of actions are probabilistic rather than certain, by taking the maximand to be expected value. In situations of uncertainty, where not even the probabilities of various outcomes are known, one view suggests that decisionmakers can simply assign subjective probabilities, converting uncertainty back into risk; other views propose other choice criteria, such as maximin (maximizing the minimum payoff). The distinctions between certainty, risk and uncertainty are orthogonal to the issue I address here; the important point is that in all of these situations, the standard model defines rational choice as maximizing choice.

In the simplest versions of the standard model, the feasible set of actions or options is just given. Here the idea of satisficing is incoherent. How could it be rational, in a static context, to choose anything other than the best available action—to do anything other than maximizing? Some philosophers suggest that it can be rational to choose less than the best, so long as the action chosen is satisfactory. But if the action chosen is less than the best and also satisfactory, than the best action is also satisfactory, and satisficing gives no reason to choose the former over the latter. In static contexts, the superiority of maximizing to satisficing is conceptually entailed by the scale of value the decisionmaker uses.

Satisficing comes into its own, however, when decisionmaking is viewed more dynamically. In many real-world decisions, the set of options is itself (at least partially) the product of earlier decisions. One of the most important questions decisionmakers face is the extent of rational search: how many options, and how much information, should be sought out and considered before an ultimate choice is made? Here satisficing is coherent; as Herbert Simon emphasized, satisficing is a constraint on further search for new information and new options. The satisficer searches only until finding a choice whose outcomes are good enough. By satisficing with respect to the particular decision at hand, the satisficer conserves time and other resources that may then be expended on other decisions.

In dynamic settings, optimizing is also best understood as a constraint on further search for new options or new information. Maximizing always takes place within budget constraints, especially time. To maximize in the simpleminded sense of searching until one finds the very best option in the case at hand, all things considered, is to neglect the opportunity costs of search. (Although simpleminded maximizing in the case at hand is indefensible, I will also claim below that it is a

common approach to interpretive decisionmaking). The antonym of satisficing, on this view, is not simpleminded maximizing. Rather it is optimizing, by which we denote maximization that takes into account constraints, such as the cost of searching for further information and options. The optimizer searches until the costs of further search equal or exceed the expected benefits of additional information, or of new options.

In what follows, then, I shall contrast three different decisionmaking strategies: maximizing, taken to mean a simpleminded effort to find the very best choice, all things considered, in the particular decisionmaking context at hand; optimizing, or maximization that takes into account the direct costs and opportunity costs of acquiring information and making decisions; and satisficing. Optimizing and satisficing are different ways of pursuing the same larger aim. Both strategies rest on an implicit recognition that to do what is best, all things considered, with respect to some particular decision in an array of decisions is to do something that may not be globally best, or best from some larger perspective. Both strategies, in other words, are second-order maximizing: to maximize globally, the decisionmaker may do best to choose in a way that is less than maximally best with respect to the local decision at hand.

It is important to be clear, however, that optimizing and satisficing are different second-order decision strategies. The two strategies employ different stopping rules, or rules for constraining further search among possible options. The optimizer stops searching when the marginal benefit of finding a better option, discounted by the probability of finding such an option, is equal to or less than the costs of further search. The satisficer stops searching when she finds an option that is good enough. Although the two strategies sometimes yield similar choices, sometimes they do not, and even if they were extensionally equivalent, the two strategies would still embody different rules for choosing.

As the reference to stopping rules suggests, the jurisprudential distinction between rules and standards is relevant here. Maximizing is the ultimate standard: the maximizer does what is best, all things considered, taking into account the totality of the circumstances relevant to the local decision at hand. Optimizing and satisficing strategies both appeal to the higher-order virtues of rules, to the idea that a decisionmaker who takes into account less than the full set of considerations that bear on a particular decision may for a range of reasons do better, over a whole array of decisions, than does the simpleminded maximizer. Yet this conceptual point about the possible virtues of rules does not purport to specify the content of the rules. Optimizing and satisficing strategies are different stopping rules that use different means to their common aim of global maximization. “An optimizing strategy places limits on how much we are willing to invest in seeking alternatives. A satisficing strategy places limits on how much we insist on finding before we quit that search and turn our attention to other matters.”

Here is an illustration of the three strategies. Three decision theorists—M, O, and S—enter the main university cafeteria, which holds no less than twenty separate stations, each of which offers a different type of cuisine. M is a simpleminded maximizer who seeks the most satisfying possible meal right now; M spends the next hour visiting each station, pondering possible choices, and so on. (By the time M is done choosing his meal, O and S have finished eating and are back in their offices working on papers). O is a second-order maximizer, who sees that maximizing her satisfaction from this particular meal is suboptimal from a overall perspective. O thus adopts a

stopping rule that is calculated to optimize her satisfaction, taking into account decision costs and opportunity costs. Calculating the marginal costs and benefits, O decides to visit five randomly-selected stations out of the possible twenty, and then to choose from within this set the station whose offering maximizes the satisfaction of O's tastes. S is also a second-order maximizer, but she employs a different stopping rule: S proceeds along the stations until he finds an offering that is good enough, and then stops searching. Although O and S may happen to light upon the same offering, there is no guarantee that they will do so; and although it is true that both O and S manage to avoid the plight of the obsessive and self-defeating M, they have used different strategies to that end.

Of course no maximizer really considers all things. M will eventually choose a meal, rather than spending an infinite amount of time evaluating micro-features of the alternatives. Yet M may spend far more time and effort on this local choice than would be justified from a second-order, globally maximizing perspective. Although the local maximizer will stop at some point, that point may lie far beyond the local optimum, as identified from the global point of view.

II. Decisionmaking and interpretation

What does all this have to do with interpretation? Many debates over interpretive practices are debates over the decision procedures interpreters should use. These debates include important questions about how much information interpreters should collect, what set of possible interpretations to consider, and what stopping rules they should use as constraints on further search for information. Examining the history and theory of statutory and constitutional interpretation in America, we can identify interpretive styles that correspond to, and implicitly draw upon, the decisionmaking strategies we have identified.

Rather than pursue this theme in the abstract, I will proceed by demonstration, examining just two of the many settings in which the contrast among these interpretive strategies shows plainly. The first involves the debate over legislative history and the plain meaning rule; the second involves the debate, in both statutory and constitutional arenas, over the weight to be accorded to statutory or constitutional clauses other than the clauses directly at issue—including, in the statutory setting, the question how much attention courts should pay to statutes *in pari materia*.

Legislative history and the plain meaning rule

Suppose interpreters are intentionalists: they subscribe to some high-level political theory, perhaps an account of representative democracy, according to which legislators' intentions make the law. Intentionalism thus supplies the value theory that defines what counts as a good or bad interpretation: a good interpretation is one that captures legislators' intentions. Suppose also that this scale of value is continuous: interpretations may capture more or less of the legislators' true intentions, and the more the better.

Maximizing. In light of this value theory, the intentionalist interpreter who is also a simpleminded maximizer will proceed to search as widely as possible for information about legislators' intentions in the case at hand. On this view, interpreters may begin by consulting statutory text for evidence of legislative intent. But there is nothing special about text; interpreters will range beyond the statute

to consult legislative history, previous and subsequent statutes, perhaps even in-court testimony from the statute's drafters. This expansive search for further information about legislators' intentions follows from the assumption that "there is no invariable rule for the discovery of legislative intention"; any source is in principle admissible, and should be given whatever probative weight it intrinsically deserves.

Despite its intuitive appeal, simplistically maximizing intentionalism of this sort is exposed to serious objections. It neglects the direct costs and opportunity costs of searching further and further afield for evidence of legislative intentions. We may sort these costs into two rough categories, decision costs and error costs. Holding constant the accuracy of decisions, simplistically maximizing intentionalism produces wasted effort (from an ex post standpoint) whenever the interpretation the court ultimately settles upon, after extensive review of collateral sources, is the same interpretation it would have reached with a more restricted set of sources. Holding the costs of decision constant, searching further and further afield might even reduce accuracy even in the decision at hand, if maximizers with constrained cognitive and information-processing capacities become bewildered by a large set of conflicting evidence. Furthermore, decision costs and accuracy interact over the whole array of cases. Even if collecting more and more evidence of legislative intentions in Case 1 increases accuracy in that case, the opportunity cost of search means that the intentionalist interpreter will have less time to spend on Case 2 than she would if decisionmaking resources were distributed more evenly over cases—which means that the interpreter will tend to perform less accurately in Case 2.

Optimizing. An important alternative to simplistically maximizing intentionalism, then, is optimizing intentionalism, which constrains the search for evidence of legislators' intentions by reference to a larger cost-benefit calculus. The optimizing intentionalist employs a stopping rule: she declines to search further afield if the expected benefits of further search are less than the costs, including among the latter both the decision costs of search and the chance that new sources will reduce accuracy by driving a fallible interpreter off a correct interpretation that a smaller set of sources would have suggested. This description of the optimizing intentionalist calculus is abstract. We have sketched the variables the optimizing intentionalist takes into account, but the specific decision-rules that result depend on what values those variables actually take—an empirical issue. My point here is just that, given certain values of the relevant variables, the notion of an optimizing-intentionalist stopping rule provides a justification for considering less than all probative information bearing on legislative intentions in particular cases—even for interpreters fully committed to intentionalism as a high-level account of statutes' authority.

For a concrete example, we might consider something like the version of intentionalism embodied in the "plain meaning rule." Under that rule, legislative intent is the ultimate object of interpretation, but clear statutory text is conclusive evidence of legislative intent. The plain meaning rule is a stopping rule: the interpreter stops searching for further evidence of legislative intentions when the statutory text is clear. Further search, say into legislative history, is permissible only when the statutory text is ambiguous or otherwise lacks a plain meaning.

Thus in *Caminetti v. United States* the Court affirmed a conviction under the White Slave Traffic Act of a man who had transported a woman across state lines to become his mistress, reasoning that the statute's prohibition of interstate transportation for any "immoral purpose" was so plain as to

obviate the need for any recourse to legislative history—even though that history, according to the dissenters, showed that the intent of the prohibition was to criminalize only prostitution or other commercialized immorality. Although the Court conceded that legislative history “may aid the courts in reaching the true meaning of the legislature in cases of doubtful interpretation,” it held that “the language being plain, it is the sole evidence of the ultimate legislative intent.”

Nothing in the optimizing-intentionalist view of Caminetti, and the plain meaning rule, excludes or needs to exclude the possibility that the dissenters in Caminetti were correct. Perhaps the plain meaning did not actually track legislators’ intentions, and the legislative history would have revealed them. From the second-order standpoint of the optimizing decisionmaker, this is just to repeat the point that the plain meaning rule tolerates results that are suboptimal or erroneous in the decision at hand, all things considered, for the sake of better results over an array of interpretive decisions.

This tradeoff can, under certain circumstances, be beneficial on net. Suppose that plain statutory text is usually excellent evidence of legislative intent, and that going beyond statutory text into voluminous and complex legislative history often produces high decision costs and opportunity costs. Given finite time and decisionmaking capacity, and an array of future cases that must be decided, the interpreter may do better overall by allocating less effort to discerning legislative intentions in the particular case at hand, while allocating more effort to other cases in the array. Moreover, simpleminded maximizing past the optimum point may produce little improvement in accuracy even in the very case at hand. It is quite possible that simpleminded intentionalist maximizing in a given case will reduce accuracy, even in that very case, if cognitively fallible interpreters are confused or led astray by legislative history that drives them off the correct interpretation that plain statutory text would otherwise indicate. For any of these reasons, simplemindedly maximizing intentionalism will do worse, from a higher-order perspective, than optimizing intentionalism that employs a stopping rule. Again, I do not argue here that the plain meaning rule is indeed the best stopping rule for intentionalists; to know that, we would have to know more than we currently do. I only mean to indicate a type of justification for the plain meaning rule that is invisible to the maximizing interpreter.

Satisficing. A different stopping rule for intentionalist interpreters would be to search until, but only until, an interpretation turns up that is satisfactory. The optimizing intentionalist, in a case like Caminetti, proceeds on the basis of a rule that is calculated to produce the best possible interpretation given the resource constraints, including limited time and limited cognitive capacity, under which the interpreter labors. The satisficing intentionalist employs a different approach. Rather than searching for the best possible interpretation, even under constraints, the satisficer sets a limit to search by accepting the first interpretation that is good enough.

A satisficing interpreter might come to the same result as did the Caminetti court, not on the ground that the costs of further search for evidence of legislative intentions would be greater than the expected benefits, but simply on the ground that the plain meaning of the text provided an account of legislators’ intentions that was internally consistent, intuitively plausible, and in that sense good enough. The idea animating this relaxed attitude is that a maximizing search for the very best account of legislators’ intentions would be a sort of local perfectionism; and local perfectionism would make the interpretive system worse off, from a higher-order point of view. As with

optimizing, satisficing intentionalism may produce a better allocation of time and effort across a whole array of cases than does simpleminded maximizing in each particular case. As compared to the optimizer, the satisficer uses a different decision rule to produce the globally-maximizing allocation, but the aim of global maximization is the same.

An interesting implication of the satisficing account is that the interpretation the court produces is sensitive to the order in which materials are considered. In our cafeteria example, which offering the satisficing consumer ends up choosing depends upon which end of the cafeteria she starts from (assuming there are satisfactory offerings at various points in the line). In *Caminetti*, a court that (1) employed a satisficing stopping rule but that (2) considered legislative history before statutory text would conclude that the legislative history offered a fully satisfactory account of legislators' intentions. The bare idea of satisficing, by itself, cannot justify a rule that intentionalist interpreters should stop with the plain (and satisfactory) meaning of the text, as opposed to the plain (and satisfactory) meaning of the legislative history. But this is no serious objection to the satisficer. In any decisionmaking setting in which less than all possible alternatives are to be searched out and considered, one must begin somewhere. The satisficing interpreter will satisfice at this higher level also, accepting any starting point that is good enough.

The starting point will thus be set by convention, within any particular legal system. In our legal system, the convention is that statutory text is the starting point, and intentionalist stopping rules constrain the search for evidence beyond the text. The convention might be otherwise, as far as the satisficing perspective goes, but it actually isn't otherwise. The satisficing judge in our legal system has no reason to lose sleep over what satisficing interpreters might do in other, possibly counterfactual legal systems.

The allure of maximizing intentionalism. A simple idea, which seems intuitive to many, is that the intentionalist interpreter should consider all relevant and probative evidence of legislative intentions in the case at hand. To deny this is to lose the rhetorical high ground (a point to which I return in Section IV): the maximizing intentionalist can always lampoon the optimizing or satisficing intentionalist by pointing to some particular case in which an optimizing or satisficing stopping rule would cause the interpreters to miss out on highly probative evidence of intentions. On either the optimizing or satisficing perspectives, however, this sort of argument is a simpleminded mistake, if it counsels neglect of the costs of decisionmaking and the risks of error (the optimizer's worry) and neglect of the perils of local perfectionism (the satisficer's worry).

If simpleminded maximizing intentionalism is mistaken, it is also a prominent and often dominant strand in American legal interpretation. Consider the famous *Holy Trinity* case, in which the Court without much ado discarded the traditional rule against consulting internal committee reports as evidence of legislative intentions. The Court said very little to justify this crucial methodological move, but it might have seemed the most natural thing in the world. After all, a famous intentionalist injunction held that "where the mind labours to discover the design of the legislature, it seizes everything from which aid can be derived." The pre-*Holy Trinity* court had combined this rule with a rule barring judicial consideration of internal legislative evidence; to the *Holy Trinity* court, this must have seemed an odd, even incoherent combination, akin to telling a jury to "consider all relevant evidence" while excluding the smoking gun the police seized from the defendant.

In both the exclusionary-rule setting and the interpretive setting, however, it is familiar that there may be good higher-order reasons to adopt seemingly conflicted rule-combinations of this sort. What I add here is a taxonomy of, and contrast between, two different interpretive decision strategies grounded in higher-order considerations. Those strategies have a common enemy or antonym, however; both are alternatives to the sort of simpleminded maximizing intentionalism that has been so prominent in the history of American interpretive theory and practice.

Clause-bound textualism and holistic textualism

The previous example assumed an intentionalist account of interpretation. Here I offer an example premised on a strictly textualist account of interpretation, to show that textualist interpreters also face the choice among maximizing, optimizing and satisficing styles of interpretation. The choice among decisionmaking strategies, I have suggested, is entirely agnostic as among various value theories different interpreters might hold; textualists as well as intentionalists must pursue their aims by means of some decision-procedure.

Suppose an interpreter believes that the aim of interpretation is to capture the ordinary meaning of statutory or constitutional text, quite apart from anyone's intentions. (I bracket here, as irrelevant for our purposes, the possible justifications for this view). The textualist interpreter faces a range of implementation issues that must be confronted to make her high-level commitment operational. How exactly should the textualist view be embodied in decision-procedures that interpreters will use? Of these implementation questions, I consider only the following: How much text should the interpreter consider? Suppose there is both a primary text, a statutory section or constitutional clause whose interpretation will settle the rule of law that applies between the parties, and collateral texts, such as other provisions of the Constitution or of the relevant statute, or other statutes. How widely should the textualist interpreter cast her net, and how much weight should be given to collateral texts?

This issue underlies important debates in both statutory and constitutional interpretation. In the statutory arena, the Court has at times adopted a strong presumption of textual coherence across whole statutes, on the view that textual similarities and differences across provisions are at least presumptively significant. Justice Scalia's opinion for the Court in *West Virginia University Hospitals v. Casey* goes even further, suggesting that interpreters should treat the whole U.S. Code as though terms are used consistently across statutes. In constitutional interpretation, Akhil Amar defends an "intratextualist" view that makes extensive use of comparisons across clauses, even to the point of insisting that words appearing in widely separated contexts be given similar meanings. Elsewhere in the constitutional arena, there are occasional hints of an even more expansive "intertextualist" view, analogous to the Casey opinion; on this view, the Constitution would be read in light of collateral legal texts, such as the Declaration of Independence and the Northwest Ordinance. What these views have in common is a (more or less expansive) commitment to holistic or coherent textualism, as opposed to the sort of clause-bound textualism that focuses principally or solely on the statutory or constitutional provisions directly applicable in the case at hand.

Maximizing. The textualist interpreter who is also a simpleminded maximizer seeks the ordinary meaning of a legal term in the provision at hand. Other provisions of the same text, or other legal

texts, use the same term or use different terms in ways that illuminate by contrast. The simplemindedly maximizing textualist reasons that interpreters should consult collateral texts as widely as possible, and give them whatever weight they intrinsically deserve, all things considered, to whatever extent they happen to provide relevant evidence of ordinary meaning. Indeed, on this construal there is nothing at all special about legal texts; any sources of ordinary meaning will do, such as dictionaries, literature or the testimony of linguists. The reductio ad nauseam of the search for ordinary meaning is not Casey. It is the debate in Muscarello v. United States over the meaning of the statutory term “carry,” as in “carry a firearm.” The case occasioned heated subsidiary debates over usage in the Bible, Melville, Defoe, and M*A*S*H, and saw Justice Breyer (or his clerks) attempt to discern ordinary meaning by collecting examples from a database of national newspapers and then “randomly” sampling the results—a procedure that fails to account for selection bias in the database itself, akin to taping speeches at a convention of spit-and-polish grammarians and then “randomly” sampling the results to see whether it is ordinary usage to split infinitives.

Optimizing and satisficing. Simplemindedly maximizing textualism fails on the same grounds that condemn simplemindedly maximizing intentionalism. The key points are familiar. First, holding constant the quality of decisions (from a textualist perspective), the direct decision costs and opportunity costs of holistic textualism are real, and perhaps quite high. The refined and comprehensive comparisons required by maximizing textualism take time. The time spent searching out and comparing usage across the whole Code, or within a database, means less time to be spent on refined textualism in other cases. Second, holding the costs of decisionmaking constant holistic textualism may do worse even according to the very value theory presupposed by maximizing textualism itself. As compared to clause-bound textualism, holistic or expansive textualism requires a more complicated and information-intensive inquiry, one that will reduce decisional accuracy whenever fallible interpreters read the comparison texts mistakenly. There is no particular reason to think that the illuminating effect of holistic textualism will predominate over its error-producing effect. Third, and related to the last, holistic textualism in the hands of fallible interpreters risks producing a holistic, highly coherent, but fundamentally mistaken analysis, one that enforces a simultaneous misreading of a whole set of related provisions. Risk-averse interpreters might prefer the limited incoherence of clause-bound interpretation to a sweeping, integrated, but erroneous universal account.

In light of these considerations, both optimizers and satisficers will hold that textualists need a stopping rule that constrains the ever-expanding search for evidence of ordinary meaning. A consequence of this is the standard tradeoff of local inferiority for global superiority. Optimizing textualists and satisficing textualists will both be willing to render interpretations in particular cases that are inferior from a textualist standpoint to those that an infallible maximizing textualist would render. The hope, for both optimizers and satisficers, is that local inferiority will prove globally superior across a set of decisions—again, superior on the same value theory that the locally-maximizing textualist holds.

But optimizers and satisficers will employ different stopping rules to promote their common aim of superior performance from a global perspective. The optimizing textualist attempts to calculate the point at which the expected marginal benefit of further evidence of textual meaning is equal to or less than the expected marginal cost, and stops there. The satisficing textualist stops when the

text(s) she has examined point to an ordinary meaning that seems satisfactory—one that is plausible to her, in light of her native linguistic competence, that suggests a rational legislative policy, and so forth. Perhaps a database search, or a deep understanding of Melville, would make the meaning on which the satisficing interpreter has alighted seem less satisfactory than it currently does. But the satisficer refuses to lose sleep over that possibility, not because she anticipates that the costs of further search equal or exceed the expected benefits, but because she believes that the perfectionist search for the very best local interpretation is often a self-defeating enterprise, from a global perspective. She believes that yielding to the siren song of local perfectionism will make her performance worse, overall, than does her resolute (even phlegmatic) disposition to take the good in preference to searching for the best.

As usual, the two stopping rules might happen to converge to similar results. An opinion that both optimizing textualists and satisficing textualists might love is Dewsnup v. Timm, in which the Court interpreted the phrase “allowed secured claim” in §506(d) of the Bankruptcy Code while ignoring holistic arguments that the phrase “allowed secured claim” in §506(a) of the Bankruptcy Code might have something to do with the matter—adopting instead what Justices Scalia and Souter, in dissent, derided as a “one-subsection-at-a-time approach to statutory exegesis.” To the intentionalist or textualist maximizer, opinions like Dewsnup seem foolishly narrow. Why not consider an obviously relevant collateral provision when interpreting the provision at hand? To the optimizer or the satisficer, however, stopping with §506(d) might be a good idea, either because fathoming the complexities of §506(a) is itself a daunting task for generalist interpreters (so daunting that the costs of a wider inquiry outrun the expected benefits), or because the Court’s straightforward interpretation of §506(d) was itself a satisfactory local solution.

Before we leave these examples, a general caution is in order. It has not been my enterprise here to defend, substantively, any particular interpretive decision-procedures or results. Whether the optimizer or satisficer should actually defend an opinion like Dewsnup will depend on the precise shape of the cost-benefit curves (to the optimizer) or the level of judicial aspiration (to the satisficer). I merely aim to illustrate two justifications for using truncated interpretive decision-procedures, justifications that must forever remain invisible to the maximizing interpreter. Nor do I mean to express any view about which interpretive strategy it would be best for judges in particular to use. A general issue is that suitable interpretive decision-procedures will vary with the institutional capacities of different interpreters occupying different roles, and with regard to the systemic interactions of different institutions. I confine myself here to mapping the conceptual terrain; choosing interpretive decision-procedures for particular interpreters turns on empirical questions, about institutional capacities and systemic effects, that I do not address.

III. Optimizing or Satisficing?

If simpleminded local maximizing is an untenable interpretive posture, which stopping rule should interpreters use? Proponents of optimizing, on the one hand, and satisficing, on the other, have each offered generic arguments that seek to knock out the rival stopping rule, leaving their preferred rule in sole possession of the field, as the only valid alternative to simpleminded maximizing. Here I offer some brief remarks on these generic arguments, only to suggest that they fail. There is no knockout argument against either optimizing or satisficing as decisionmaking strategies in the service of global goals. The right assessment is more catholic: both strategies have virtues and vices

that vary across contexts, and the two strategies can often be usefully combined. In general, and thus in the interpretive setting as well, optimizers and satisficers can join forces to reject simplemindedly maximizing interpretation without having to stage a second-round runoff to choose a single winner. The choice between optimizing and satisficing will turn on particular features of the decisionmaking context, especially the nature and size of the stakes.

Proponents of satisficing tax optimizing with a problem of infinite regress. Optimizers stop searching for better alternatives, including better interpretations, when the costs of further search exceed the expected gain in new information (about legislative intentions, ordinary meaning, and so on). Yet information is a good with unusual properties. One cannot know what the value of new information will be until one has it. "To maximize subject to the constraint of information cost one would have to know the expected value of information, but this is not in general possible." On this view, satisficing is the only alternative to simpleminded maximizing, because optimizing under conditions of limited information is a contradiction in terms.

The infinite-regress problem, however, is best understood as a contextual caution rather than a foundational difficulty. In some settings, especially when making a string of similar decisions, decisionmakers can confidently expect that the next bit of information to turn up will be of limited value, even if they do not know what that information will be. Searching a very large database of employment opportunities, I might rationally calculate that the costs of further search exceed the benefits. Even if I cannot know what the next opportunity I turn up will be—perhaps it is something far superior to any of the ones I have seen so far—I can rationally assess the chances of its being superior, because I have observed a long series of similar entries, know something about how the database was compiled, and thus have no good reason to expect anything different in the future. Experience supplies prior probabilities from which decisionmakers can generate rational expectations about the benefits of future information.

In the interpretive setting, interpreters or other interpreters can draw upon prior experience to form rational assessments of the value of the next bit of information about, say, legislative intentions, even if they do not know the content of that information. Much will depend upon the conditions under which optimizing is deployed. Some interpretive decisions are of a sort that are repeated frequently, some are not. Where experience suggests that new options or new information can be expected to differ somewhat but not radically from previous options or information, so that the expected value of further search is predictably a medium-sized quantity that can be meaningfully compared to the costs of further search, it would be hard to deny that optimizing interpretation is possible.

So the argument from infinite regress cannot knock optimizing out of the ring. Is there a general argument against satisficing? The most common criticism is that satisficing is arbitrary, because the bare idea of satisficing says nothing about where exactly the satisficer's aspiration level should be set. Why should interpreters be satisfied with this much rather than that much evidence of legislative intentions, or ordinary meaning? Again, however, the force of this point varies across settings. Aspiration levels need not be exogenously fixed once and for all; they can be endogenously formed, as decisionmakers acquire experience with the relevant settings. In many domains, experience will suggest that the band within which an aspiration level might plausibly be set will be quite narrow. In the cafeteria example, no one would continue the search until a gourmet

meal turns up, simply because it is quite predictable that no such meal will ever turn up; the context will not support any such aspiration. Although such practical considerations are not conceptually satisfying, it is not clear that satisficing rules themselves have or need have any aspiration to be conceptually satisfying. So long as they fulfill their primary mission of truncating local search at a point that is globally maximizing, they are good enough.

These mutual criticisms by optimizers and satisficers point to a commonsensical conclusion. Neither optimizing nor satisficing is universally or acontextually best; the choice between stopping rules will and should vary as the features of the interpretive setting vary. A large challenge, one I cannot take on here, would be to delineate the conditions under which one stopping rule or the other is superior. As a preliminary matter, however, a useful idea is to focus on the size of the stakes at hand in the decision. On one hand, with very low-stakes decisions, satisficing appears most attractive, relative to optimizing. “The less we care about the gap between satisfactory and optimal toothpaste, for example, the more reason we have to satisfice—to look for a satisfactory brand and then stop searching when we find it.” On the other hand, where the stakes are very high, especially where decisions involve a choice between seemingly incommensurable values or life plans, satisficing also appears useful, perhaps inevitable. In the choice whether to become a doctor or a soldier, it is quite unclear what it would mean to collect an optimal amount of information before making a decision, and the best we can hope for is to make a satisfactory choice. On this view, optimizing makes most sense for decisions with medium-sized stakes. In those settings, the expected benefit of collecting more information or generating new options is appreciable but also limited, and thus meaningfully comparable to the costs of further search.

Many interpretive decisions are plausibly medium-stakes problems, in which any of the possible interpretations will be consequential for the legal system but none will be of overwhelming importance. An even larger fraction of interpretive decisions, however, are the law’s equivalent of picking toothpaste in a store, at least from the interpreter’s point of view. These are cases important to the parties at hand, but with few broader consequences for other cases, actors, or problems. Such cases probably dominate in the interpretive work of administrative agencies and lower courts; some might claim that they are not uncommon on the Supreme Court’s docket. To move beyond the conceptual mapping of interpretive strategies I offer here, we would need to think about which interpretive strategies might be best at different levels of the legal system, when used by different interpreters with distinct roles and institutional capacities. We might find, for example, that satisficing is best for interpreters in the low-stakes settings that predominate in the ordinary work of law, while interpreters working in more consequential settings might do better to optimize. Here I merely indicate the importance of these institutional questions, without offering answers to them.

IV. The prevalence of maximizing: some mechanisms

I have criticized maximizing interpretation and offered qualified praise, varying across contexts, for the optimizing and satisficing alternatives. But if maximizing interpretation is so simpleminded, why does it have so many defenders—both on the bench and in the academy? Here I will survey some mechanisms that tend to push interpreters past the point at which optimizers or satisficers will abandon the hope of local perfection and cut off the search for further evidence.

Self-interested motivation. For Bentham, the mechanism was obvious: self-interested collusive behavior on the part of judges, lawyers and other legal actors. Maximizing interpretation is, above

all, complex interpretation. The simplemindedly maximizing local interpreter is far more likely to take into account a rich and complex array of considerations, sources and evidence bearing on the interpretive problem at hand, whereas. For Bentham, complexity was the fruit of conspiracy. Judge & Co. benefit jointly and severally from a highly complex, even mysterious legal system in which laypeople must pay rents to lawyers to steer through the thickets that the lawyers themselves have created. Although the claim rests on far too simpleminded an account of the motivations of lawyers and judges, it had better institutional foundations when Bentham wrote; consider that judges were commonly paid from the fees of litigants, and that one of Bentham's principal reforms was to urge that judges be put on regular government salaries.

Cost-externalization. Peter Schuck puts Bentham's argument in more modern terms, and less conspiratorial ones, by observing that the costs and benefits of complex legal interpretation are unequally distributed over different actors and groups. If maximizing interpretation is complex interpretation, as I have claimed, then Schuck's argument also suggests that maximizing interpreters may push the search for evidence beyond the bounds that would limit an optimizing or satisficing interpreter, just because maximizing interpreters do not themselves internalize all the costs—including opportunity costs and decision costs—of maximizing interpretation. Consider a possible claim that maximizing intentionalism, with its exhaustive search through legislative history and other sources, is feasible for judges only because many of the resulting decision costs are externalized onto litigants and other legal actors.

Neglect of opportunity costs. Bentham's account, as updated by Schuck, points to self-interested motivations on the part of lawyers, judges, and other legal actors. But there is a cognitive mechanism operating as well, one that plausibly affects even public-spirited actors. I refer to the neglect of opportunity costs, relative to more visible direct costs. Neglect of opportunity costs tends to make decisionmaking excessively intensive and complex, in the maximizing style, rather than brisk and mechanical, in the optimizing and satisficing styles. Decisionmakers focus to excess on the costs of getting this decision wrong, while overlooking the costs that a protracted process of decisionmaking itself creates. "The neglect of the opportunity costs that are created by the fact that decision making takes time is . . . an important and pervasive source of irrationality."

An ethical analogy, and the distorting force of particulars. An analogy to ethics is useful here. The choice between maximizing, on the one hand, and optimizing or satisficing on the other, has some of the same intellectual structure as the choice between rule-utilitarianism, which counsels that ethical subjects act in accordance with rules whose observance will maximize overall utility, and act-utilitarianism, which counsels ethical subjects directly to choose whichever action maximizes overall utility. The rule-utilitarian acknowledges that the utility-maximizing rules may sometimes call for actions that, when viewed in isolation, are locally suboptimal from the utilitarian point of view. The rule-utilitarian, then, will sometimes be placed in the awkward position of defending acts whose immediate effect is, when viewed in isolation, socially detrimental. So too, it is the easiest thing in the world for maximizing interpreters to emphasize specific cases in which stopping rules produce interpretive blunders, relative to an all-things-considered approach that would (if applied by an infallible interpreter with infinite time to make decisions) have taken into account relevant information excluded by a stopping rule.

A corollary is that second-order interpretive strategies suffer from the distorting force of particulars. Maximizers will always be able to point to lurid examples in which second-order strategies produce suboptimal results, even if second-order strategies are best from some overall perspective. Such examples pack a strong rhetorical punch, but the problem goes deeper than that. The vivid costs in particular cases may trigger cognitive failings in the audience that is to evaluate the competing decisionmaking strategies, causing them to overreact to specifics while ignoring the crucial question of overall justification. A crucial mechanism here is saliency, a heuristic that causes decisionmakers to overweight the importance of vivid, concrete foreground information and to underweight the importance of abstract, aggregated background information.

Conclusion

I conclude with some mixed notes of pessimism and optimism. If simpleminded local maximizing is intellectually untenable, it is also remarkably persistent in the theory and practice of legal interpretation, by virtue of the mechanisms discussed in Part IV. We can at least hope, however, to make interpreters and students of interpretation aware that the maximizing style is not inevitable, that it represents a particular choice among local decisionmaking strategies, and that local maximizing is by no means obviously desirable from the globally-maximizing perspective. This point is only preliminary, for it does nothing to indicate what interpretive strategies are actually best in various settings. The subsequent questions are empirical, and involve the institutional capacities of interpreters. But in a legal system that still respects the claims of maximizing interpretation, it is a good place to start.