

STATIC VERSUS DYNAMIC DISCLOSURES, AND HOW NOT TO JUDGE THEIR SUCCESS OR FAILURE

Richard Craswell*

Abstract: Disclosure laws can serve many different purposes. This Article is the first to distinguish two of those purposes, which I call *static* and *dynamic* disclosures. In brief, static disclosures aim to improve consumers' choice from among the set of products that are already available on the market. By contrast, dynamic disclosures aim to improve the range of products from which consumers must choose, by sharpening sellers' incentives to improve the quality of their products.

The Article also discusses the various ways in which the effects of static and dynamic disclosures might be measured and evaluated. In doing so, it examines and mildly criticizes the position recently advanced by Professors Omri Ben-Shahar and Carl Schneider, who argue (approximately) that disclosure almost *never* works, and that it should not even be considered as a policy option. While I agree with much else that Professors Ben-Shahar and Schneider say, their claim that disclosures almost *never* work is far too broad.

INTRODUCTION	334
I. BACKGROUND	335
II. THE PURPOSE(S) OF DISCLOSURE	337
III. SOME SIMPLE ECONOMICS	340
IV. CRITERIA FOR SUCCESS: STATIC DISCLOSURES	345
A. Measuring Static Effects	346
B. Lessons from the Data	347
C. Identifying the Costs and Benefits	348
D. Do Static Disclosures Always Fail?	350
V. CRITERIA FOR SUCCESS: DYNAMIC DISCLOSURES	354
A. Some Examples	355
1. Credit Contract Scores	355
2. Restaurant Grades	356
3. Automobile Fuel Economy Ratings	357

* William F. Baxter – Visa International Professor of Law, Stanford Law School. Although I formerly worked for the Federal Trade Commission (1977-1983), I have no affiliation with any of the parties discussed here.

This paper was prepared for the symposium on “The Disclosure Crisis” held at the University of Washington Law School, February 28, 2013. A distant ancestor of this paper benefited from helpful comments from participants in that symposium, as well as from Andrew Daughety, Ryan Calo, Barbara Fried, Daniel E. Ho, A. Mitchell Polinsky, Jennifer Reinganum, Carl E. Schneider, and members of a faculty workshop at the Stanford Law School.

4. The Cigarette “Tar and Nicotine Derby”	358
5. Food Labeling and Salad Dressing	358
B. Evaluating Success or Failure	359
1. Gathering the Data	359
2. Eliminating Exogenous Causes	362
3. Isolating the Effect of the Legal Requirements	365
4. Designing More Effective Disclosures	370
5. Identifying the Costs and Benefits	371
VI. CRITERIA FOR SUCCESS: DISCLOSURE VERSUS	
QUALITY REGULATION	372
A. Restaurant Grades Revisited	373
1. Disclosure Versus No Regulation	374
2. Disclosure Versus Direct Quality Regulation	375
3. Asking the Right Question	376
CONCLUSION	379
MATHEMATICAL APPENDIX	380

INTRODUCTION

In this Article I have three ambitions. First, I hope to show that we cannot evaluate the success or failure of any disclosure law without considering the possible goals that law might have had. I do not take this point to be hugely controversial.

My second ambition, however, is to improve our understanding of two particular purposes that disclosures might serve. To this end, I distinguish here between what I will call *static* and *dynamic* disclosures. Static disclosures take a consumer’s existing range of choices as more or less given, and aim merely to improve a consumer’s choice from among the existing choice set. By contrast, dynamic disclosures seek to improve the existing choice set by creating incentives for sellers to improve the quality of their offerings. This distinction has not yet been discussed in the disclosure literature, but I hope to show that it has important implications for how the success or failure of disclosures can best be measured.

Finally, my third ambition is to illustrate the perils of trying to analyze disclosure laws *without* paying attention to the specific purposes that different disclosure laws might serve. To illustrate these perils, I use as a recurring example the recent and provocative Article, “The Failure of Mandated Disclosure,” by Professors Omri Ben-Shahar and Carl E. Schneider.¹ It is perhaps unfair to single out this Article in such a way,

1. Omri Ben-Shahar & Carl E. Schneider, *The Failure of Mandated Disclosure*, 159 U. PA. L. REV. 647 (2011).

for (as I discuss below) there is much that is good in the Article, and I agree with many of its conclusions. However, the systematic way in which that Article analyzes a wide range of disclosure laws (which is one of the Article's strengths) also makes it an ideal Article in which to find an occasional cautionary example showing what happens when the purposes of disclosure laws are not properly understood.

With that in mind, let us proceed. Section I, below, provides some necessary background by describing Professors Ben-Shahar and Schneider's views in slightly more detail, and by relating their Article to the rest of the academic literature on disclosures. Sections II and III then develop at more length the distinction I wish to draw between static and dynamic disclosures, with Section III providing an economic interpretation of that distinction. Finally, Sections IV through VI discuss in more detail some criteria for evaluating the success or failure of various kinds of disclosure, to show how those evaluations should differ depending on whether the disclosure has static or dynamic aims.

I. BACKGROUND

Mandatory disclosure, we are told, is a regulatory technique that is “much used but little remarked.”² However, while I fully agree that disclosures are “much used,” the case for them being “little remarked” is doubtful. Professors Ben-Shahar and Schneider themselves cite dozens of studies of the effects of various disclosure regimes; and they could easily have added more, had they not worried about overburdening their readers.³ For example, mention should certainly be made of the 2007 book by Archon Fung, Mary Graham, and David Weil, which compiled eighteen case studies of various disclosure regimes.⁴ Unlike Professors Ben-Shahar and Schneider, these authors came to a more mixed view of the efficacy of disclosures (some worked well; others didn’t). Similarly mixed views can be found in the marketing literature, in several recent survey Articles and meta-analyses.⁵

2. *Id.* at 649.

3. *Id.* at 652.

4. ARCHON FUNG, MARY GRAHAM & DAVID K. WEIL, *FULL DISCLOSURE: THE POLITICS AND PERILS OF TRANSPARENCY POLICIES* (2007). This book is cited only once, and briefly, by Ben-Shahar & Schneider, *supra* note 1, at 679 n.166.

5. E.g., Jennifer J. Argo & Kelly J. Main, *Meta-Analyses of the Effectiveness of Warning Labels*, 23 J. PUB. POL'Y & MARKETING 193 (2004); Anthony D. Cox, Dena Cox & Gregory Zimet, *Understanding Consumer Response to Product Risk Information*, 70 J. MARKETING 79 (2006); Eli P. Cox III, Michael S. Wogalter, Sara L. Stokes & Elizabeth J. Tipton Murff, *Do Product Warnings Increase Safe Behavior? A Meta-Analysis*, 16 J. PUB. POL'Y & MARKETING 195 (1997); Baruch

Indeed, this concern with the possible effects of disclosure (both good effects and bad) is not even particularly new. Over thirty years ago, an Article in the *Journal of Law and Economics* began its discussion by observing that:

Consumer protection regulation has come under increasing fire from Congress, the courts, and the business community. . . . One response to these charges has been a movement away from traditional forms of regulation and toward interventions that are more compatible with consumer and seller incentives. In particular, there has been increased interest in techniques which ensure that consumers have sufficient information to protect themselves against unsafe products or unfair seller behavior. . . . [However,] these simple prescriptions . . . mask many of the complexities involved in the ways in which information is communicated to consumers and the ways that consumers (and the market) respond.⁶

For all these reasons, I think a better interpretation of the “little remarked” comment is that it is a statement about what gets remarked *in law school classes*. On that point, Professors Ben-Shahar and Schneider are absolutely correct. Most law school curricula allot little if any time to issues of how (and whether, and with what effect) consumers’ information might be improved. Indeed, I am struck by the fact that today, in 2013, we still make it very hard for any law student to graduate without getting at least a brief exposure to the rule against perpetuities, or to the important problems that arise if there are two ships in the world that are both named *Peerless*.⁷ But it is entirely possible (and not even very difficult) to graduate from most law schools today without being exposed to *any* serious study of the costs and benefits of consumer disclosures. Thus, even if Professors Ben-Shahar and Schneider succeed only in bringing this issue to the attention of a new generation of law professors, that would be a real contribution. Of course, I hope they will succeed in other ways, too—but even if that were *all* they could do, it

Fischhoff, Donna Riley, Daniel C. Kovacs & Mitchell Small, *What Information Belongs in a Warning?*, 15 PSYCHOL. & MARKETING 663 (1998); David W. Stewart & Ingrid M. Martin, *Intended and Unintended Consequences of Warning Messages: A Review and Synthesis of Empirical Research*, 13 J. PUB. POL’Y & MARKETING 15 (1994); William L. Wilkie, *Affirmative Disclosure at the FTC: Objectives for the Remedy and Outcomes of Past Orders*, 4 J. PUB. POL’Y & MARKETING 91 (1985).

6. Howard Beales, Richard Craswell & Steven Salop, *The Efficient Regulation of Consumer Information*, 24 J.L. & ECON. 491, 491 (1981).

7. If my assertion in the text is accurate, no lawyer or law student is likely to need this citation, but I supply it anyway: *Raffles v. Wichelhaus*, (1864) 159 Eng. Rep. 375 (Ex.).

would still be impressive.

As should by now be apparent, I agree with roughly 85% of what Professors Ben-Shahar and Schneider have to say. In particular, I agree that disclosure requirements are used too often, are designed by institutions with little understanding of their potential pitfalls, and often produce small (or even negative) effects. I also agree that the apparent attractions of disclosure regimes sometimes offer lawmakers a convenient excuse for ducking the more difficult choices they would have to face if they were serious about choosing an optimal consumer policy. I have written about these issues elsewhere, so I will not repeat that analysis here.⁸ But to the extent that there are “sides” in the debate over consumer disclosures, I am mostly on the side of Professors Ben-Shahar and Schneider.

Still, the mutual praise of academics is usually of interest only to the individual academics being praised. In an attempt to interest a wider audience, in this paper I take up some issues on which I may disagree with Professors Ben-Shahar and Schneider; or at least (to put it more fairly) some issues that were not those authors’ principal focus. Specifically, I consider the following questions: By what criteria should we judge whether a disclosure succeeded or failed? And should those criteria be different for different kinds of disclosure?

II. THE PURPOSE(S) OF DISCLOSURE

Professors Ben-Shahar and Schneider do not claim that they know, or even that they can estimate, the exact percentage of disclosure laws that “fail.” For example, they do not claim that the disclosure laws they describe exhaust the entire universe of disclosure laws. Nor do they claim their study includes even a random sample from that universe, for they explicitly exclude non-mandatory disclosures, which (I will argue below) skews their sample in important ways. They also exclude even some successful disclosures, without really explaining why.⁹

In spite of these limits on their analysis, though, Professors

8. See, e.g., Beales et al., *supra* note 6; Richard Craswell, *Taking Information Seriously: Misrepresentation and Nondisclosure in Contract Law and Elsewhere*, 92 VA. L. REV. 565 (2006); Richard Craswell, “Compared to What?” *The Use of Control Ads in Deceptive Advertising Litigation*, 65 ANTITRUST L.J. 757 (1997) [hereinafter Craswell, *Control Ads*]; Richard Craswell, *Regulating Deceptive Advertising: The Role of Cost-Benefit Analysis*, 64 S. CAL. L. REV. 549 (1991) [hereinafter Craswell, *Cost-Benefit Analysis*].

9. Ben-Shahar & Schneider, *supra* note 1, at 678–79 (“We could go on. . . . We could also point to some moderate successes of disclosure regimes, particularly those that rely on rating systems. But enough. Let us move on. . . .”) (footnote omitted).

Ben-Shahar and Schneider are not shy about extrapolating from their limited sample, leaving readers to believe that disclosures almost always fail. At various places in the Article, they claim that disclosure laws “regularly” and “chronically” fail, and that the costs of those laws “generally swamp” their benefits.¹⁰ At one point, they even describe the failure of disclosure laws as “inevitable”—though they may later retreat from that particular description.¹¹

Moreover, Professors Ben-Shahar and Schneider present these conclusions not as their own speculations, but as uncontested empirical fact. For example, the section of the Article where they review the empirical evidence is not given some tentative or diffident title, like “Some Possible Failures of Mandated Disclosure.” Instead, their section heading is emphatic “THE DOCUMENTED FAILURE OF MANDATED DISCLOSURE.”¹² Elsewhere in their Article, they are even more explicit: “the empirical evidence *show[s]* that mandated disclosure regularly fails in practice.”¹³ As I said at the outset, I happen to agree with Professors Ben-Shahar and Schneider that many disclosure laws fail. I do not, however, agree with their repeated implications that almost *all* disclosures fail. More fundamentally, I also do not agree that the kinds of evidence relied on by Professors Ben-Shahar and Schneider are sufficient to let us decide whether any particular disclosure law has failed. As a result, I cannot agree that their verdict (chronic and all-but-inevitable failure) is supported by the available empirical evidence.

At bottom, all of my disagreements with Professors Ben-Shahar and Schneider ultimately stem from their failure to defend (or even to identify) any criterion for judging the success of a disclosure law. This omission is critical, because any conclusion that some disclosure law *fails* must depend, at least implicitly, on the standard we set for judging *success*. Set the bar for success high enough, and every disclosure is a failure. Set the bar low enough, and every disclosure succeeds.

Now, to some readers, and perhaps to Professors Ben-Shahar and Schneider, it might seem that the purpose of disclosure laws is obvious; and that this purpose is indeed the same for all disclosures across the board. After all, disclosure (by definition) is about giving people information, and surely one of the main purposes for giving people

10. *Id.* at 651.

11. *Id.* (“Not only does . . . mandated disclosure regularly fail[] in practice, but its failure is *inevitable*”) (emphasis added). *But see id.* at 679 (disclosure “regularly—*though not inevitably*” fails) (emphasis added).

12. *Id.* at 665 (emphasis added).

13. *Id.* at 651 (emphasis added).

information is to help them make better decisions. When Professors Ben-Shahar and Schneider do talk about the possible purposes of disclosure laws, this goal—helping people make better decisions—is the one they mention most frequently.¹⁴ This is a good start, though when it is stated at that level of generality, it leaves open a number of quantitative questions about where to draw the line. For example, *how many* people must be helped in order for a disclosure to count as a success? And how much should we be willing to spend to improve the decision making of 14% of consumers?

In any event, at one point in their Article Professors Ben-Shahar and Schneider do recognize that disclosure laws might have other purposes in addition to the open-ended goal of helping people make better decisions. In particular, they identify at least four possible purposes that are sometimes advanced in support of disclosure requirements:

[Mandated disclosure] rests on a plausible assumption: that when it comes to decision-making, more information is better than less. More information [1] helps people make better decisions, thus [2] bolstering their autonomy. Since people can no longer customize most transactions, disclosure [3] helps restore some individual control. It may also [4] induce enterprises to behave more efficiently.¹⁵

In this paper, I will say little about goals [2] and [3], since the application of autonomy theories to disclosure policy is underdeveloped, and in any case it is not my area of expertise. But the other two goals, numbered [1] and [4] in this list, correspond at least approximately (or could be read as corresponding) to what I refer to as static and dynamic disclosures. Recall: static disclosures take the options consumers face as given, and try to help consumers make a better choice from among those options. In the list quoted above, this is item [1]: disclosure that “helps [some] people make better decisions.” By contrast, dynamic disclosures aim to give sellers better incentives to offer improvements over whatever is currently available. In Professors Ben-Shahar and Schneider’s list, this is an example of item [4]: disclosures that “induce [some] enterprises to behave more efficiently.”

Of course, these two goals are not the only ones disclosure might serve, so this is not meant as an exhaustive list. For example, some disclosures aim to explain how to use a product, or tell buyers where to return a product if it is defective. Disclosures with either of these goals

14. *E.g.*, *id.* at 649, 729, 746.

15. *Id.* at 650 (numbers added).

in mind may look very different from the ones I discuss here, with different implications for everything from the timing of the disclosure (pre-sale or post-sale?) to what medium ought to be used (package insert or a permanent page on the seller's website?).

It is also true that any given disclosure might have more than one goal. For example, some disclosures might attempt simultaneously to improve consumers' choices from the currently available set (the static goal, in my terms); *and* to give sellers an incentive to improve that set (the dynamic goal). Even when disclosures have mixed purposes, though, it is still useful to distinguish those goals, for purposes of evaluating a disclosure's success or failure. After all, if different disclosures have a different *mix* of goals, then the criteria for success or failure will usually have to be different too.

Section III, below, presents a theoretical economic model (the technical mathematics of which is left to a short Appendix) to show what a disclosure regime would have to do to achieve these two goals, and to show why the requirements for success at one goal are not the same as the requirements for success at the other. Sections IV and V then discuss some concrete examples of disclosures that may or may not have been successful, to show how various criteria for success might be applied. Finally, Section VI considers a different criterion for success—are the disclosure's dynamic effects strong enough to eliminate the need for direct regulation of product quality?—to show how this criterion requires still a different set of inquiries.

III. SOME SIMPLE ECONOMICS

Consider a hypothetical market in which sellers produce a relatively complex product—an automobile, say, or a wireless Internet router or a nonprescription drug. Each of these products has multiple attributes that might vary along several dimensions. To keep the analysis simple, though, I will focus on a single attribute (s) which could conceivably have different values for different brands of the product (that is, some sellers' products may have more s than others'). Those who prefer more concreteness can think of s as standing for "safety," as measured by (say) the probability that the product will not have any dangerous defects. But the substance of my analysis would still apply if s represented some other, non-safety attribute that consumers valued—"speed," in the case of a wireless router; or (the absence of) "side-effects" for a drug.

Rather than assuming perfect information, I instead assume that every consumer has some estimate—possibly an inaccurate one—of each seller's level of s . I will refer to that estimate as r (for "reputation"). A

market with *perfect* information is thus a market where r is always and everywhere equal to s . If some buyers instead overestimated the product's safety, that would imply $r > s$ for those buyers. If buyers underestimated the product's safety (or overestimated its risks), that would mean $r < s$ for those buyers.¹⁶

Viewed in these terms, it should be clear that while there is only one way to have a market with perfect information ($r = s$ at all times and for all consumers), there is an almost infinite number of ways for a market to fall short of perfection, by having one or more consumers with $r \neq s$. As Tolstoy might have said (but didn't), "markets with perfect information are all alike, but every imperfect market is imperfect in its own way."¹⁷ This diversity of possible imperfections poses a challenge to economists, for there is no theoretical consensus (much less conclusive empirical data) on just how closely consumers' beliefs (r) track the true level of any attribute (s).¹⁸

The problem becomes even more complex once we introduce disclosures into the picture. To see this, imagine a spectrum of disclosures, running from those that disclose very little to those that disclose a lot. Let d stand for the quantity of information that sellers currently disclose (measured any way you like). We now have two questions of interest: (1) How, if at all, are buyers' estimates of r influenced by changes in the true level of safety, s ? And (2) how, if at all, are buyers' estimates of r influenced by changes in the quantity of disclosure, d ?

Unfortunately, neither question can be answered purely as a matter of theory.¹⁹ As a result, in this Article I make no assumptions about how consumer beliefs get formed, or about the exact way in which those beliefs are influenced by either s or d . Instead, I work backwards from

16. In standard economic notation, s and r would each be written with subscripts or superscripts—say, r_{ij} —to designate the belief that some particular consumer i held about the safety produced by some particular seller j . For ease of reading, I have omitted those superscripts here.

17. LEO TOLSTOY, ANNA KARENINA 3 (Leonard J. Kent & Nina Berberova eds., Modern Library 1993) (1918) (my paraphrase).

18. This aspect of the model draws heavily on Michael Spence, *Consumer Misperceptions, Product Failure, and Producer Liability*, 44 REV. ECON. STUD. 561 (1977).

19. For other theoretical models, each making a different assumption about consumer beliefs, see for example George A. Akerlof, *The Market for "Lemons": Quality Uncertainty and the Market Mechanism*, 84 Q.J. ECON. 488 (1970); Andrew F. Daughety & Jennifer F. Reinganum, *Products Liability, Signaling and Disclosure*, 164 J. INSTITUTIONAL & THEORETICAL ECON. 106 (2008); Xavier Gabaix & David Laibson, *Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets*, 121 Q.J. ECON. 505 (2006). Professor Akerlof's model of the "market for lemons" is probably the one that is best known in the legal academy. The Appendix discusses some of the similarities between that model and the one I use here.

the other end of the problem, by asking *what would have to be true about the influence of s or d* for disclosure to achieve either a static or dynamic goal?

First, consider the goal of static disclosures: disclosures that aim to help consumers make better choices from among the products currently available on the market. Of course, in most cases consumers have to make many *different* choices—not just whether to buy the product at all, but which brand to buy, how much of it to buy, how often (and how carefully) to use the product, and so on. To keep the discussion simple, though, I will focus on only a single consumer decision: the decision about how much of each brand to buy.

If brands differ in their level of safety, then ideally consumers should reserve the riskier brands for use only in extreme cases, when whatever benefits the product yields are great enough to outweigh the extra risk. However, if consumers have imperfect information, they may not know which brands are the riskiest, so they won't know which brands they ought to buy and consume in lower qualities. In particular, if any consumers *overestimate* a brand's safety ($r > s$), they will tend to buy more of that product than they ideally should. (In a more detailed economic model, they might also use the product too frequently, or use it less carefully than they should, or make inaccurate decisions in other respects as well.) Conversely, if a consumer *underestimates* a brand's safety ($r < s$), that consumer will be *more* cautious than he or she should be, and thus will buy too little of that brand.

The key question for static disclosures, therefore, is whether requiring more disclosure (that is, higher values of d) can succeed in nudging consumers' beliefs in the right direction. Borrowing some mathematical notation, let $\partial r / \partial d$ stand for the actual change in r (if any) that a change in d would produce. Obviously, static disclosures could not work if $\partial r / \partial d = 0$, for that would mean that consumer beliefs were not affected at all by disclosures. Instead, for static disclosures to succeed, it would have to be the case that $\partial r / \partial d < 0$ for $r > s$, or possibly $\partial r / \partial d > 0$ for $r < s$. In plain English, static success requires that a disclosure succeed in *reducing* consumers' estimates of a brand's safety if their current estimates are too high, or possibly *increasing* consumers' estimates of a brand's safety if their current estimates are too low.

While this conclusion might seem obvious, and not worth all this mathematical superstructure, my next conclusion may be less so.

Suppose we turn from static disclosures to *dynamic* disclosures, whose goal is to strengthen sellers' incentives to improve the mix of products that is available on the market. As the Appendix demonstrates, sellers' incentives to produce safer products in the first place do not

depend on whether buyers over- or underestimate the absolute levels of risks that are offered in equilibrium. In other words, sellers' incentives to improve their level of s do not depend on whether r is greater or less than s , either for any individual or for buyers generally. Instead, sellers' incentives to produce better safety (or to make any other improvements in product quality) depend on how much buyers' estimates will *change* in response to any change the seller makes.

While this conclusion, too, is rarely emphasized in law school classes, the intuition behind it is not hard to understand. In a nutshell, if most buyers would underestimate the effect of any *change* in s (that is, if a seller's reputation r changes by less than the change in that seller's true safety level s), this means that any seller who makes such a change will not get full credit for it because consumers will underestimate the true impact of the change. Obviously, this will reduce sellers' incentives to make such changes in the first place, thus leading sellers to produce lower levels of product quality.

On the other hand, if most buyers *overestimate* the effect of the change (that is, if r tends to change by *more* than s does), that will exaggerate the market rewards for such a change, and will give sellers an incentive to make even more improvements than are truly optimal.²⁰ Thus, the ideal buyer would be one who revised her estimate by exactly the same amount as the actual change in the seller's safety, as this buyer would increase her purchases of the seller's product by exactly the amount that the improvement warranted. In mathematical terms, the ideal buyer would be one for whom $\partial r/\partial s$ was always equal to 1 .

My reason for distinguishing these two aspects of buyers' information is that each is relevant to one (but only one) of the two kinds of disclosure I have distinguished. Static disclosures aim to help buyers make the best use of any product *given* the safety currently built into it, perhaps by not using that product as often or in inappropriate ways. Dynamic disclosures, on the other hand, aim to give sellers the incentive to build the right level of safety into their products in the first place. This means that static disclosures can succeed only by moving buyers' estimates, r , closer to the true value, s , so buyers no longer over- or underestimate the product's current level of risks. But dynamic disclosures can succeed only by moving $\partial r/\partial s$ closer to 1 , meaning that

20. For earlier discussions of this point, see Beales, Craswell & Salop, *supra* note 6, at 586–87; Alan Schwartz, *Proposals for Products Liability Reform: A Theoretical Synthesis*, 97 YALE L.J. 353, 374–78 (1988); and Jon D. Hanson & Douglas A. Kysar, *Taking Behavioralism Seriously: The Problem of Market Manipulation*, 74 N.Y.U. L. REV. 630, 710–14 (1999). For a formal mathematical model, see the Appendix.

they must change the rate at which consumer beliefs *change* when sellers change the true levels of risk.

Now, it may seem at first as if these two goals should usually overlap. That is, it could be the case that buyers who were accurately informed about current levels of safety would also be accurately informed about any *changes* in the current levels. In fact, the Appendix shows that this would indeed be the case in a world of *perfect* information. If buyers are at all times perfectly informed about the absolute level of any product's risks, for these buyers r would always equal s . Moreover, it follows that these buyers will also accurately estimate any change in the level of risks, so $\partial r/\partial s$ should always equal 1. Thus, if we ever achieve absolutely perfect information, we will not have to worry about the distinction between static and dynamic goals.

All real markets, however, fall to some degree short of perfect information. Further, in markets with less than perfect information, it is possible for buyers to be well-informed about one of these dimensions, but not so well-informed about the other. For example, if buyers would under-respond to changes in the level of safety ($\partial r/\partial s < 1$), that could give sellers an incentive to produce too low a level of safety, for the reasons discussed above. If that suboptimal level of safety persisted for a long time, however, buyers might well come to expect that level of safety, so in equilibrium, r would indeed equal s . It thus is entirely possible to have $r = s$ in equilibrium, but $\partial r/\partial s < 1$ for values of s that no seller is currently offering, or values of s that are off the equilibrium path.

In addition, in some cases static and dynamic goals may do worse than merely fail to coincide. In some cases, the two goals may actually conflict, if disclosures that produce positive static effects turn out to interfere with the possible dynamic effects (or vice versa). For example, suppose that consumers underestimate the health risks of sodium in foods; and suppose that, as a result, the Food and Drug Administration (FDA) requires all food packages to carry a disclosure reading "Caution: High sodium is dangerous to your health." It is certainly possible that such a warning might have a positive static effect—perhaps a very small one—on the number of consumers who previously underestimated the risks of sodium. But it is also possible that, with every product displaying *the same* health warning, some consumers would then be less likely to be aware of any *differences* in sodium content across different foods, or across different brands of the same food. In that event, we could have a disclosure that simultaneously improved consumers' understanding of the absolute level of risks (mathematically, moving r closer to the true value of s); but also worsened consumers' sensitivity to

any changes or differences in those risks (mathematically, moving $\partial r/\partial s$ even farther below its ideal value of one).

The opposite problem is also possible, where dynamic effects can interfere with static ones. For example, if the disclosure instead required all food manufacturers to disclose their *relative* sodium content, then it might enhance the disclosure's dynamic effects by steering more business toward those sellers whose products were lowest in sodium. At the same time, though, it might also lead consumers to think that "low" sodium products were actually safer than they really are ("If everybody's making a big deal about low sodium foods, maybe it's okay for me to eat a lot of them."). In that case, any benefits that were produced by the improved dynamic effects would have to be balanced against the costs produced by making consumers' static beliefs less accurate.

IV. CRITERIA FOR SUCCESS: STATIC DISCLOSURES

Having distinguished between two of the possible goals of disclosure, we can now talk about how to tell whether disclosures succeed or fail at those goals. I begin this section with static disclosures, then turn to dynamic disclosures in Section V.

Static disclosures aim to correct buyers' information about one particular fact or feature of a product that is currently available on the market. For example, the Surgeon General's warning that appeared (by law) on U.S. cigarette packages had, as one of its goals, improving buyers' information about the health risks associated with smoking. Expressed in the notation of the preceding section, this was a case where some consumers' estimates of r were believed to be too high ($r > s$), meaning that these consumers overestimated the safety of smoking. As a consequence, one aim of the warning requirement was to reduce some consumers' estimates of r , thereby bringing r and s closer together (or bringing consumers' own beliefs closer to the truth).

Of course, there are many techniques that might be used to alter consumers' beliefs, as well as many possible legal or doctrinal bases for those techniques. Sometimes disclosure is ordered after a finding that the seller's ads would otherwise be false or misleading without the disclosure of additional facts, or after a finding that the seller's previous ads had been deceptive.²¹ In other cases, disclosure is required without

21. Disclosures in this latter category are sometimes referred to as "corrective advertising." Studies of particular instances—which typically show small but non-zero effects on consumer beliefs—include Gary M. Armstrong, Metin N. Gurol & Frederick A. Russ, *A Longitudinal Evaluation of the Listerine Corrective Advertising Campaign*, 2 J. PUB. POL'Y & MARKETING 16 (1983); Peter R. Darke, Laurence Ashworth & Robin J.B. Ritchie, *Damage from Corrective*

any finding of deception if Congress passes a statute requiring certain information to be disclosed. While these distinctions may be important legally, the empirical *effect* of any disclosure is unlikely to vary based on the legal rationale that its drafter had in mind, so I will ignore these doctrinal differences here.

A. Measuring Static Effects

There are also many different ways to look for the effects of advertising (or the effects of any other disclosure) on consumers' beliefs and actions. In most cases, though, researchers do not have data about consumers' eventual purchase decisions, so the dependent variable that is most often studied is consumers' own self-reporting about their beliefs, or their own predictions of what product they might purchase. Such predictions are not always accurate, of course, but they are usually better than nothing.

The independent variable may also cause problems, since most researchers' budgets do not allow them to fund an \$80 million advertising campaign in order to see if consumer beliefs would change. To be sure, real advertising campaigns can sometimes be used as a natural experiment, as when the American Cancer Society introduced its own anti-smoking advertisements in the 1960s. But most studies of static disclosures take the easier and less expensive approach of using laboratory experiments, where each member of a group of subjects is artificially exposed to the disclosure being studied—a television ad, a consumer brochure, or a product label—and is then asked questions to elicit his or her beliefs about the advertised product.²² At the same time, similar questions are asked of a control group of subjects who were not exposed to that disclosure. If the experiment is properly designed, it may be possible to attribute any *differences* between the two groups to the effect of the disclosure itself.

As this brief description will suggest, questions about generalizability arise at nearly every turn. For example, care must be taken to make sure that the experimental subjects are similar (in demographic characteristics, previous experience with the product, etc.) to the

Advertising: Causes and Cures, 72 J. MARKETING 81 (2008); Michael B. Mazis, FTC v. Novartis: *The Return of Corrective Advertising?*, 20 J. PUB. POL'Y & MARKETING 114 (2001); Andrea Heintz Tangari, Jeremy Kees, J. Craig Andrews & Scot Burton, *Can Corrective Ad Statements Based on U.S. v. Philip Morris USA Inc. Affect Consumer Beliefs About Smoking?*, 29 J. PUB. POL'Y & MARKETING 153 (2010); William L. Wilkie, Dennis L. McNeill & Michael B. Mazis, *Marketing's "Scarlet Letter": The Theory and Practice of Corrective Advertising*, 48 J. MARKETING 11 (1984).

22. For an example of one such study, see *infra* Section IV.A.

consumers whom the disclosure is ultimately aimed at. It is also important that the subjects' *exposure* to the disclosure message is similar to the way they would be exposed in real life. In testing the effects of print advertisements, for example, subjects may be given a reprint of a short magazine article, which happens to have an advertisement inside the front cover of the reprint – but the subjects will not be told until later, after they have finished reading the article, that the *ad* is what they will be asked questions about. This allows subjects to either attend to the advertisement or not, just as they could if they were reading an actual magazine.

In short, there are many different studies that might be relevant to static disclosures, and questions might be raised about any or all of those studies. Still, the steady growth of this kind of study (paralleling the steady growth of disclosure requirements themselves) seems to be leading toward a rough consensus about how such studies should be conducted.²³ In this Article, rather than describing each of the many studies in more detail, I will instead offer four observations that summarize what these studies can teach us.

B. Lessons from the Data

The first three observations will not be surprising, but they are worth stating explicitly. (1) Disclosure never results in 100% of consumers holding completely accurate beliefs. (2) Some disclosures result in approximately 0% of consumers changing their beliefs. In other words, some static disclosures appear to have no effect at all on the goal they aim to achieve. (3) Most disclosures fall somewhere between these extremes—that is, most have effects that are somewhere between 0% and 100%.

As I said, none of these conclusions is surprising. It would be unreasonable to expect *any* disclosure to alter the beliefs of every consumer, for even Madison Avenue does not enjoy that much success. Nor is it surprising that some static disclosures appear to have no effect

23. See for example the exchange between the experts who testified on each side of the Federal Trade Commission's suit *In re Kraft, Inc.*, 114 F.T.C. 40 (1991); Jacob Jacoby & George J. Szybillo, *Consumer Research in FTC versus Kraft: A Case of Heads We Win, Tails You Lose*, 14 J. PUB. POL'Y & MARKETING 1 (1995); David W. Stewart, *Deception, Materiality and Survey Research: Some Lessons from Kraft*, 14 J. PUB. POL'Y & MARKETING 15 (1995); and Seymour Sudman, *When Experts Disagree: Comments on the Articles by Jacoby and Szybillo and Stewart*, 14 J. PUB. POL'Y & MARKETING 29 (1995). Other useful discussions include Debra K. Owen & Joyce E. Plyler, *The Role of Empirical Evidence in the Regulation of Advertising*, 10 J. PUB. POL'Y & MARKETING 1 (1991); and Janis K. Pappalardo, *The Role of Consumer Research in Evaluating Deception: An Economist's Perspective*, 65 ANTITRUST L.J. 793 (1997).

at all (even Madison Avenue whiffs from time to time). Some of these failures involve the disclosure of information that consumers simply have no interest in. Others may have been disclosures that were just badly designed—an outcome entirely understandable for disclosures designed by government employees, who don't get Madison Avenue salaries, and who tend to lack Madison Avenue's copy-writing skills.

Here is a fourth proposition, almost as unsurprising as the first three: (4) Most disclosures entail costs. If the disclosure is presented in dedicated advertisements of its own, part of the cost will be the cost of those ads' air time or print space. In other cases, though, especially if the disclosure is instead made to appear in the targeted company's own advertisements, there may also be an "interference cost" if the required disclosure crowds out other information that would itself be useful to consumers. Interference costs can result even if the other information is still physically present (i.e., is not literally crowded out) but is communicated less effectively now, as a result of having to compete with the disclosure for consumers' attention spans. Indeed, often the very disclosures that are most effective (in correcting the false belief we are worried about) are also the most costly (in interfering with other beliefs) as well. To turn the point around: often the best way to *avoid* undesirable interference effects is to make the disclosure so small, or so unobtrusive, that it does not produce any of its intended benefits, either.²⁴

I emphasize the possibility of both costs and benefits because both are crucial in deciding whether any static disclosure has succeeded or failed. Of course, if a static disclosure has produced literally *no* effects, by failing to change *any* consumer's belief, it should not be controversial to label that disclosure a failure. But what about all the static disclosures—in advertising cases, the vast majority—whose effects fall somewhere between 0% and 100%? If, say, only 32% of all current consumers have correct beliefs about a particular seller's level of safety (*s*); and if, after a disclosure is implemented, that number increases only to 37%, is the gain of five percentage points enough to call the disclosure a success? Or is five percentage points so small that it should really be labeled a failure?

C. Identifying the Costs and Benefits

Lest these questions seem fanciful, here is an illustration based on a

24. For further discussion of this trade-off, see the exchange between Ivan L. Preston & Jef L. Richards, *Consumer Miscomprehension and Deceptive Advertising: A Response to Professor Craswell*, 68 B.U. L. REV. 431 (1988); and Craswell, *Cost-Benefit*, *supra* note 8, at 567-72.

study by Cornelia Pechmann.²⁵ UPS had advertised its overnight delivery service by claiming (among other things) that its rate for delivery by 10:30 a.m. the next morning was three dollars less than Federal Express's rate for that service. This was perfectly true, but 37% of all subjects who were shown the UPS ad believed that UPS also had lower rates for delivery by 5:00 p.m. the next day. This belief was false: UPS charged more than FedEx did for delivery by 5:00.

Professor Pechmann then designed some hypothetical disclosure language that potentially could correct consumers' misimpression about which company was cheapest for delivery by 5:00. For instance, one version left UPS's original language intact, but added a disclaimer: "Rate comparisons refer only to prices for packages delivered by 10:30 A.M." When consumers were shown this version of the ad, only 29% formed the incorrect belief that UPS's rates were lower for delivery by any time the next day. This meant that the percentage of consumers who were incorrect on this point had been reduced by only eight percentage points. (The disclosure was in relatively small print at the bottom of the ad, which may be why the size of the reduction was small.)

If a disclosure reduces consumers' false beliefs by just eight percentage points, should that be considered a success or a failure? In answering that question, it may help to know that Professor Pechmann's test also showed a possible interference effect, for the addition of her disclosure language ("Rate comparisons refer only to prices for packages delivered by 10:30 A.M.") also reduced the percentage of consumers who *correctly* understood that UPS did indeed have the lower rate for 10:30 delivery. Specifically, 47% of her subjects formed the correct belief about 10:30 a.m. delivery when they were shown the original ad, without the disclosure language. However, only 33% of the consumers who were exposed to the disclosure ad formed the correct belief about 10:30 a.m. delivery, which is a decline of fourteen percentage points. In other words, Professor Pechmann's hypothetical disclosure reduced, by eight percentage points, the share of consumers who formed one false belief (about 5:00 p.m. delivery); but it also increased, by fourteen percentage points, the share of consumers who formed a different false belief (concerning 10:30 a.m. delivery). By what criteria, then, should we judge whether this disclosure succeeded or failed?

While there are several criteria that might be employed, my own suggestion—developed at more length elsewhere²⁶—is to employ a kind

25. Cornelia Pechmann, *Do Consumers Overtgeneralize One-Sided Comparative Price Claims, and Are More Stringent Regulations Needed?*, 33 J. MARKETING RES. 150 (1966).

26. See especially Craswell, *Cost-Benefit Analysis*, *supra* note 8; and Craswell, *Control Ads*,

of cost-benefit analysis. That is, the total number (or percentage) of consumers who form an incorrect belief is certainly one relevant factor; but so too is the seriousness of any consequences they might suffer as a result of their false belief. For example, if the best disclosure available could only give us a reduction of seven or eight percentage points in the number of consumers who held false beliefs *about some fatal side effect of a drug*, that disclosure should still be counted as a success, even if that disclosure also produced a very large increase in the share of consumers who held incorrect beliefs *about some other, less important feature* of the product. But if the numbers were reversed, a mere eight percentage-point reduction in the incidence of a false belief about some relatively unimportant product feature might not be enough to count as a success. And if that same disclosure also produced an interference effect, which increased the number of consumers who held a false belief about some other, more important feature, then the disclosure would almost surely fail in cost-benefit terms.²⁷

Of course, in closer cases, the use of cost-benefit analysis may require difficult or controversial judgments about the value of competing effects. For example, how we value any improvement (or any decline) in the accuracy of consumers' beliefs is a question that cannot be answered by mathematical calculations. Instead, it requires a fundamental value judgment about the importance of a better- or worse-informed citizenry compared to the value of other uses to which the money might otherwise have been put. Cost-benefit analysis is a useful way to highlight the importance of these value questions, but it does not always provide an uncontroversial way of answering them.

D. Do Static Disclosures Always Fail?

I will close this section, and my discussion of static effects generally, by returning to Professors Ben-Shahar and Schneider. As we saw earlier, their Article asserts (or strongly implies) that disclosures almost always "fail." Their Article also implies that they are not merely speculating about this failure, but that disclosures' failures have been confirmed by empirical evidence, time and time again.

How might this conclusion be justified? On my view, the best evidence would consist of studies of a number of actual disclosures,

supra note 8.

27. For an argument that this is how the Federal Trade Commission has often exercised its own authority over advertising disclosures—though sometimes without explicitly being aware it was doing so—see Craswell, *Cost-Benefit Analysis*, *supra* note 8.

showing (1) the number or percentage of consumers whose views were made more accurate by the disclosure, (2) the number or percentage of consumers whose views were made less accurate, and (3) the number or percentage of consumers whose views were not affected at all. Taking those numbers as a starting point, the study should then address the question of whether the disclosure's positive effects were important enough to outweigh any costs or negative effects that were present. Then, if a number of these studies all showed that the net effects were indeed negative, that would provide strong support for Professors Ben-Shahar and Schneider's pessimism.

As near as I can tell, however, Professors Ben-Shahar and Schneider do not cite even a single study that attempts to evaluate the costs and benefits in this way. To be sure, they cite a good many studies, so it is possible that I may have missed one here or there. But going by their own descriptions of the studies they rely on, few (if any) of those studies even *address* the question of whether the total cost of some disclosure does or does not outweigh any total benefits; and none of the cited studies addresses that question and concludes that the net effects were negative.

How, then, can Professors Ben-Shahar and Schneider conclude that the costs of mandatory disclosures "generally swamp its benefits"?²⁸ And how can they claim that the empirical evidence proves that disclosure "generally fails to achieve its goals"?²⁹ While the structure of their argument is not always clear, they seem to place the most weight on what can best be described as *their own theoretical prediction* about what is required for the effects of a disclosure to be positive.

Specifically, Professors Ben-Shahar and Schneider assert that no disclosure (or almost no disclosure) can be effective unless the parties involved successfully navigate "an impossibly long series of unlikely achievements," making the right decision at every step. For example, legislatures must (1) correctly identify a problem that needs a regulatory solution, and must then (2) correctly decide that mandated disclosure is the best regulatory response. Legislatures must also (3) decide just *which* information should be disclosed, as well as (4) design an effective enforcement system to make sure those disclosures are carried out.³⁰ Even then, the disclosure will still not be effective (according to Professors Ben-Shahar and Schneider) unless the firms that are required

28. Ben-Shahar & Schneider, *supra* note 1, at 651.

29. *Id.* at 652.

30. *Id.* at 679–91.

to disclose (5) correctly understand the legal requirements, figure out which information they should disclose, and in what format.³¹ Finally, disclosures will also be ineffective unless the intended audience of the disclosures successfully navigate a number of hurdles of their own. In particular, disclosures will not work (or so the argument runs) unless consumers are able to (6) locate the disclosed information, in whatever format it might be disclosed; (7) understand that information; and (8) remember that information until they are ready to make a purchase decision, or to act on the disclosure in some other way. In addition, if consumers are to make the *right* purchasing decision, they must have been able to (9) analyze the information to understand its implications for their choice, without being led astray by any heuristics and biases.³²

Viewed in terms of this framework, the claim by Professors Ben-Shahar and Schneider that empirical evidence supports their conclusions is really a claim about the empirical evidence bearing on each of these nine steps addressed individually. That is, some studies may show that consumers never even get exposed to the disclosed information (step 6), while other studies may show that consumers who see the disclosed information do not understand its implications (step 9). In this way, Professors Ben-Shahar and Schneider can (quite correctly) say that there are many, many empirical studies that cast doubt on how well each of these nine hurdles is likely to be overcome. And in their minds, casting doubt on how well any hurdle is overcome is equivalent to showing that the disclosure in question “failed,” because their premise is that disclosures cannot possibly succeed unless every hurdle is successfully navigated.

As I have said, though, this premise represents Professor Ben-Shahar and Schneider’s own theoretical prediction. That is, as far as I know, there has never been any empirical test of the claim that no disclosure (or hardly any disclosures) can succeed unless each of those nine steps is successfully surmounted. Instead, Professors Ben-Shahar and Schneider presumably regard this as a matter of simple logic that requires no defense or justification. After all, how could a disclosure law *possibly* succeed if, say, consumers didn’t understand the information that was disclosed?

What this simple logic overlooks, however, is that most disclosure

31. *Id.* at 691–704.

32. *Id.* at 704–28. My division of the various steps into nine is to some extent arbitrary, as some of them could be combined with others (leaving only seven or eight steps), while others could be broken down more finely (leaving ten, eleven, or twelve). However, nothing turns on the exact number of different steps, so I will continue to count them as nine.

laws (and most of the nine steps or hurdles) can succeed or fail as a matter of degree. As the surveys discussed earlier in this section show, many disclosure requirements improve the accuracy of some consumers' information, while having no effect (or even a negative effect) on the accuracy of other consumers. Indeed, such cases are sufficiently common, and the issues that arise in such cases are sufficiently well known, that there is even a name for that set of issues: "the *n percent* problem." Spelled out in more detail, if *n* represents the percentage of consumers whose information has been adversely affected by some disclosure, the "*n percent* problem" is a way of asking how great *n* should have to be in order for the disclosure to be judged a failure.³³

Unfortunately, this is where the simple logic of Professors Ben-Shahar and Schneider ceases to be useful, for failing to address one version of the *n percent* problem. Suppose an empirical study shows that one of the nine steps is *partially* satisfied—say, if 80% of consumers do not remember the disclosed information (step 8), but 20% of them do. Is a 20% memory rate enough to make the disclosure a success? What if the memory rate were 30%, or 40%? Once we recognize that disclosures can be partial successes or failures—in other words, once we recognize that there are plenty of numbers between 0 and 1—we can no longer use simple logic to infer what degree of failure with respect to one of the nine different steps is enough to make the disclosure as a whole unsuccessful.

Professors Ben-Shahar and Schneider are clearly aware of the possibility of intermediate results, for even if the *n percent* problem was not itself widely known, many of the studies they cite do show intermediate success rates (i.e., success rates between zero and one). For example, they cite studies showing that 90% of consumers don't understand how annual percentage rates (disclosed under the Truth in Lending Act) relate to simple annual interest rates.³⁴ Other studies that they cite conclude that only 16% of Medicare beneficiaries had adequate knowledge to choose between traditional Medicare and an HMO, while 67% "did not have a good grasp" of the differences between HMO plans and traditional fee-for-service arrangements.³⁵ Yet another study, this time dealing with food nutrition labels, found that only 22% of consumers could correctly calculate the net carbohydrates in two slices

33. For a more extended discussion, including other versions of the "*n percent* problem" that are functionally similar to the one discussed here, see Ivan L. Preston, *The Definition of Deceptiveness in Advertising and Other Commercial Speech*, 39 CATH. U. L. REV. 1044–46 (1990).

34. Ben-Shahar & Schneider, *supra* note 1, at 666.

35. *Id.* at 673.

of low-carb bread.³⁶

To Professors Ben-Shahar and Schneider, it must seem obvious (so obvious as to not require any defense) that these studies all give an empirical backing to their claim that disclosures rarely if ever work. Notice, though, that these same studies could just as well be read to say that these disclosures may have *improved* the information of 10% (i.e., one minus 90%), or 16%, or 33%, or 22% of the affected consumers. Moreover, most of these studies do not themselves attempt to balance the possible offsetting effects to see if the disclosure should be labeled a net failure—and neither do Professors Ben-Shahar and Schneider. Instead, Professors Ben-Shahar and Schneider simply ignore most of the range between 0 and 1, as if the success of any disclosure (or of any of their nine intermediate steps) was always a binary yes/no matter, rather than a matter of degree. As a consequence, the lengthy parade of evidence that they cite simply does not do what they say it does, for that evidence does not yet show that *any* disclosure with mixed effects has failed.

Of course, there is a possible way to rehabilitate Professors Ben-Shahar and Schneider's use of this evidence. If we believed that the proper standard for judging any disclosure required a 100% success rate, with anything less than that being deemed a failure, Professors Ben-Shahar and Schneider would then be right to treat every intermediate case as a failure as well. As I noted earlier, if the standard for success is set high enough then every disclosure will "fail," and that may be what Professors Ben-Shahar and Schneider are implicitly (or subconsciously) doing here. But if that is in fact what they are doing, then they are using a crazy standard; one that, at the very least, would require some argument in its justification. In any event, the very possibility of such an absurd standard for success provides further support for the main point of this paper, which I identified in Section I. That is, it is always a bad idea to try to evaluate disclosures without being careful and explicit about what would count as a disclosure's success.

V. CRITERIA FOR SUCCESS: DYNAMIC DISCLOSURES

I now turn to the criteria for evaluating dynamic disclosures, or disclosures that aim to improve the average quality of the products available on the market. To see whether that goal was achieved, we need good measurements (both before and after the disclosure) of the average

36. *Id.* at 676. I address the related question of whether consumers' ability to make such calculations is even necessary for a disclosure to work in notes 50–53 and the accompanying text.

quality of products sold on the market to see whether the average quality improved. Unfortunately, though, this data is not always available. And even when good data is available, it may still be hard to tell whether the improvement in product quality (if any) was caused by the disclosure, or by some other causal factor that happened to change at the same time.

In this section of the Article, I address these and other difficulties in more detail. Section V.A tries to make the discussion more concrete by describing four industries in which dynamic disclosures have been (or could have been) implemented. Section V.B then elaborates on the difficulties scholars have had in evaluating the effects of dynamic disclosures. In the end, I conclude that at least some of the disclosures described in Section V.A probably did succeed. However, I also suggest that the evidentiary and other difficulties should make us skeptical of anyone who claims that dynamic disclosures always produce a particular result. In other words, we should be skeptical of anyone who claims that dynamic disclosures always succeed, but we should be equally skeptical of those who imply that dynamic disclosures always (or nearly always) fail.

A. Some Examples

1. Credit Contract Scores

To make the discussion less abstract, it will help to have in mind some actual examples of dynamic disclosures, before discussing how to evaluate their success. In this spirit, I begin with an example that was never actually put into effect, though it may still have something to teach us. In the 1980s, before the Federal Trade Commission (FTC) decided to prohibit certain debt collection practices, it at one point considered the feasibility of a disclosure requirement instead. For example, the FTC might simply have required each lender to disclose any terms of its standard loan contract that bore on that lender's collection practices.³⁷ Alternatively (and more ambitiously), the FTC might even have rated creditors' consumer loan contracts to come up with a numerical index for each creditor, based on factors like whether the loan contracts required debtors to pay the creditor's attorney's fees, or whether they required debtors to confess judgment in advance of any dispute. Conceivably, disclosure of each creditor's numerical score might have made it easier for borrowers to shop around for their loans, and that, in

37. FTC Trade Regulation Rule: Credit Practices, 49 Fed. Reg. 7740-01, 7745–47 (Mar. 1, 1984) (Statement of Basis and Purpose).

turn, might have led some creditors to try to improve their numerical rating.

However, the FTC decided that lenders would be unlikely to compete for business along this dimension, because of an adverse selection problem.³⁸ That is, if a favorable score succeeded in signaling to borrowers that one creditor was more lenient toward its delinquent debtors, that creditor might then attract mostly the worst credit risks: those who knew they were likely to default, and therefore had the most reason to be concerned about each creditor's debt-collection practices. In the end, the FTC decided to directly prohibit what it saw as the most egregious collection practices (in other words, direct quality regulation rather than disclosure), while requiring no disclosure at all about the collection practices that it decided not to ban.³⁹

In this example, then, there are theoretical reasons to expect that dynamic disclosures would not have succeeded in improving the collection practices of the average lender on the market. However, because this disclosure regime was never adopted, there is no way to be certain what its effects would have been.

2. *Restaurant Grades*

A second example involves restaurant health and sanitation grades in Los Angeles County and elsewhere. Like the "creditor index ratings" described above, restaurant grades attempt to combine a number of attributes into a single letter grade ranging from A to C. Proponents of the system hoped both that the grades would let consumers make a better choice from among existing restaurants (a static goal), *and* that they would encourage restaurants to compete to improve their grades (a dynamic goal). According to some measures, the Los Angeles program may have achieved that dynamic effect, producing a reduction by as much as 20% in the number of hospitalizations for food-borne illnesses.⁴⁰

Others, though, have questioned the robustness of that finding.⁴¹ More

38. *Id.* at 7747.

39. Compare Ben-Shahar & Schneider, *supra* note 1, at 681 (quoting another FTC decision as evidence that lawmakers have bought into "[t]he more-information-is-better mantra").

40. Ginger Zhe Jin & Phillip Leslie, *The Effect of Information on Product Quality: Evidence from Restaurant Hygiene Grade Cards*, 118 Q.J. ECON. 409 (2003). Professors Ben-Shahar & Schneider discuss this example briefly and interpret it as a case of disclosure succeeding at its goal. Ben-Shahar & Schneider, *supra* note 1, at 743.

41. Clifford Winston, *The Efficacy of Information Policy: A Review of Archon Fung, Mary Graham, and David Weil's "Full Disclosure: The Perils and Promise of Transparency,"* 46 J. ECON.

recently, Daniel Ho examined the effects of restaurant grading programs in New York, San Diego, and several other cities, finding basically no positive effect in any of them.⁴² In other words, the combined results of a number of studies of restaurant grades can only be described as “mixed.” Possibly this means that restaurant grades have sometimes been effective (as in Los Angeles) but not always (as in other cities). Or possibly it means that the difficulty of assessing changes in average restaurant sanitation means that, at the present state of our knowledge, we cannot yet determine where restaurant grades have been successful and where they are not. In other words, possibly the success of restaurant grades has been mixed, or possibly it is only our data that is mixed—but neither interpretation supports an assertion that restaurant grades never (or almost never) work.

3. *Automobile Fuel Economy Ratings*

My third example involves the Environmental Protection Agency’s (EPA) rating of automobile fuel economy, in terms of miles driven per gallon of gas.⁴³ This, too, is a composite statistic, reflecting a variety of different driving habits and conditions. In fact, average fuel economy did steadily improve following the introduction of the EPA ratings, which might indicate a dynamic effect of the disclosures.⁴⁴ The principal complication here is that at least some of this improvement may have been caused by other factors, including (a) direct quality regulation, which required manufacturers to improve the overall fuel economy of their fleet, and (b) the imposition of higher taxes on certain low-economy vehicles. I will return to this issue later, in Section V.B.3, when I discuss causation issues more directly.

LITERATURE 704 (2008).

42. Daniel E. Ho, *Fudging the Nudge: Information Disclosure and Restaurant Grading*, 122 YALE L.J. 574 (2012). I discuss Professor Ho’s study at more length *infra* in Section V.D.

43. For the most recent ratings, see U.S. Department of Energy and U.S. Environmental Protection Agency, Model Year 2013 Fuel Economy Guide (2013).

44. See Model Year 2013 Fuel Economy Guide (presenting figures for fuel economy in the model year 2013) cited *supra* in note 43; and U.S. Department of Energy & U.S. Environmental Protection Agency, Model Year 1999 Fuel Economy Guide (presenting figures for fuel economy in the model year 1999). For data from earlier years going all the way back to 1978, see U.S. Dep’t of Energy, *Download Fuel Economy Data*, FUELECONOMY.GOV, <http://www.fueleconomy.gov/feg/download.shtml> (last visited May 20, 2013) (providing links to data from earlier years going back to 1978).

4. The Cigarette “Tar and Nicotine Derby”

A fourth example comes from the history of cigarette advertising.⁴⁵ Beginning in the early 1950s, there were several periods when *Consumer Reports* and/or *Reader’s Digest* collected and published information on the levels of “tar” and nicotine delivered by each brand of cigarette. (The federal government did not begin its own tar and nicotine testing until 1967.) There were also some periods when cigarette manufacturers themselves were allowed to emphasize their brand’s tar and nicotine levels in their advertising, if they chose to; and other periods when this was forbidden.

Perhaps significantly, some researchers have found that average tar and nicotine levels declined rapidly during those periods when tar and nicotine information was readily available to consumers. By contrast, tar and nicotine levels declined much more slowly—or, in some cases, not at all—during periods when that information was not as readily available.⁴⁶ To be sure, some of these changes may have been due to other factors; and some of the disclosures might better be described as voluntary rather than mandatory. I will discuss the possible significance of this distinction in Section V.B.3.

5. Food Labeling and Salad Dressing

Finally, my last example involves nutritional labels on foods, as required by the 1990 Nutrition Labeling and Education Act (NLEA).⁴⁷ Prior to passage of the NLEA, makers of low-fat salad dressings had often disclosed their calories and fat content voluntarily, to trumpet their own brand’s advantage in that regard. However, after the NLEA required sellers of high-fat foods to disclose information about their brands, too, the market share of high-fat dressings began to decline, while the market share of low-fat dressings began to rise.⁴⁸

In other words, in this case too the disclosure seems (at first glance) to

45. For an overview of the long and colorful history of tobacco regulation, see SUSAN WAGNER, CIGARETTE COUNTRY: TOBACCO IN AMERICAN HISTORY AND POLITICS (1971). The specific disclosure issues that I discuss in the text are described in more detail in John E. Calfee, *Cigarette Advertising, Health Information, and Regulation Before 1970* (Fed. Trade Comm’n Bureau of Econ., Working Paper No. 134, 1985); John E. Calfee, *The Ghost of Cigarette Advertising Past, REGULATION*, Nov./Dec. 1986, at 35; and Robert McAliffe, *The FTC and the Effectiveness of Cigarette Advertising Regulations*, 7 J. PUB. POL’Y & MARKETING 49 (1988).

46. See for example the two studies by Calfee *supra* cited in note 45.

47. Pub. L. No. 101-535, 104 Stat. 2353 (codified as amended at 21 U.S.C. § 343-1 (2006)).

48. Alan Mathios, *The Impact of Mandatory Disclosure Laws on Product Choices: An Analysis of the Salad Dressing Market*, 43 J. L. & ECON. 651 (2000).

have produced a dynamic change, by improving the average quality of salad dressings available on the market. But to decide whether this change was an improvement, and to decide whether the disclosures actually contributed to the change, we need to look at these examples more closely.

B. Evaluating Success or Failure

1. Gathering the Data

As I am using the term, dynamic disclosures aim to improve the average quality of the products that are available in the market. This means that the data required to evaluate a dynamic disclosure will be very different than the data available for static disclosures. As Section IV.A discussed, static disclosures can be studied experimentally in a laboratory, by exposing typical consumers to the disclosure and then seeing whether their beliefs changed as a result. By contrast, laboratory measures of consumer beliefs will not normally tell us whether a disclosure produced any dynamic effects, because dynamic effects require us to know how sellers *responded* (if at all) to any change in buyers' beliefs, and to know what new equilibrium resulted from sellers' combined responses. Unfortunately, though, it is hard to simulate the behavior of entire markets in a laboratory setting, so most studies of dynamic disclosures rely on a relatively small number of "natural experiments." What researchers look for, though cannot always find, are cases where there is good data about the average level of quality that sellers produced before and after the introduction of some new disclosure. If sellers' average quality did in fact improve during that period, we can then proceed to try to determine how much of that improvement (if any) was caused by the newly-introduced disclosure.

Unfortunately, if researchers fail to distinguish between dynamic and static effects, they will almost surely fail to gather the right kind of data for the kind of effect they are evaluating. A convenient example can be found in Professors Ben-Shahar and Schneider's comments on the food nutrition labels that were introduced by the NLEA.⁴⁹ Recall that when Professor Alan Mathios studied the effects of that disclosure, he found that the market shares of high-fat salad dressing declined (and the market shares of low-fat dressings rose) after the introduction of nutritional labels. In other words, Professor Mathios found some evidence of a dynamic effect, in which the average fat content of all

49. This disclosure program was described in the text *supra* at note 48.

salad dressings on the market improved.

In response, Professors Ben-Shahar and Schneider argue that the evidence is actually “mixed.”⁵⁰ They do not, however, mean that the data is mixed with respect to *high-fat dressings’ market shares*, which was the dynamic effect that Professor Mathios was studying. Instead, what Professors Ben-Shahar and Schneider actually argue (though apparently without realizing it) is that there is data that might be relevant to showing that NLEA food labels may not have succeeded at a *static* goal. Here is their argument, which is worth quoting at length:

Another study discovered that people infrequently consult nutrition labels, and that those who do often find it difficult to comprehend and use the information provided. Subjects in the study particularly struggled when trying to gauge whether nutrient contents comprised a “low, medium, or high amount.” Likewise, a review of 103 studies “found that although some consumers could understand some of the information on nutrition labeling, in general they reported finding nutrition labeling confusing, especially the use of some technical and numerical information.” . . . So, for example, when 200 patients were asked twenty-four questions about actual labels, only 22% of them “could determine the amount of net carbohydrates in two slices of low-carb bread, and only 23% could determine the amount of net carbohydrates in a serving of low-carb spaghetti.”⁵¹

Now, if nutrition disclosures were intended to provide basic consumer education, by teaching consumers the rudiments of nutrition, these studies might be strong evidence that the disclosures had failed that purpose. These studies might also be evidence of failure at what I would call a static purpose—for example, if the nutrition disclosures had been intended to correct the mistaken beliefs of consumers who over- or underestimate the amount of net carbohydrates in currently available brands of low-carb bread or low-carb spaghetti.⁵² In such a case, it might be disappointing to learn that only 22% or 23% of consumers now held accurate beliefs, even after the disclosures were put into effect. Even in that case, though, if we were trying to achieve static goals, we would still need to decide whether correcting the views of 22% or 23% of

50. Ben-Shahar and Schneider, *supra* note 1, at 675.

51. *Id.* at 675–76 (footnotes omitted).

52. Note that each of the studies cited by Professors Ben-Shahar and Schneider used experimental techniques, with consumer responses to surveys as their dependent variable. As discussed in Section IV.A, this is a perfectly appropriate method for evaluating the effects of *static disclosures*.

consumers was worth the costs associated with that correction. Oddly, Professors Ben-Shahar and Schneider do not discuss this cost-benefit trade-off at all.

Be that as it may, in this section of the Article my concern is not with static goals but with the dynamic goals of disclosure—that is, with whether the disclosures improved the general nutrition of products available on the market, so that consumers had a healthier set of options from which to choose. Success or failure at that goal cannot be determined merely by asking how many consumers *understand* some topic or phrase. Instead, dynamic disclosures require us to look at how the market actually responds to consumers' perhaps imperfect understanding—that is, did the nutrition level of products available on the market actually improve?

To put the point slightly differently, it may well be true that many consumers do not understand technical nutrition terms like “net carbohydrates.” It may also be true that many consumers do not know what kind of net carbohydrate score would be high, medium, or low. But I suspect it is equally true that many consumers could not tell you what chemical compounds are measured in cigarette tar and nicotine ratings, or what an average rating might be. Nor could many automobile buyers explain exactly how the EPA’s “highway” and “city” miles-per-gallon (mpg) ratings are measured. For example, how many times must the auto being tested be stopped, and then restarted, while the test is going on? And how many of those stops should be “cold restarts,” in which the engine is allowed to cool down before the car is restarted?

My point, of course, is that consumers don’t always need to understand the underlying science in order for ratings to have a dynamic effect. Indeed, this was the entire point of distinguishing between static and dynamic effects (this paper’s very *raison d’être!*), as well as the point of the mathematics presented earlier in section III. To recapitulate that analysis: achieving dynamic success, or improving the products available on the market, does not require that consumers become more accurately informed about the *absolute* level of some product attribute or risk (in mathematical terms, it does not require $r = s$). Instead it requires only that consumers more accurately perceive any *change* in a product’s attributes (in mathematical terms, it requires $\partial r/\partial s = I$). And because dynamic and static disclosures have different requirements for success, failure at one of those goals does not necessarily imply failure at the other.

A similar failure to distinguish between static and dynamic goals

limits the brief discussion by Professors Ben-Shahar and Schneider of the possible effects of cigarette disclosures.⁵³ The authors focus on what may well be the most important effect—whether the disclosures reduced the overall quantity of smoking—and they note the very real difficulty of distinguishing the effects of disclosures from the effects of all the other changes that have affected smoking, such as public education campaigns, or taxes on cigarettes, or other anti-smoking legislation. (I return to these difficulties at more length in the following subsection.)

For now, my only point is that achieving reductions in the overall level of smoking is a static goal. It takes the safety of existing cigarettes as given, and aims only to get consumers to choose the proper level of consumption *given* the existing products' safety. A dynamic goal, by contrast, aims to improve the quality or safety of the products available on the market; in this case, by reducing tar and nicotine levels. I do not claim that dynamic effects are necessarily superior to static effects, for that is surely not the case here (after all, even low-tar cigarettes are still dangerous). My claim is merely that dynamic effects are *different* from static effects, as each depends for its success on different factors. And since Professors Ben-Shahar and Schneider did not consider the possible dynamic goals that tar and nicotine disclosures might have achieved, they cannot be said to have shown that those disclosures were a dynamic failure.

2. *Eliminating Exogenous Causes*

We have seen, then, that dynamic disclosures will usually have to be assessed using real-world data about the average quality in the market as a whole. However, once we leave the laboratory for the real world, we encounter additional problems. For one thing, real-world markets are influenced by all kinds of factors, not just disclosures. Thus, even if market data shows that average quality did indeed improve, we cannot always be sure whether that improvement was due to the disclosures, or whether it was due to some other change that was taking place at the same time.

Indeed, this difficulty is especially likely to arise when we are studying disclosure laws. At bottom, the problem is that lawmakers rarely enact disclosure requirements unless the public has already become concerned about some issue. But if the public has, in fact, become more concerned about some issue, it is always possible that any subsequent change in the products available on the market will have

53. Ben-Shahar & Schneider, *supra* note 1, at 734.

been caused by the public's heightened concerns, not by the newly-enacted disclosure. In other words, the apparent coincidence in time between the disclosure and the change in products available on the market may be merely a spurious correlation.

Still, some natural experiments are more problematic than others in this respect. In the case of salad dressing disclosures, public interest in nutrition and health grew steadily throughout the 1980s and 1990s, so if public concern alone was enough to get sellers to improve their products, we ought to have seen improvements in salad dressings even before the NLEA was passed in 1990. To an extent, this is what Professor Mathios found: high-fat salad dressings were slowly losing market share even before the NLEA went into effect.⁵⁴ However, Professor Mathios also found that the market share of high-fat salad dressings subsequently began to fall at a much steeper rate—and this steeper decline began not in 1990, when the NLEA was passed, but four years later, in 1994. This is consistent with that decline being produced by NLEA's disclosure requirements, since 1994 was the year that those requirements finally went into effect (after the FDA finished its implementing regulations). To be sure, it is always possible that there was some other exogenous spike in 1994, one that had nothing to do with the disclosure laws. So far, though, no plausible candidate for that 1994 spike has been identified.

The cigarette tar and nicotine disclosures provide another good example, both of the difficulty of testing causation, and of how those difficulties can sometimes be overcome. During the 1950s and 1960s, when tar and nicotine disclosures were introduced (and sometimes were subsequently withdrawn), there were many other changes that affected the cigarette market more profoundly than when the market for salad dressings was altered in the 1990s. In the case of cigarettes, the other influences include everything from higher taxes and better medical studies linking smoking to lung cancer, to (eventually) a massive realignment of public attitudes toward smoking in general.⁵⁵ Thus, if we knew only that average tar and nicotine levels declined during that period, it would be a real challenge to know how much (if any) of that decline was due to increased disclosure of tar and nicotine, and how much was due to these other, more fundamental changes.

In the case of tar and nicotine levels, though, we know more than that. For one thing, we know that the underlying public attitudes changed

54. Mathios, *supra* note 48.

55. See sources cited *supra* note 45.

more or less steadily throughout this period, without ever reversing themselves. Fortunately for scholars, though, the applicable disclosure rules were not at all steady. Instead, they changed several times during this period, ranging from non-mandatory disclosures in popular magazines (roughly 1950 to 1954); to no disclosure at all (1954 to 1957); to non-mandatory disclosures again, this time mostly in cigarette advertisements (1958 to 1960); back to another period of no disclosure (1960 to 1963); and eventually to mandatory disclosures of varying kinds (1964 to the present).⁵⁶ Significantly, tar and nicotine levels typically declined in the years when some form of disclosure was permitted (or required), but typically stopped declining when those disclosures were prohibited or discontinued. Given such a back-and-forth pattern, it is hard to attribute both changes—that is, the declines in tar and nicotine during the disclosure years, *and* the subsequent leveling off of those declines in years when disclosure was absent—to long-term changes in the public's attitudes toward smoking, unless we think that those long-term changes themselves reversed direction every three or four years. A more parsimonious explanation of the difference in market tar and nicotine levels (the difference between disclosure years and non-disclosure years, that is) suggests that the disclosures themselves had some causal effect.

Finally, even when the change in market behavior has clearly been influenced by factors other than disclosure, that does not mean that the disclosure had no effect at all. The automobile mpg example illustrates this point, for at least part of the increase in average fuel economy is surely due both to changes in the marketplace (i.e., fuel prices were high, so consumers had more reason to consider high-mileage cars) and to other legal regulations. In particular, the Energy Policy and Conservation Act imposed a legal requirement that auto manufacturers improve the average fuel efficiency of their entire fleet of cars,⁵⁷ and the Energy Tax Act of 1978 penalized low-mileage cars in another way, by subjecting them to higher taxes.⁵⁸ As a consequence, manufacturers already had plenty of other incentives to improve their fuel economy, so we cannot simply assume that any improvement must have been due to the disclosures of mpg ratings.

However, it would be equally unjustified to simply assume that *all* of the improvement in average fuel economy was due to these other factors,

56. See sources cited *supra* note 45.

57. 49 U.S.C. § 32902 (Supp. V. 2011).

58. 26 U.S.C. § 4064 (2006).

and *none* of it was due to the disclosures. Even if we somehow knew that the disclosures contributed only 2% of the improvements in fuel economy, that would still count as a positive contribution, perhaps enough of a contribution to prevent us from saying that these disclosures had “failed.” In other words, the test for “success” (where dynamic disclosures are concerned) cannot require that *every bit* of any subsequent quality improvement must have been caused by the disclosure alone. Instead, as long as the disclosure contributed *something* to the improvements in market quality, that should count as a positive benefit whenever net success or failure is being assessed.

Indeed, in the fuel economy example, we have relatively good evidence that the changes in average fuel efficiency were not caused *entirely* by exogenous factors. If the exogenous factors—the taxes, the federal fleet requirements, and so on—had been enough *by themselves* to determine the resulting market equilibrium, there would have been no need for car makers to spend advertising dollars to play up their own high-mileage models. After all, even advertising is a form of disclosure, in the sense that it represents an attempt (not always successful) to convey information and other messages to the target audience. To be sure, advertising is usually a voluntary rather than a legally mandated disclosure,⁵⁹ but the very fact of its being voluntary lets us draw at least a limited inference from the auto manufacturers’ own choices. That is, car manufacturers apparently believed that it was worth it (to them) to spend money on this private or voluntary disclosure, so obviously the car manufacturers did not believe that *their* disclosures inevitably failed. In other words, even when consumer decisions are already affected by lots of other, more influential factors, there can still be some causal role for changes in the mix of information that consumers are exposed to.

3. *Isolating the Effect of the Legal Requirements*

The previous subsection’s discussion of “voluntary disclosures” raises an important point, for in practice it is very common to find voluntary and mandatory disclosures coexisting in the same market. The automobile mpg ad campaigns provide one example, for manufacturers were required to disclose their mpg ratings only on their cars’ sales sticker (a mandatory disclosure).⁶⁰ In addition, however, manufacturers were free to say what they liked about fuel economy in their advertisements, and some manufacturers chose to say a good deal (a

59. I discuss this distinction at more length in the following subsection.

60. See 49 U.S.C. § 32908(b)(1) (2006).

voluntary disclosure).

A similar mixture of mandatory and voluntary disclosures was also present in both the cigarette and the food nutrition examples. The NLEA is partly mandatory, for it tells sellers what they must say on their labels, and in what format they must say it (a mandatory disclosure). Like auto manufacturers, though, sellers remain free to give their nutritional information even more publicity in their print and TV ads, and many manufacturers did exactly that (a voluntary disclosure).⁶¹ The cigarette example is more complicated, as the mandatory aspect of cigarette disclosures changed several times during the period in question, and in some years cigarette manufacturers were indeed barred from saying anything more about tar and nicotine.⁶² In other years, however, manufacturers were entirely free to say more than what was minimally required. In those years (not surprisingly), companies whose cigarettes scored low in tar and nicotine were happy to display that fact in their ads, giving it far greater prominence than was required by any government mandate.⁶³

In these cases, therefore, the coexistence of mandatory and voluntary disclosures can make it hard to know how much of any dynamic effects were produced by the mandatory disclosures, and how much were produced by the sellers' own voluntary efforts. Of course, from one standpoint it might not matter which form of disclosure was responsible for the dynamic effects: perhaps all that matters is that those dynamic effects were in fact achieved. Still, Professors Ben-Shahar and Schneider titled their Article "The Failure of *Mandatory Disclosure*" (emphasis mine). This suggests that they might be perfectly willing to concede that *voluntary* disclosures (like advertising) could be effective, while still asserting that *mandatory* disclosures almost always fail. And from a policy standpoint, if it were true that the only benefits from disclosure were all produced by voluntary disclosures, that might argue that the government should never *require* disclosures. Instead, on this view, the only sensible policy for the government to take toward disclosure would be to get out of the way, relying thereafter on advertisers' own voluntary

61. See Mathios, *supra* note 48.

62. See *supra* text accompanying notes 57–59.

63. I will mention in passing that, historically, the low-tar cigarette manufacturers (who emphasize tar and nicotine ratings in their advertising) have usually been fringe firms or new entrants who were challenging the larger and more established tobacco companies (whose advertising usually ran more to beautiful women, rugged cowboys, or cool mountain lakes). Calfee, *supra* note 45. Compare Ben-Shahar & Schneider, *supra* note 1, at 738 (predicting that the fixed costs of complying with disclosure laws will "hurt[] small companies trying to enter and compete in the market").

disclosures to give us all the benefits that disclosures can possibly provide.

There are, however, two problems with positing such a fundamental difference between voluntary and mandatory disclosures. The first problem is that Professors Ben-Shahar and Schneider have given no reason why voluntary and mandatory disclosures should be expected to differ in such a fundamental way. To the contrary, almost all of the reasons that Professors Ben-Shahar and Schneider offer when they argue that mandatory disclosures can hardly ever be effective imply with just as much force (or nearly as much force) that voluntary disclosures will not be effective, either.

For example, consider the “quantity” problem that Professors Ben-Shahar and Schneider identify as one of the main theoretical reasons why they expect mandatory disclosures to fail.⁶⁴ In their view, consumers are overwhelmed by all the disclosures they are bombarded with, to the point where most consumers would be unable to make good use of those disclosures even if they wanted to. Professors Ben-Shahar and Schneider also emphasize the cognitive limitations that consumers are subject to, which may further reduce the amount of information they can process.⁶⁵ While I agree that these limits are real, I see no reason to limit their effects to consumers’ receptiveness to *mandatory* disclosures. Rather, it seems to me that if consumers are so overwhelmed with information that they cannot usefully take in any more, that limit should reduce the effectiveness of voluntary and involuntary disclosures alike.

In any event, the second problem with the argument that “voluntary disclosures will produce all the benefits that are there to be attained” is that neither theory nor data supports this prediction. At least as a matter of economic theory, there are several reasons to think that markets for information (even more than other markets) will not always work perfectly.⁶⁶ As a matter of empirical observation, it is not hard to find cases where sellers were unable to promote some attribute effectively until a disclosure requirement was introduced, at least in a quasi-voluntary way (if not stronger). In the market for salad dressings, for instance, makers of low-fat dressings could and did promote their own dressings even before passage of the NLEA. But those promotions did not have nearly as much effect on the market shares of high-fat dressings until after the NLEA made information about the fat content of *all* brands

64. Ben-Shahar & Schneider, *supra* note 1, at 686–91.

65. *Id.* at 719–29.

66. For more extended discussions of this point, see Beales et al., *supra* note 6, at 501–03, and Fung et al., *supra* note 4, at 30–33.

publicly available for use by consumers or by competing advertisers.⁶⁷ In these cases, then, mandated and voluntary disclosures may be not so much substitutes as complements, meaning that the adoption of mandatory disclosure rules may increase (rather than decrease) the amount and the effectiveness of voluntary disclosures.

Why might the adoption of mandatory disclosures have this effect? This is one of the least-studied aspects of disclosure, so conclusive answers are unattainable. However, several possibilities have been suggested in the literature. For example, in some industries there may be no commonly-accepted method of measuring product quality, and it may be in no single seller's interest to construct one.⁶⁸ If different sellers use different measurements—for example, if different auto companies measure mpg in different ways or under different conditions—that may make it hard for consumers to make apples-to-apples comparisons when evaluating different makes of cars. In addition, the widespread use of different measurement systems may also slow the rate at which consumers become familiar with any one of those systems. But if the government were to require sellers to use one single way of measuring fuel efficiency (as the EPA has done), over time that might increase consumers' understanding of the particular metric that was adopted.

Another possibility is that, when the government requires the disclosure of information, that decision itself signals to (some) consumers that the issue is important enough to worry about, thus making sellers' later voluntary disclosures more salient to consumers than they would have been if the government had not acted. Still another possibility is that when information is disclosed voluntarily, consumers may not always trust the advertiser who provides it. Some may fear the advertiser is telling outright lies; others may simply worry that the advertiser (without actually lying) has managed to exploit imperfections in the measuring system, thus making his own brand look better than it really is.⁶⁹ In these cases, if the government itself is the one who devises the measurement system and who conducts the tests, some consumers may be more willing to rely on the resulting information.

More generally, once we consider all the different roles that the

67. Mathios, *supra* note 48.

68. For more extensive discussions of the "public good" aspect of many disclosures, see Beales et al., *supra* note 6, at 503–05, and Fung et al., *supra* note 4, at 30–33.

69. For a recent allegation that Hyundai has in fact misstated its mpg ratings, see Harvey Rosenfield, *Hyundai – False Advertising of Miles Per Gallon (MPG)*, CONSUMER WATCHDOG, <http://www.consumerwatchdog.org/case/hyundai-false-advertising-miles-gallon-mpg> (last visited May 10, 2013).

government might play, we can see that the simple distinction between voluntary and mandatory disclosures does not exhaust the spectrum of possibilities.⁷⁰ In some cases, for example, the government might announce a standardized method for measuring fuel economy (thus solving any difficulties in coordinating on a single measure), without actually requiring anybody to disclose that measurement (thus leaving it to firms who score well on that measure to handle all the publicizing). Other possibilities involve triggered requirements—for example, if sellers are not required to say anything at all about fuel economy; but if they do make fuel economy claims, they must use the government's approved method for measuring fuel economy. Still another possibility would be to mandate disclosure during a short initial period (while sellers and consumers got accustomed to that information), but then turn it into a purely voluntary disclosure at the end of two or three years.

As should be apparent, none of these examples would be entirely mandatory (except perhaps for a two- or three-year period); but they are not completely voluntary, either. In effect, these examples fall into a third category—government-aided disclosures (GADs)—that deserves study of its own. As Professors Ben-Shahar and Schneider do not discuss these, I do not know whether they would say that GADs would fall within the scope of their argument that “mandatory” disclosures almost inevitably fail.

For now, I will merely note that the pessimism toward disclosures that Professors Ben-Shahar and Schneider display may in part arise from this same error: the error of excluding government-assisted disclosures from their analysis. After all, most disclosure laws operate as floors rather than as ceilings, in that they allow any seller who chooses to give the disclosed information even greater prominence in their ads, so most disclosures have at least the possibility of becoming partly voluntary. To be sure, in some cases no seller has any incentive to give the information even greater publicity; so in those cases sellers will do no more than the bare minimum that the law requires. In other cases, however, some sellers will have an incentive to publicize the information even more than the law requires, usually with more effective materials than government agencies could produce, and a bigger advertising budget than government agencies typically command. Thus, when agencies design a dynamic disclosure, they should usually *hope* that the information that they mandate is so attractive to consumers (and to sellers whose

70. For further discussions of each of these possibilities, see Beales et al., *supra* note 6, at 521–31.

products score well on it) that it becomes, in effect, a mere government-aided disclosure.

Of course, at the design stage, regulators can rarely be sure they can achieve this effect. Just as Broadway producers cannot be sure which of their plays will become hits, the designers of disclosure laws cannot be sure which of theirs will achieve GAD status. I think it is clear, though, that the disclosures most likely to reach GAD status are the ones that really mean something to consumers. Advertisers rarely waste their own money on promotions that consumers will ignore or will not understand. If I am right in this, then when Professors Ben-Shahar and Schneider exclude from their analysis all disclosures that are merely government aided, they systematically exclude the very disclosures that are most likely to make a dynamic impact. This, in turn, leads them into what may be a simple error about the direction of causation. That is, it may be approximately true to conclude, as they do, that mandatory disclosures are rarely successful. But it is equally true, and produces perhaps a more accurate picture, if we run the causation in the other direction: disclosures that are successful rarely remain entirely mandatory.

4. *Designing More Effective Disclosures*

Finally, the exclusion of all voluntary and government-aided disclosures from their analysis may also lead Professors Ben-Shahar and Schneider to overlook some of the design features that can contribute to a disclosure's success or failure. For example, when Professors Ben-Shahar and Schneider acknowledge the "moderate successes" (their words) enjoyed by some disclosures that employ rating systems, they conclude that the main prerequisite for such success is the *simplicity* of the rating system.⁷¹ Of course, this recommendation fits well with their earlier diagnosis of the problems that most disclosures face. If consumers are being bombarded with information, and are not even capable of processing most of it, then presumably only the simplest pieces of information can ever make it through the bombardment to affect consumers' decisions.

While I agree with Professors Ben-Shahar and Schneider that simplicity is important, I think this overlooks another important factor, which is whether the information being disclosed is different for different competing firms, or whether all firms are instead required to make the same disclosure. The five examples that I discussed earlier—

71. Ben-Shahar & Schneider, *supra* note 1, at 678–79, 743–45. A similar conclusion was reached by Fung et al., *supra* note 4.

Los Angeles restaurant grades, cigarette tar and nicotine ratings, automobile mpg ratings, salad dressing fat labels, and the FTC's proposed rating of credit contracts—all disclosed information that varies across different firms, depending on their product's fat content or mpg rating or other feature. By contrast, the standard warning that appeared on all cigarette packages beginning in 1970 ("Warning: The Surgeon General Has Determined that Cigarette Smoking Is Dangerous to Your Health") is a uniform disclosure in that it does not vary across firms, since all cigarette companies must disclose the exact same warning.

This distinction is significant because competing sellers, if they have any incentive at all to embrace a disclosure and make it partly voluntary, are more likely to do that if their disclosure is different from their rivals'. For instance, no cigarette company ever went out of its way to publicize the Surgeon General's warning in its own advertising campaigns—and why should they, since that warning gave consumers no basis for choosing their brand of cigarettes rather than some other? However, some cigarette companies did go out of their way to emphasize tar and nicotine levels in their advertising campaigns.

This quasi-voluntary embrace of tar and nicotine information cannot be attributed to its greater simplicity, for if there is any difference at all, the Surgeon General's warning is probably the simpler of the two pieces of information. (Compare "Cigarettes are bad for your health" with "Cigarettes contain something called tar and nicotine, and our brand contains them in the following quantities.") This suggests that the key fact in this case was not just the disclosure's simplicity, but the fact that tar and nicotine disclosures were not *uniform* across different sellers (while the Surgeon General's warning was entirely uniform, down to its smallest details). As a result, tar and nicotine levels—unlike the Surgeon General's warning—gave sellers a basis on which to differentiate their brands from their rivals', and thereby gave sellers some reason to spend their own money to publicize the new rating system.

5. *Identifying the Costs and Benefits*

Finally, as is often the case, we are left with costs and benefits. The preceding subsections have argued that it is possible for dynamic disclosures to produce benefits, but they can also produce costs. For example if the EPA's method of calculating mpg is actually inferior to some other method, then making that method mandatory (or even favoring that method in one of the "softer" ways discussed above) could entrench an inferior method of calculation, possibly making consumers' information worse rather than better.

Successful disclosures could also distort consumer choices in another,

more subtle way, if (say) the easy availability of information about fuel economy led consumers to place too much weight on that decision and too little weight on other, arguably more important factors, such as crashworthiness or lifetime maintenance costs. Viewed broadly, this is another version of the “teaching to the test” problem, if cars that score well on mpg are not necessarily the cars that are best on other, less easily observable dimensions.⁷²

Indeed, this was the ostensible basis for *restricting* sellers’ use of tar and nicotine information during two periods in the 1950s. The concern was that consumers might pay so much attention to the relative tar and nicotine ratings (since those were what low-tar sellers emphasized in their advertising) that they might come to see that dimension as more important than it truly was, by not understanding that even low-tar cigarettes are still not very safe.⁷³ In the terms used in this Article, this would be an example of information that might have positive dynamic effects (if it brought down average tar and nicotine levels on the market) but negative static effects (if it led consumers to underestimate the riskiness of those cigarettes that remained on the market).

In short, just as we saw earlier in connection with static disclosures, dynamic disclosures too can sometimes produce conflicting effects. To put it another way, both static and dynamic effects are usually a matter of degree, for some effects will be large while others may be small. As a consequence, it will often be impossible to label any disclosure a success (or a failure) without making the value judgments needed to say whether one set of effects outweighs the others. Hopefully, Professors Ben-Shahar and Schneider’s future work will contribute to our understanding of these trade-offs.

VI. CRITERIA FOR SUCCESS: DISCLOSURE VERSUS QUALITY REGULATION

So far, I have argued that any criteria for evaluating the success or failure of disclosures should recognize the different effects that might be produced by static versus dynamic disclosures. In this final section, I argue that our criteria for success or failure may also have to be adjusted depending on what the best alternative is to the disclosure that is being evaluated.

72. *See supra* text accompanying note 43.

73. Calfee, *supra* note 45. An alternative hypothesis is that the established cigarette companies were beginning to lose market share to the smaller, low-tar companies, so the established firms were happy to have comparative tar and nicotine advertising prohibited.

More specifically, sometimes when we evaluate a disclosure, we are implicitly asking: “Is this disclosure effective enough that we would be better off with the disclosure than without it—better off, that is, than we would be with no disclosure requirement at all?” At other times, though, we may implicitly be asking a different question: “is this disclosure effective enough that there is no need to regulate product quality directly in this market—that is, would we be better off (a) with the disclosure requirement plus no quality regulation, or (b) without the disclosure requirement but with quality regulation?” As I discuss below, each of these questions implies a somewhat different criterion for saying whether a disclosure has “failed.”

A. Restaurant Grades Revisited

To illustrate, consider again the system of restaurant hygiene grades adopted in Los Angeles County and many other jurisdictions.⁷⁴ One study of their effects in Los Angeles County suggested that the grades had produced large and positive dynamic effects by increasing average hygiene levels at restaurants throughout the county. On the other hand, Daniel Ho’s thorough study of restaurant grading in New York, San Diego, and several other cities, found that those regimes had produced little or no dynamic effect.⁷⁵

Now, it is *possible* that these varied effects might result from problems with the available data. Inadequate restaurant hygiene is most likely to make itself felt in the incidence of food poisoning or other food-borne disease, but it is hard to get good data on the incidence of those health problems, and even harder to link that data to any particular restaurant whose food might have caused the problem. As a consequence, analysts have had to be creative in finding proxies to estimate the health effects, like the total number of hospital admissions for food poisoning (even though many of those will have been caused by home-cooked foods),⁷⁶ or the number of Google searches for terms and phrases like “food poisoning” or “listeria.”⁷⁷

In New York and San Diego, though, Professor Ho found that the problems with food grading ran far deeper than a lack of data. In San Diego, for example, the cut-offs for the different letter grades had been drawn in such a way that 99% of all restaurants got A’s, thus depriving

74. See *supra* section V.A.2.

75. Ho, *supra* note 42.

76. Jin & Leslie, *supra* note 40, at 435–36.

77. Ho, *supra* note 42, at 645.

that grade of any useful comparative content. And in New York, Professor Ho found a perfect storm of problems, including (among others) these:

- (1) Inspectors are not paid well, so turnover is high, and reports of bribery are not uncommon.
- (2) Inspectors do not specialize in restaurant scoring, but are drawn from a pool of general-purpose inspectors.
- (3) The forms used by inspectors are complex, and inspectors have a great deal of discretion about how to score any given violation.
- (4) The scoring categories listed on the form focus on problems other than the ones that modern health experts consider most serious.
- (5) Restaurant hygiene is volatile, so a restaurant that is perfectly clean on one day may be dangerously unsafe a few days later (or vice versa).
- (6) A restaurant that (initially) receives a low grade can apply to be retested, and only the later grade must be disclosed to the public.
- (7) Inspectors spend quite a lot of their time re-inspecting “B” restaurants (who hope re-inspection will raise their grades to an “A”); and relatively little time re-inspecting “C” restaurants, who are presumably the least hygienic of the bunch.

In light of all these problems, it is not surprising that New York’s restaurant grading program does not appear to have produced any positive dynamic effects. If the grades bear only a slight correlation, or possibly no correlation at all, to actual restaurant food safety; and if “teaching to the test,” in this context, means bribing the local restaurant inspector, it would be both surprising and discouraging if consumers *did* pay any attention to New York’s restaurant grades.⁷⁸ Thus, by any of the criteria I have discussed so far, New York’s restaurant grades have to be labeled a dynamic failure.

1. Disclosure Versus No Regulation

However, even this conclusion (that New York’s restaurant grades have failed) may have to be modulated by the question to which that conclusion is offered as an answer. Suppose, for example, that New York’s city council is considering repealing the restaurant grading system; and suppose (for now) that if the repeal passes, the result will be

78. Note, too, that restaurants that get a good grade in New York do not generally build advertising campaigns around that fact, the way sellers of low-tar cigarettes or high-mileage automobiles have done. In terms of the discussion *supra* in section V.A.4, New York’s restaurant grades have failed to achieve GAD status.

no disclosure requirement at all, and no other change in any of the other restaurant licensing laws. In other words, this is not a case where the council is considering replacing the letter grades with some other form of disclosure; nor are they considering replacing their failed disclosure policy with much stricter licensing laws, or with any other form of direct safety regulation. Instead, in this case the council's choice is "restaurant grades or nothing." If the choice is narrowed to those two options, clearly Professor Ho's study would support those who preferred "nothing," since New York does not seem to be getting any benefits at all from its restaurant grades. This, then, is one criterion by which New York's grades might be deemed a failure. On this view, they are not merely bad; they are actually worse than no regulation at all.

2. *Disclosure Versus Direct Quality Regulation*

However, now suppose that the city council is considering a different proposal. Rather than continuing to operate an ineffective disclosure system, one council member has proposed converting it to a form of direct regulation, under which all restaurants that failed to achieve the highest possible grade would be shut down. In other words, rather than relying on a low grade to shame a restaurant into improving its hygiene; and rather than hoping such grades would lead consumers voluntarily to other, less risky restaurants; this proposal would take that decision out of consumers' hands (and out of the hands of individual restaurateurs) by prohibiting *any* consumer from eating at a restaurant whose safety and hygiene were judged too low. In effect, this is the restaurant equivalent of what the FTC did with respect to credit contracts, when it decided that no disclosure was likely to achieve the desired dynamic effects, so it prohibited some contracts as a matter of law.

Now suppose that, at a meeting where this proposal is being discussed, one of the council members argues that disclosure is generally superior to direct regulation, because it allows people to choose where to eat. In response, another council member argues that disclosure remedies often fail, citing either Professors Ben-Shahar and Schneider's work on disclosure in general, or Professor Ho's work on New York restaurant grades in particular. Clearly, these two council members are disagreeing about the effectiveness of disclosures. But is the second council member correct, as Professors Ben-Shahar and Schneider believe? That is, do the problems Professor Ho identified mean that restaurant grades must have "failed" by the criteria that are relevant here, when the choice is between a problematic disclosure system and direct regulation of restaurant hygiene?

It might seem that the answer is obvious. If the disclosure system has

failed so badly that it is worse than no regulation at all (as we concluded in the previous subsection), then it might seem that disclosure must also be worse than the proposed licensing system that is on the table now. However, that conclusion involves a questionable leap, for it ignores the possibility that the proposed licensing system might also be worse than no regulation at all.

To put the point more clearly, most of the problems that Professor Ho identified do more than limit the gains we can expect from the letter grade disclosure system (though they certainly do that). In addition, most of those problems also limit the gains we could expect from direct quality regulation, such as the proposal to shut down any restaurant that fails to achieve the highest grade. After all, if inspectors are corrupt or poorly trained, that will cause some good restaurants to be shut down under the new licensing system, while allowing bad restaurants to survive (if they pay the right bribe). Shutting down low-scoring restaurants is also problematic if the criteria on which the scores are based do not correspond to the problems that most modern health experts consider serious, or if the relation between inspector scores and actual health risks is essentially random. True, in theory it might be possible to improve New York's inspection system, to reduce or eliminate some of these problems. But if that is the case, it should also be possible to adopt those same improvements while retaining the letter grade disclosure system, now slightly less problematic because improvements have been made.

In short, the problems that Professor Ho identified are not problems that plague disclosure laws alone. Instead, most are problems that would plague any form of regulation, whether or not any disclosure is involved. As a result, in cases where the relevant choice is not "disclosure or nothing" but rather "disclosure or direct quality regulation," the problems identified by Professor Ho do not necessarily tip the latter choice in either direction. By what criteria, then, should we say *in this context* that New York's restaurant letter grades have failed? Must we say that we are using the term "failed" in a descriptive rather than a normative sense, which is not meant to imply anything about whether a "failed" system might nevertheless be the best one to vote for (depending on the other choices available)?

3. *Asking the Right Question*

A similar point can be made about many of the factors that Professors Ben-Shahar and Schneider point to when they explain why (in their view) most disclosure laws are likely to fail. As their Article is already a long one, the authors understandably do not try to evaluate any of the

possible regulatory alternatives to each of the disclosures they discuss.⁷⁹ The consequence, though, is that many of the factors they identify are similar to those identified by Professor Ho, in the sense that they increase the difficulty of direct regulation by just as much as they increase the difficulty of disclosures. This gives some of their discussion an unintended “Tea Party” flavor, as if what they were really trying to show was the failure of *all* regulation, not just disclosure requirements.

For example, Professors Ben-Shahar and Schneider note that disclosure laws are often adopted because of “trouble stories,” or publicized disasters where someone dies or loses their home because of some choice that might have been prevented by better disclosures.⁸⁰ In these cases, the public often demands that lawmakers “do something” to prevent a recurrence, which may lead to the hasty adoption of an ill-considered disclosure law. But while I agree that this can indeed result in too many disclosures (or in not enough effective ones), it seems to me that much the same problem also plagues direct regulation, where lawmakers are also pressured to act by highly publicized but unrepresentative trouble stories. Thus, if this is a problem (and it surely is), it is not a problem that changes the *relative* attractiveness of disclosure laws by making them any *less* likely to be successful than direct regulation.

Or consider one of Professors Ben-Shahar and Schneider’s reasons for dismissing the possibility that food nutrition labels had contributed to the decline in market share of high-fat salad dressings. After making various other arguments, Professors Ben-Shahar and Schneider conclude by noting the “sobering” fact that, even when people reduce their consumption of fat from some sources, they often compensate by increasing their fat intake from other foods.⁸¹ This is a valid concern, for it suggests that the net benefits may perhaps be smaller than would be shown by evaluating the changes in fat content from any one product alone.⁸²

My point, of course, is that this same fact also suggests that the net benefits of direct quality regulation may also be smaller than would otherwise be supposed. That is, if Congress or the FDA, perhaps after being persuaded that disclosures rarely if ever work, were to abandon

79. Ben-Shahar & Schneider, *supra* note 1, at 742 (“That is too large a question to squeeze into the compass of this Article.”).

80. *Id.* at 679–80.

81. *Id.* at 677.

82. I will set aside the possibility that some people might prefer to consume whatever fat they permit themselves in the form of, say, ice cream rather than salad oil.

nutrition labels and simply *require* salad dressing manufacturers to lower their fat content, presumably consumers would still be just as inclined to compensate by increasing their intake of other fatty foods. If so, then this means consumers' compensation habits will reduce the benefits of disclosures and direct regulation equally, without in any way increasing the *relative* drawbacks of disclosure.

Indeed, much the same could be said of the other hurdles that Professors Schneider and Ben-Shahar identify:

For mandated disclosure to work, lawmakers must succeed at several tasks. First, they must correctly identify a problem that needs a regulatory solution. Second, they must correctly decide that mandated disclosure is the appropriate regulatory method. Third, they must correctly decide what disclosure to mandate. Fourth, they must correctly and comprehensively articulate the standard of disclosure. Each step is problematic; tackling all four successfully is uncommon.⁸³

While each of these tasks is indeed difficult, they do not seem notably harder than the tasks that any direct regulation must accomplish. For a lawmaker to get direct quality regulation right, she must first identify a problem that needs a regulatory solution, then correctly decide that quality regulation is the appropriate regulatory method. After that, the lawmaker must decide just what quality the regulation ought to require, and then articulate a comprehensive standard defining that quality. While it may indeed be hard for lawmakers to successfully navigate all four of these tasks, I see no reason why that combination of tasks should be any harder (or any easier) when disclosures are involved than when they are not.

In short, there are some contexts where our (often implicit) comparison is not between disclosure and no regulation at all, but between disclosure and some other form of direct quality regulation. In those contexts, the most important aspects of disclosure laws may be those that *differentially* affect disclosure laws and direct quality regulations, thus altering the balance between those two regulatory strategies. Some of the factors that Professors Ben-Shahar and Schneider mention do indeed have such a differential effect—for example, information overload is likely to limit the effectiveness of disclosure laws by more than they limit the effectiveness of direct regulation, since direct quality regulation (by definition) spares consumers from processing additional information. Hopefully, Professors Ben-Shahar

83. Ben-Shahar & Schneider, *supra* note 1, at 679.

and Schneider's future work will illuminate this aspect of the disclosure trade-off as well.

CONCLUSION

H.L. Mencken once asked his readers to imagine:

[H]ow distressing fatherhood would become if prospective fathers were all taught that the human infant radiates an aroma like the rose—if the truth came constantly as a surprise! Each fresh victim of the deception would feel that he had been basely swindled—that his own child was somehow bogus. Not infrequently, I suppose, he would be tempted to make away with it in some quiet manner, and have another—only to be shocked again.⁸⁴

Mencken's target when he wrote this was politicians, and the expectation of many Americans that their politicians should indeed smell like roses—"only to be shocked again" each time a new one was elected. I quote him here only because what he says about Americans' attitude toward politicians may also be true of their attitudes toward disclosure laws.

That is, people who expect disclosure laws to solve almost every problem—quickly, easily, and with very little cost—are doomed to have their expectations crushed. The truth is that a successful disclosure law, like any other form of regulation, requires lots of difficult work, including a careful specification of the purposes the disclosure might serve, and careful assessment of all the various effects (both good and bad) that most disclosures produce. As a result, those who expect disclosures to succeed easily will usually be doomed to disappointment. They may feel that they have been "basely swindled," and they may even be tempted to do away with disclosure laws entirely.

However, one lesson to be drawn from Mencken is the possibility that sometimes our disappointment in something is not the fault of the object's inherent unattractiveness, but is rather the fault of the standards by which we judge those objects. If we expect good babies to smell like roses, then every baby will be judged a failure; and if we expect disclosure laws to make people perfectly informed, then every disclosure law will fail. The solution, though, is not to do away with disclosures (or with babies), but rather to revise the standards by which we judge them, employing more realistic expectations. Indeed, as I have argued here,

84. H.L. Mencken, *The Politician*, in PREJUDICES: FOURTH, FIFTH, AND SIXTH SERIES 72, 75 (1955).

different disclosure laws are likely to have *different* purposes, which will place different demands on anyone trying to evaluate the disclosure's success or failure.

For example, disclosures will succeed at some static purpose if, but only to the extent that, they can reduce the number of consumers who misperceive the benefits and risks of products presently available on the market. By contrast, disclosures with a dynamic purpose will succeed only if they increase the number of consumers who accurately perceive the differences between those products that are currently available, and those products that *might* be available if market conditions were to improve.

In short, understanding the different possible goals of disclosure is essential to any decision about which disclosures "work." Professors Ben-Shahar and Schneider have already illuminated many other aspects of disclosure laws. If they can one day articulate (and defend) their own criteria for what would count as a "success," they will advance our understanding by even more than they already have.

MATHEMATICAL APPENDIX

As noted in the text, this model draws heavily on one used by Michael Spence.⁸⁵ Sellers and buyers are both risk-neutral, and sellers produce a single product that carries some risk of a defect or injury. Let s represent the probability of *no* defect in any given unit of the product ($0 < s < 1$), so the probability of a defect is $1 - s$. I will refer to s as "safety," though it could also refer to non-safety attributes like performance or durability. For simplicity, I let s depend only on sellers' decisions about how to design or build the product, so consumers' care decisions can be eliminated from the analysis. Consumers can, however, affect the total losses (though not the per-unit losses) by their decision of how much of the product to purchase.

Sellers choose both the level of safety they build into their products (s) and the per-unit price they charge (p). Sellers' per-unit production costs do not depend on how much they produce (constant marginal costs), but production costs do depend on the product's safety (safer

85. Michael Spence, *Consumer Misperceptions, Product Failure, and Producer Liability*, 44 REV. ECON. STUD. 561 (1977). It is also similar to one part of Steven Shavell's model of tort rules—specifically, when he models the case where (1) sellers are not liable for injuries their customers suffer from product defects; (2) customers are less than perfectly informed about the risk of a defect; and (3) customers cannot themselves affect the probability of a defect, except by altering the quantity of goods that they purchase. Steven Shavell, *Strict Liability Versus Negligence*, 9 J. LEGAL STUD. 1, 12–17 (1980).

products are more expensive to produce). Formally, let sellers' per-unit costs be given by $c = C(s)$, with $C'(s)$ and $C''(s)$ both positive to reflect positive but decreasing returns from expenditures on safety. I assume a competitive market (sellers are price-takers, and consumers have perfect information about each seller's price). Thus, the equilibrium price will be competed down until it just barely covers sellers' production costs, using a circumflex to indicate the level of safety \hat{s} (or the level of any other variable) that sellers choose in equilibrium:

$$(1) \ p - C(\hat{s}) = 0$$

Consumers respond to sellers' choices by choosing the quantity they purchase (q). Consumers then learn whether any of the units that they purchased are defective. If they are, the consumer who purchased that unit suffers a loss L , which is the same for every consumer. (For convenience, I assume consumers are also identical in all other respects, so it is sufficient to analyze a single representative consumer). In consumers' utility functions, the loss L is additively separable from the benefits they get if there is no defect, so those benefits equal the area under the inverse demand curve $p = P(q)$.

These assumptions let us express overall welfare W as a function of s and q :

$$(2) \ W(s, q) = \int_0^q P(z) dz - qC(s) - q(1-s)L$$

The integral on the right side of the equation is a consumer's expected benefit from consuming q units; while the second term, $qC(s)$, is the total cost of producing those units. The final term is the expected loss from defects, given the number of units consumed. Differentiating over s and q yields the following first-order conditions, using asterisks to denote socially optimal values:

$$(3) \ C'(s^*) = L$$

$$(4) \ P(q^*) = C(s^*) + (1 - s^*)L$$

These conditions both have familiar interpretations. Equation (3) says that sellers, ideally, should raise the level of safety s until the marginal cost of further expenditures on safety just equals the marginal gain from those expenditures. That marginal gain is simply the loss, L , whose probability is reduced by further expenditures on safety. Meanwhile, equation (4) says that consumers, ideally, should increase their purchases of this product until the marginal benefit from an additional purchase just equals the full cost of that purchase. That full cost includes the expected losses from defects, $(1 - s)L$, as well as the direct costs of production, $C(s)$.

In perfect information models, consumers know the exact value of every relevant variable. I follow those models in two respects: by assuming that consumers know the exact price they pay (p) and the

quantity they purchase (q). To introduce imperfect information, I assume that consumers do not know the probability that the product will be non-defective (s). Instead, they know only some estimate r (for “reputation”), which may or may not be accurate.

I will express r as a function $r = R(s, d)$, where d is the level of sellers’ disclosure (measured any way you like). However, this function is merely a black box: initially, I place no restrictions on the form this function can take, or on the exact way in which r varies with s or d . Indeed, I mean to include the possibility that $R_1(s, d) = R_2(s, d) = 0$, meaning that consumer beliefs about safety are completely oblivious to both the true level of safety s , and the amount of seller disclosures d . I also include the possibility that buyers are always *perfectly* informed, meaning $R(s, d) = s$ for all possible values of s and d .

If consumers take r (rather than s) to be the probability of no defect, they will estimate their net benefits V as:

$$(5) \quad V(s, q, d) = \int_0^q P(z) dz - qC(s) - q(1 - R(s, d))L$$

This is analogous to equation (2) except that, in the final term on the right, consumers estimate the probability of no defect as r rather than s . Differentiating with respect to q shows that consumers will choose the quantity they purchase according to the following first-order condition, which is analogous to equation (4):

$$(6) \quad P(q) = p + (1 - R(s, d))L$$

At the same time, competition over safety should lead sellers to choose the level of s that maximizes the value that buyers *think* they are getting, keeping in mind that imperfectly informed buyers may be wrong about what that value is. Formally, sellers choose s to maximize equation (4) subject to the constraints represented by equations (5) and (1).

Differentiation now yields the following, slightly different first-order conditions:

$$(7) \quad C'(\hat{s}) = R_1(\hat{s}, d)L$$

$$(8) \quad P(\hat{q}) = C(\hat{s}) + (1 - r)L$$

By comparing these with the conditions for social optimality given earlier in equations (3) and (4), we can see two distinct ways in which this equilibrium may depart from optimality.

First, if $r > s$, a comparison of equations (4) and (8) indicates that consumers will purchase too large a quantity of the good. Recall that both r and s refer to the probability of *no* defect, so $r > s$ means that consumers underestimate the product’s risks (or overestimate its safety). As a consequence, consumers underestimate the true costs associated with using this product, and they therefore buy too much. If the opposite held ($r < s$), consumers would overestimate the risks and would buy too

little.⁸⁶

Second, a comparison of equations (3) and (7) shows that sellers will choose too low a level of safety in equilibrium if but only if $R_I(s,d) < 1$. Recall that $r = R(s,d)$ is the function that describes how consumers' beliefs about safety, r , respond to changes in the true level of safety, s . If consumers were completely ignorant about the true level of s , we would have $R_I(s,d) = 0$, and sellers would have no incentive at all to improve their products' safety. If, instead, consumers overreacted to changes in safety, we would have $R_I(s,d) > 1$, and sellers would have an incentive to produce too much safety.

These equations also include the special case of perfect information, where $R(s,d) = s$ for all values of s and d , implying $R_I(s,d) = 1$ and $R_2(s,d) > 0$. In that case, equations (6) and (7) would be identical to equations (2) and (3), thus eliminating both of the distortions referred to above.

Notice, though, that these distortions do not both disappear whenever consumers accurately perceive the equilibrium levels of risk—that is, they do not disappear whenever $\hat{r} = \hat{s}$. True, any distortion of *the quantity purchased* will disappear if $\hat{r} = \hat{s}$, because that would mean consumers accurately perceived whatever level of safety was built into the products currently available on the market. However, $\hat{r} = \hat{s}$ does not preclude the possibility of $R_I(\hat{s},d) < 1$, because it is still possible that consumers might underestimate any *change* above or below the level of safety that sellers currently offer. In that case, sellers would choose inappropriately low (or high) level of safety, in spite of the fact that no consumer would be fooled in equilibrium, because $\hat{r} = \hat{s}$ implies that consumers end up with perfectly accurate estimates of the level of safety offered by all sellers in equilibrium.

Indeed, this is similar in spirit to George Akerlof's "market for lemons,"⁸⁷ where consumers accurately perceive the average level of safety in the market at any given time ($\hat{r} = \hat{s}$, in the notation I have used here), even though they would not be able to tell if any individual seller departed from that equilibrium level. If consumers know only the average level of safety, then if one of n sellers increased the safety of its products, consumers would react no differently than they would if all sellers had increased their safety by a fraction $1/n$ of that amount. For any individual seller, then, Akerlof (in effect) posited $R_I(\hat{s},d) = 1/n < 1$, which satisfies the condition for sellers to underinvest in safety.

86. In the torts literature, this corresponds to what is sometimes called the "level of activity" effect. See Shavell, *supra* note 85.

87. Akerlof, *supra* note 16.

Finally, we can now address the effects of increased disclosures. Differentiating equation (5) with respect to d shows that, for disclosures to correct consumers' tendency to purchase too much (or too little) of the product, the R function would need to have one or both of the following properties:

$$(9) \begin{aligned} R_2(s,d) &< 0 \text{ for } r > s \\ &> 0 \text{ for } r < s \end{aligned}$$

In other words, it must be possible for disclosures to reduce consumers' estimates r in cases where those estimates are too high, or to reduce r in cases where r is too low.

By contrast, different characteristics are required for disclosures to be capable of correcting sellers' tendency to produce too low (or too high) a level of safety. To correct those problems, the R function will need to have one or both of the following properties:

$$(10) \begin{aligned} R_{12}(s,d) &> 0 \text{ for } R_1(s,d) < 1 \\ &< 0 \text{ for } R_1(s,d) > 1 \end{aligned}$$

In other words, if most consumers currently under-respond to changes in product safety, disclosures must somehow succeed in magnifying their response to any change. And if, instead, consumers currently over-respond to any change, disclosures must somehow succeed in dampening that reaction.