ISSN 1936-5349 (print)
ISSN 1936-5357 (online)

HARVARD

JOHN M. OLIN CENTER FOR LAW, ECONOMICS, AND BUSINESS

'THEY RUINED POPCORN': ON THE COSTS AND BENEFITS OF MANDATORY LABELS

Cass R. Sunstein

Forthcoming in *Law & Marketing*, Jacob Gersen & Joel Steckel, eds.
Cambridge University Press

Discussion Paper No. 949

02/2018

Harvard Law School
Cambridge, MA 02138

Very preliminary draft 12/22/17

Forthcoming, Law & Marketing, Jacob Gersen & Joel Steckel, eds., Cambridge University Press.

"THEY RUINED POPCORN": ON THE COSTS AND BENEFITS OF MANDATORY LABELS

Cass R. Sunstein*

Abstract

Do consumers benefit from mandatory labels? How much? These questions are difficult to answer, because assessment of the costs and benefits of labels presents serious challenges. In the United States, federal agencies have (1) claimed that quantification is essentially impossible; (2) engaged in "breakeven analysis"; (3) projected various endpoints, such as health benefits or purely economic savings; and (4) relied on private willingness-to-pay for the relevant information.

All of these approaches run into serious normative, conceptual, and empirical objections. Approach (3) will exaggerate what consumers gain, because many people suffer welfare losses when they see labels, whether or not they end up making different choices. (Part of that loss is captured in one reaction to mandatory calorie labels: "They ruined popcorn!") In principle, approach (4) is usually best, but people may lack the information that would permit them to say how much they would pay for (more) information, and sometimes tastes and values shift over time, which means that willingness to pay may fail to capture welfare effects. These points raise fundamental conceptual, normative, and empirical questions about welfarist approaches to public policy.

When should government mandate labels? When would mandatory labels have desirable consequences for social welfare?

* Robert Walmsley University Professor, Harvard University. This essay draws on Cass R. Sunstein, On Mandatory Labeling, With Special Reference to Genetically Modified Foods, 165 U. Pa. L. Rev. 1043 (2017). I am grateful to Oren Bar-Gill for indispensable help and for collaborative work that has had a large impact on my presentation here.

How can those consequences be measured? When would labels do more good than harm?

These questions arise in many contexts, involving (for example) calorie labels, mortgage disclosures, energy efficiency labels, fuel economy labels, credit card disclosures, labels for genetically modified food, nutrition facts panels, country of origin labels, graphic warnings for cigarettes, and much more. Some of these labels are designed to enable consumers to protect themselves from risks, involving money or health. Some of them attempt to protect third parties of respond to moral concerns – as, for example, when labels offer information that bears on animal welfare. Some of them respond to some kind of consumer (or interest-group) demand for government action, whether or not risks are involved.

In all of these cases, assessment of welfare effects can be challenging. Sometimes government agencies know far too little to make any kind of projection of likely effects, and they simply confess that fact. Sometimes they can engage in "breakeven analysis," explaining that if the benefits reach a certain level, the costs of labels will be justified. Sometimes agencies are able to quantify the benefits and costs of mandatory labels, or at least significant subsets of them, either by using endpoints (economic savings or health benefits) or by measuring private willingness-to-pay for labels. Sometimes they can point to human dignity, equity, or distributional concerns.

As we shall see, the costs of labels may be higher than is readily apparent, because they may produce subtle decreases in consumer welfare -- as, for example, when calorie labels lead people to buy goods that are lower-calorie but less tasty, or when energy efficiency labels lead people to purchase appliances that cost less to operate but are less attractive. The point was captured in a reaction of one government official to mandatory calorie labels: "They ruined popcorn!"

As we will see, private willingness-to-pay is the best approach in theory, because it should capture everything that consumers stand to gain from labels. But obtaining a useful measure raises serious empirical, normative, conceptual challenges. A central reason is that to be worth using, willingness to pay should be informed, and often consumers lack the information that would enable them to decide how much to pay for (more) information. Another reason is that in some of the relevant contexts, preferences may be labile and endogenous. Once informed about health risks associated with

certain foods, for example, people might (begin to) develop different tastes. On optimistic assumptions, for example, salt and sugar labels can lead to transformations in tastes. Ex ante willingness to pay figures will be insufficiently informative on that count, which creates serious problems for welfare analysis.[1]

PRODUCT LABELING IN GENERAL

Market Failure?

When should government require products to be labeled? Suppose that we care about social welfare, suitably specified, and answer that labels should be required when they would do more good than harm. It is easy to imagine labels that are unnecessary, that are costly to impose, that are widely ignored by consumers, that mislead consumers, or that promote the interests of powerful private groups, not of the public as a whole. It is also easy to imagine labels that help consumers to save money, to avoid serious risks, to protect third parties, or to register their deepest moral commitments. Under the standard economic approach, the initial question is whether there is a market failure. In many cases, we expect the market to produce the necessary information on its own.[2] In other words, sellers are expected to disclose relevant information voluntarily. Mandatory disclosure is needed only when voluntary disclosure fails.

When offering accounts of market failure under the requirements of prevailing executive orders, agencies usually ask about what consumers are likely to demand. A standard market failure, often invoked by agencies themselves, involves *incomplete information*. Sometimes consumers lack the information that would enable them to make (sufficiently informed) choices, and government provides that information in order to make the market work efficiently.

It is true, of course, that consumers sometimes insist on product-related information, and hence the market will provide it; there is no need for a mandate. But consumers might not have the information that would put them in the position to demand disclosure of (further) information, and it might not be rational for them to attempt to acquire that information. Consider the health risks posed by trans

[1] For an important treatment of the general question, involving larger changes than those discussed here, see L.A. Paul. Transformative Experience (2014).

[2] *See* Howard Beales et al., *The Efficient Regulation of Consumer Information*, 24 J.L. & ECON. 491, 502 (1981) ("The economic incentive for consumers to gather information is strong. Increases in the efficiency of purchase decisions made are equivalent to increases in real income, and, given the diversity of choices available in a modem economy, improved choices can lead to a large gain. In many markets, price dispersion is substantial for identical or similar products.").

fats, which raise highly technical questions. Rational ignorance on the part of consumers might lead them not to acquire information from which they would ultimately benefit. Without that information, they might lack the knowledge that would lead them to even ask for labels. For that reason, a government response might be appropriate.

A further problem stems from the fact that information has the characteristics of a public good, which means that the market will not generate enough of it. Acting on his or her own, each consumer might not seek information from which all or most consumers would benefit. Mandatory labels overcome a collective action problem.

Yet another problem arises when the point of disclosure is to protect third parties. Often consumers want to know whether products are harming people, but even if they do not, disclosure might be required in order to reduce that harm. Suppose, for example, that disclosure of information is designed to reduce the risks of second-hand smoke, to prevent harms to animals (such as elephants or dolphins), to protect vulnerable groups (as with disclosure of "conflict minerals"), or to protect American jobs (as with "country of origin" or "made in America" labels). If third parties are at risk, we have a standard argument for government intervention. To the extent that GM food is thought to pose risks to the environment, a market failure seems to be involved. It is true, of course, that the preferred response to such risks is some kind of corrective tax, not disclosure. But if a tax is unavailable, for political or other reasons, then disclosure might seem to be a reasonable second-best.

There are behavioral issues as well. If risks are not sufficiently salient, then consumers might not demand relevant information about them, even if those risks are not exactly trivial. In principle, disclosure could therefore increase consumer welfare.[3] Or suppose that health risks are long-term; if so, then "present bias" might lead consumers not to demand information about them.[4] It is true that in the face of present bias, disclosure might not do much good; present-biased consumers might not care about what they learn.

[3] *See* Xavier Gabaix & David Laibson, *Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets*, 121 Q.J. ECON. 505, 511 (2006) (arguing that when consumers have access to more information, they are able to "make more informed choices among the available goods").

[4] *See generally* Ted O'Donoghue & Matthew Rabin, *Present Bias: Lessons Learned and To Be Learned*, 105 AM. ECON. REV. 273, 274-75 (2015) (discussing the impact of present bias on economic decisionmaking).

But perhaps information could be provided in a way that would reduce present bias. For example, labels might be graphic or specifically focus people on what might happen in the long-term.

Producer Behavior

Notwithstanding these points, a standard argument predicts voluntary disclosure even if consumers do not demand it. Assume that for whatever reason (rational or not), consumers would choose non-GM foods if they were given the information that would enable them to do so. Specifically, assume that consumers are willing to pay $10 for GM salmon and $20 for non-GM salmon. Further, assume that GM salmon costs $5 to produce, whereas non-GM salmon costs $7 to produce. Finally, assume that, initially, half the salmon on the market is GM and half is not. Without any labeling, the consumer would not know what kind of salmon she is buying and would, therefore, be willing to pay $15 (= 0.5*$10 + 0.5*$20). This state of (consumer) ignorance benefits the producers of GM salmon and harms the producers of non-GM salmon.

But this state of ignorance is not an equilibrium. The non-GM sellers will voluntarily add a "No GMOs" label so that they can charge $20, rather than $15 per salmon (as long as the cost of adding such a label is less than $5 per salmon). The GM salmon will not be labeled, but GM labeling would not be necessary—rational consumers would infer that non-labeled salmon is GM. As Bar-Gill and Board explain, "An implication of this result is that mandatory disclosure of product-attribute information is often unnecessary."[5]

In the example just given, the relevant quality dimension is binary (GMO or non-GMO). A similar argument predicts voluntary disclosure when the relevant quality dimension is continuous. Assume that different microwave ovens in the market emit radiation in the range of 0–10 mW/cm², with levels of radiation distributed uniformly (such that, for example, the number of microwave ovens emitting no radiation is equal to the number of ovens emitting 1 mW/cm² of radiation, and equal to the number of ovens emitting 2 mW/cm² of radiation, and so on). Without any labeling, consumers would not be able to distinguish low-radiation ovens from high-radiation ovens and would attribute the average radiation level, 5 mW/cm², to any oven they consider purchasing. Producers of low-radiation ovens, with radiation levels below 5 mW/cm², would be harmed by this state of consumer ignorance. These producers would voluntarily disclose their ovens' radiation levels.

Now consumers would know the radiation levels of all ovens with levels below 5 mW/cm². And when considering a non-labeled oven,

5 Oren Bar-Gill & Oliver Board, *Product-Use Information and the Limits of Voluntary Disclosure*, 14 AM. L. & ECON. REV. 235, 237 (2012).

the consumer would assume an average radiation level of 7.5 mW/cm². But then producers with radiation levels between 5–7.5 mW/cm² will voluntarily disclose. Only producers with radiation levels between 7.5–10 mW/cm² will remain silent, and so consumers would attribute an average radiation level of 8.75 mW/cm² to a non-labeled oven. Now producers with levels between 7.5–8.75 mW/cm² will voluntarily disclose. And so on, until complete unraveling is achieved and all information is voluntarily disclosed.

As a real-world example analogous to the question of GM food, consider the example of gluten free foods. Some people (including those with celiac disease) are allergic to food that contains gluten. At least to date, we do not observe statutory disclosure requirements ("Warning: this product contains gluten."). Instead we see voluntary labels, saying (for example) that products are "gluten free." The FDA has issued guidance for such labels.[6] On admittedly optimistic assumptions, voluntary labels provide sufficient information.

Markets that Do Not Unravel

This happy unraveling story, however, does not always play out. Failure of voluntary disclosure occurs for several reasons—some neoclassical and some behavioral. Starting with the standard, neoclassical reasons, note that the unraveling result assumes that voluntary disclosure is truthful. But imperfect enforcement might lead to false disclosures, which government must correct—and once government is in the business of correction, it may be essentially mandating a label.

In addition, voluntary disclosure might fail when there is no standardized format or metric for disclosing information. Without standardization, consumers might not be able to make the required distinctions, in which case voluntary disclosure will be insufficient. And if the point of disclosure is to protect third parties, the unraveling story might not work because consumers might not care enough about third party effects to respond to the various informational signals. True, consumer indifference would also mean that mandatory labels would be ineffective. But it is plausible to think

6 For a summary of gluten-free labeling from the FDA itself, see *Gluten and Food Labeling*, U.S. FOOD & DRUG ADMIN., http://www.fda.gov/Food/GuidanceRegulation/GuidanceDocumentsRegulatoryInformation/Allergens/ucm367654.htm [https://perma.cc/VVP8-98YG] (last updated May 2, 2016).

that consumers care *some*—enough to make mandatory labels work but not enough to promote unraveling.

Behavioral economics suggests an additional and perhaps stronger reason for skepticism about voluntary disclosure. The unraveling result assumes that consumers attend to and draw rational inferences from silence—from the absence of a label. But attention is limited,[7] and such inferences can be quite difficult to draw, especially when consumers are receiving numerous signals at the same time (as is true for food) and when there are multiple quality levels or continuous quality dimensions. Suppose, for example, that some products come with labels saying "low fat" or "low sugar." Would consumers necessarily infer that products lacking such labels are high in fat or sugar? Or would many consumers not think much or at all about the question of fat or sugar?

A standard neoclassical argument is that in a generalization of the "lemons equilibrium,"[8] competition might occur over easily observed characteristics, such as price, and less or not at all over less observable characteristics, such as ingredients.[9] The behavioral suggestion (or exclamation point) is that in view of the scarcity of attention, this limited kind of competition is highly likely. And even if consumers pay attention to the relevant ingredient (salt, sugar, fat), they might be unable to draw a fully rational inference from the absence of disclosure.

For example, those who are purchasing cereal or milk might attend to a variety of product attributes, and unless high fat or high sugar content is brought to their attention, many of them might not consider those ingredients at all. If many consumers would not pay attention or draw a negative inference (or a sufficiently negative inference) from the absence of a label, voluntary disclosure might fail. Such failure justifies the consideration of mandatory disclosure, at least in principle. The Affordable Care Act, for example, mandates calorie labels, and there is a plausible argument on their behalf based on the considerations just sketched.

7 *See* DANIEL KAHNEMAN, ATTENTION AND EFFORT 13-17 (1973) (discussing constraints on consumer attention in relation to informed decision-making). *See generally* SENDHIL MULLAINATHAN & ELDAR SHAFIR, SCARCITY: WHY HAVING TOO LITTLE MEANS SO MUCH (2013) (examining how cognitive scarcity and "limited bandwidth" affect choices).

8 *See generally* George A. Akerlof, *The Market for "Lemons": Quality Uncertainty and the Market Mechanism*, 84 Q.J. ECON. 488 (1970) (discussing the relationship between product price and quality and consumer demand).

9 *See* Beales et al., *supra* note 2, at 510 ("By generalizing the concept of the 'lemons' equilibrium, we can show that, if price is more easily observed than quality, competition may be skewed toward less expensive, lower-quality products." (internal citation omitted)).

"Does Not Contain" Labels vs. "Contains" Labels

There are many differences between a system in which products without some characteristic say "Does Not Contain *X*" and one in which products with some characteristic say "Contains *X*." As we have seen, "Contains *X*" offers far more salient information to consumers with bounded attention. In addition, "Contains *X*" might offer a distinctive signal, suggesting that private and public institutions think that something is wrong with *X*.

"Does Not Contain *X*" might also promote a desirable form of sorting. Suppose that ten percent of the population is troubled by *X*, whereas ninety percent is not; suppose that both groups are informed and rational. If so, there is no need for "Contains *X*." Those who want to avoid *X* can easily do so, and those who have no interest in avoiding *X* need not be troubled by the issue.

On a certain view of the facts, "Does Not Contain *X*" is the right approach both to gluten-free and to GM food. People who are allergic to gluten should know what to look for. The principal problem is that if they are inattentive, they might become sick simply by virtue of the fact that the issue has not been brought to their attention. (Compare labels saying "Contains peanuts" or "Contains shellfish," which may be especially important if consumers are inattentive or if it is not self-evident that the relevant food contains either.) With "Does Not Contain" labels, consumers can easily avoid GM food if that is what they want to do. But this approach is not a solution if GM food has harmful systemic effects or threatens to cause environmental harm (or if relevant interest groups want to stigmatize GM food).

Costs and Benefits

Even if there is a market failure, the question remains: do the benefits of labels justify the costs? If it would be expensive to comply with a labeling requirement—say, $800 million annually—the question whether the benefits are sufficient would be put in stark relief. We could easily imagine disclosure requirements that do little good, perhaps because consumers pay no attention to them. If so, such requirements would be unjustified on cost–benefit grounds.

We could also imagine disclosure requirements from which consumers and third parties would benefit greatly.[10]

As we will see, agencies have not always responded well to the difficulty of quantifying the costs and benefits of disclosure requirements. In fact, they have adopted four distinctive approaches, imposing increasingly severe information-gathering demands on agencies. It is not always easy to explain why they choose one or another in particular cases.

The first approach—and it may be the most candid—is to confess a lack of knowledge by acknowledging that, in light of existing information, some costs and (especially) benefits simply cannot be quantified.[11] The problem with this approach is that it suggests that the decision to proceed is essentially a stab in the dark. When the stakes are large, that seems unacceptable, certainly for policymakers.

The second approach involves "breakeven analysis," by which agencies describe what the benefits would have to be in order to justify the costs—and suggest that the benefits are indeed likely to be of the requisite magnitude. In principle, this approach is better than a simple confession of ignorance, and it is often the best path forward. But it involves a high degree of guesswork, and it may be a mere conclusion, a kind of ipse dixit, masquerading as an analytic device. Without a great deal of discipline, it too may not be so different from a confession of ignorance.

10 *See* Partha Deb & Carmen Vargas, *Who Benefits from Calorie Labeling? An Analysis of Its Effects on Body Mass* 14 (Nat'l Bureau of Econ. Research, Working Paper No. 21,992, 2016), http://www.nber.org/papers/w21992 [https://perma.cc/B695-RL98] (finding that the "mandatory calorie labeling laws implemented over the past few years in a number of states and counties appear to be having substantial effects in terms of decreased BMI following implementation of such laws"); *see also* FUNG ET AL., *supra* note **Error! Bookmark not defined.**, at 1-10 (discussing the virtues and power of disclosures).

11 For an important decision upholding a refusal to quantify benefits, on the ground that quantification was not feasible, see Investment Co. Institute v. Commodity Futures Trading Comm'n, 720 F.3d 370, 372-75 (D.C. Cir. 2013). In the context of disclosure, the leading decision is National Ass'n of Manufacturers v. SEC, which upheld against arbitrariness review a regulation that would require disclosure of the use of "conflict minerals":

> An agency is not required "to measure the immeasurable," and need not conduct a "rigorous, quantitative economic analysis" unless the statute explicitly directs it to do so. Here, the rule's benefits would occur half-a-world away in the midst of an opaque conflict about which little reliable information exists, and concern a subject about which the Commission has no particular expertise. Even if one could estimate how many lives are saved or rapes prevented as a direct result of the final rule, doing so would be pointless because the costs of the rule—measured in dollars—would create an apples-to-bricks comparison. Despite the lack of data, the Commission *had* to promulgate a disclosure rule.

748 F.3d at 369 (quoting Inv. Co. Inst. v. Commodity Futures Trading Comm'n, 720 F.3d at 379).

The third approach is to attempt to specify outcomes in terms of (say) economic savings or health endpoints. The advantage of this approach is that it actually points to concrete benefits, and it attempts to measure and to monetize them. But it too runs into difficulties. The first is that agencies may lack anything like the information that would enable them to venture such a specification. The second and more interesting is that, for reasons I will explore, even an accurate specification will not give a complete picture of the actual benefits, and, in crucial respects, it will almost certainly overstate them. In brief, the problem is that people might experience significant losses as well as gains as a result of the label (for example, if they switch to a product that is inferior along certain dimensions), and an account of endpoints will ignore those losses.

The fourth approach is to identify consumers' willingness-to-pay. As a matter of abstract principle, that approach is (mostly) the right one, because it should capture the full universe of losses and gains from the label. At the same time, it runs into serious and perhaps insuperable normative, conceptual, and empirical challenges. As we shall see, the most obvious problem is that it is difficult to elicit people's *informed and unbiased* willingness-to-pay for labels. The most interesting problem involves the potentially labile character of some preferences.

Costs

On the cost side, some of the questions are relatively straightforward. Regulators may well be able to learn the total cost of (for example) producing fuel economy labels and placing them on new vehicles. The principal difficulty arises *when the information itself imposes costs on consumers*. It is a mistake to ignore those costs, even if they prove difficult to quantify, and even if consumers benefit on net.[12] Those costs come in several different forms. Some of them will usually be low—but not always.

12 For a useful discussion in an especially controversial area, see Helen Levy et al., *Tobacco Regulation and Cost-Benefit Analysis: How Should We Value Foregone Consumer Surplus?* (Nat'l Bureau of Econ. Research, Working Paper No. 22,471, 2016), http://www.nber.org/papers/w22471.pdf [https://perma.cc/G6RD-9ZHD].

A Small Cognitive Tax

First, a cost is involved in reading and processing the information. For each consumer, that cost is likely to be quite low, but across a large number of purchasers, it might turn out to be significant. Information disclosure is, in a sense, akin to a paperwork burden. To be sure, consumers are not compelled to read and process what is disclosed. But even for those who seek to ignore it, its very presence may operate as a kind of cognitive tax. Because people have limited bandwidth, that tax may not be safely ignored. (If there is a Hell, it may well be filled with warnings.)

Ruining Popcorn, 1: A Hedonic Tax on Those Who Do Not Change Their Behavior

Second, and more importantly, the cost may be hedonic, not cognitive. Suppose that smokers are given information about the adverse health effects of smoking or that visitors to chain restaurants are given information about the caloric contents of food. Many members of both groups will suffer a hedonic loss. Consider, for example, smokers who cannot quit and customers who decide to choose high-calorie foods notwithstanding the labels. In hedonic terms, such people will lose, rather than gain, if they are miserable or at least sadder at the time of purchase.

To be sure, there is a normative question whether regulators should count, as costs, the adverse hedonic effect of truthful information. Is it a cost, or a benefit, if people learn, truthfully, that they have diabetes or cancer? On net, that might well be a benefit, at least if they can do something about the problem. But there is a cost as well, and a large one, even if the net effect is positive. If we are operating within a welfarist framework, the hedonic loss must be treated as a cost. It might turn out to be low, but regulators should not ignore it (as they typically do).

Compare: Many people do not want to get blood tests, even if doctors advise them to do so, because they do not want to bear the hedonic cost of less-than-good results. The failure to get the tests might be a product of a behavioral bias (for example, present bias), but it might also be a product, in part, of a rational aversion to negative information.

Ruining Popcorn, 2: A Hedonic Tax on Those Who Do Change Their Behavior

Even if people might be able to quit smoking or end up choosing lower-calorie items, and will hence benefit greatly on net, they will incur a cost by seeing something that inflicts pain. In principle, that cost should also count, even if it is greatly outweighed by benefits. The point, then, is not that the hedonic cost is necessarily a trump

card; if people make different choices once they are informed, the presumption should be that they are better off. But *by how much*?

To answer that question, the hedonic cost must be taken into account. For many people, a calorie label imposes a serious cost, simply because it informs them that the delicious cheeseburger they are about to eat is also going to make their belly bulge.

A Consumer Welfare Loss

There is a fourth loss, in the form of foregone consumer surplus. Suppose that people decide that on balance, they should have a salad rather than a cheeseburger, on the ground that the latter has many more calories. If they choose the salad because of the label, they are probably better off on balance—and in a sense, they are sadder but wiser (and healthier). They are sadder to the extent that they enjoy their meal less. Assessment of the magnitude of the loss poses serious conceptual and empirical challenges, but there is no question that it exists, and that it might turn out to be a significant fraction of the benefits. In principle, a decision to forego the hamburger might make people only modestly better off, if the hedonic loss is almost as high as the health gain.

Suppose, for example, that consumers are choosing between two essentially equivalent cars; that the more fuel-efficient one would cost $2000 less annually to operate because of its fuel efficiency; that the less fuel-efficient one would cost $500 upfront; and that because of the fuel economy label, they select the fuel-efficient car. For each such consumer, we might be tempted to say that the label has produced $1500 in gains. But in actual practice, the effects of a fuel economy label will be much more complicated to assess. Some consumers will end up purchasing cars that are more fuel-efficient but inferior along some dimension, so that they will gain $1500 minus *X*, where *X* refers to the desirable features of the unchosen car that they otherwise prefer. It is hard for public officials to know whether *X* is, on average, $100, or $1000, or $1450.

The Problem of Endogenous Preferences

All this assumes that preferences are consistent and exogenous. In some contexts, however, that assumption is not correct.[13] This

[13] *See* PREFERENCE CHANGE: APPROACHES FROM PHILOSOPHY, ECONOMICS AND PSYCHOLOGY 4 (Till Grüne-Yanoff & Sven Ove Hansson eds., 2009) (discussing the work

point complicates the foregoing analysis and creates a risk that analysis of costs will ignore shifts in tastes that are induced by labels themselves.

Suppose that at Time 1, people enjoy hamburgers a lot and enjoy salads only a little. Now suppose that having seen the labels, people switch at Time 2 because they want to make healthier choices. At Time 2, they suffer costs as a result of the switch; they miss hamburgers (delicious!) and they do not much like salad (boring!). But at Time 3, people might come to dislike hamburgers (disgusting!) and to love salad (fresh!). In principle, preference change must be taken into account by the considered cost–benefit analysis, though doing so presents serious challenges: it might be difficult to know the magnitude of the change and even the sign (perhaps those who switch to salad will crave hamburgers and grow to despise salad).

Benefits

On the benefits side, the assessment is even more challenging.[14] If the government mandates a fuel economy label, agencies should project the economic and environmental benefits from the mandate. But to do that, they have to *know the effect of labels on behavior*. In principle, a randomized controlled trial would be valuable and perhaps necessary for that purpose. If one group sees a particular label and a similar group sees a different label (or no label), regulators should be able to specify the effect of the label

of economists to incorporate individuals' changing preferences into models of supply and demand).

14 For example, according to the Environmental Protection Agency and the Department of Transportation, speaking of new fuel economy labels,

> The agencies recognize that Executive Order 13563 directs agencies "to use the best available techniques to quantify anticipated present and future benefits as accurately as possible." In this context, however, quantitative information is not available, and the agencies have therefore chosen instead to continue with a qualitative assessment of benefits. It is difficult to develop a good baseline for the fleet using the existing label, partly because the existing label is not designed to incorporate advanced technology vehicles. It is even more difficult to develop a comparison for the fleet with the new labels, because the effects of label designs on vehicle purchases are not known. Thus, any assessment of quantitative effects of label design on vehicle sales involves a great deal of speculation. The agencies believe that informed choice is an end in itself, even if it is hard to quantify; the agencies also believe that the new labels will provide significant benefits for consumers, including economic benefits, though these benefits cannot be quantified at this time.

Revisions and Additions to Motor Vehicle Fuel Economy Label, 76 Fed. Reg. 39,478, 39,517 (proposed July 6, 2011) (to be codified at 40 C.F.R. pts. 85, 86, 600; 49 C.F.R. pt. 575) [hereinafter Fuel Economy Labels Rule]. In short, "The primary benefits associated with this rule are associated with improved consumer decision-making resulting from improved presentation of information. At this time, EPA and NHTSA do not have data to quantify these impacts." *Id.*

on purchasing decisions. Armed with that information, they could estimate economic and environmental consequences (at least if they could generalize from the trial).

Unfortunately, it is sometimes difficult or impossible to run randomized controlled trials. In these circumstances, making any kind of projection of how consumers will react to a label is exceedingly difficult. An additional problem is that for the reasons given thus far, the projection would not give an adequate estimate of the (net) benefits. We have seen that if people are buying cars that are more fuel-efficient but otherwise highly undesirable, there will be a welfare loss. For that reason, regulators might explore the issue from another direction.[15] Rather than asking about the economic savings from the fuel-efficient car, they might ask an entirely different question: *how much would consumers be willing to pay for a fuel economy label*?

Under ideal conditions and under plausible assumptions, and bracketing the endogeneity issue, the right question for regulators to ask involves willingness-to-pay; they should not focus on the economic benefits that consumers might receive if (for example) they purchase more fuel-efficient cars. The reason is that on optimistic assumptions, the willingness-to-pay question ought to capture everything that matters to consumers. (Of course it is true that the question will not fully capture third-party effects, nor will it capture welfare effects if preferences are endogenous.)

As an empirical matter, however, it is not easy to obtain a reliable answer to that question, or anything close to it. We might simply ask people, but for their answers to be relevant, it would be important to provide pertinent information—for example, about the potential benefits (purely economic and otherwise) of labels. Providing that information is no simple endeavor, not least because offering some numbers about those potential benefits would be important, and any numbers might "anchor" consumers and hence bias their answers. Suppose that the problem of anchoring could be overcome and that informed consumers would be willing to pay (say) $10, on average, for fuel economy labels. If so, we might have some sense of the

15 *See* Hunt Allcott & Judd B. Kessler, *The Welfare Effects of Nudges: A Case Study of Energy Use Social Comparisons* 2 (Nat'l Bureau of Econ. Research, Working Paper No. 21,671, 2015), http://www.nber.org/papers/w21671 [https://perma.cc/9MD3-8QX2] (noting that nudges "can affect behavior without changing prices or choice sets").

benefits, at least if behavioral biases are not distorting people's answers.

Unfortunately, however, such biases might well produce distortions; consider present bias and optimistic bias, which may lead to unduly low willingness-to-pay. In any case, survey evidence is imperfectly reliable, in part because of the familiar problems with contingent valuation studies, in part because of the immense difficulty of informing consumers in a sufficiently neutral way.

For health-related disclosures, the problem is even harder. One goal of calorie labels, for example, is to reduce obesity, which causes an assortment of health problems, including premature mortality. Regulators have established ways to turn health-endpoints into monetary equivalents. For example, a statistical death is now valued at about nine million dollars.[16] But how many premature deaths would be prevented by calorie labels? And what would be the effect of such labels on adverse health outcomes short of death?

To answer such questions, regulators have to undertake two tasks. First, they must begin by making some prediction about the effect of calorie labels on what people choose to eat. Second, they have to follow that prediction by specifying the health consequences of lower levels of caloric intake. At least it can be said that if they can accomplish those tasks, they will have some sense of the benefits of the labels, once (and this is a third task) they turn the various consequences into monetary equivalents. After undertaking all three tasks, regulators will have specified endpoints—but for the reasons given, a specification of endpoints will overstate benefits because it will not include various cognitive and hedonic losses.

Alternatively, we could (again) ask how much people would be willing to pay for calorie labels.[17] As before, asking that question is, in principle, preferable to an effort to assess health-states, because the answer will capture all variables that matter to consumers.[18] Also, as

16 *See, e.g.*, Memorandum from Kathryn Thomson, Gen. Counsel & Carlos Monje, Assistant Sec'y for Policy, U.S. Dep't of Transp., to Secretarial Officers & Modal Adm'rs, U.S. Dep't of Transp. 2 (June 17, 2015), https://www.transportation.gov/sites/dot.gov/files/docs/VSL2015_0.pdf [https://perma.cc/C6RQ-4ZXR] ("On the basis of the best available evidence, this guidance identifies $9.4 million as the value of a statistical life."); *see also* CASS R. SUNSTEIN, VALUING LIFE: HUMANIZING THE REGULATORY STATE 85 (2014) (providing the underlying theory and a discussion of how "agencies . . . assign monetary values to the human lives that would be saved by a proposed regulation").

17 *See* Maria L. Loureiro et al., *Do Consumers Value Nutritional Labels?*, 33 EUR. REV. AGRIC. ECON. 249, 263 (2006) (finding that "on average, consumers are willing to pay close to 11 per cent above the initial price to obtain cookies with nutritional labelling"); *see also id.* at 249 ("Consistent with prior expectations, our results also indicate a difference between the [willingness-to-pay] of individuals suffering from diet-related health problems (estimated mean 13 per cent) and those who do not suffer any diet-related health problems (estimated mean 9 per cent).").

18 In the words of the FDA,

before, there are formidable challenges in using surveys to elicit reliable numbers free from biases of various kinds. And if preferences are endogenous and labile, willingness to pay numbers might greatly understate the welfare gain from labels. Recall that people might develop tastes for the products to which they shift. (I am also bracketing the questions raises by addictive goods, such as cigarettes, for which labels might be beneficial on welfare grounds precisely because they help break the hold of the addiction. Note that cigarette taxes appear to make smokers happier.[19])

In light of these challenges, regulators have two imperfect options. First, they can work on the two relevant tracks to try to produce answers: exploring end-points and enlisting surveys. On prominent occasions, they have tried the former.[20] Second, they can acknowledge the difficulties, confess that they cannot surmount them, and use "breakeven analysis," by which they ask what the

> To our knowledge, Abaluck (2011) is the only study that translates the potential effect of increasing nutrition information on consumption into estimates of welfare gains using willingness-to-pay based on revealed preferences (Ref. 43). This study uses the variation in nutrition information generated by Nutrition Labeling and Education Act (NLEA) as a method to determine how changes in individuals' beliefs about nutrient content affect consumption decisions. The differential changes in nutrition information across food categories, measured in units of calories per gram, allow the study to identify a general model of food demand as a function of nutrient characteristics that accounts for the total daily diet, prior beliefs about nutrient content, and preferences, including willingness to substitute across food categories.

CALORIE LABEL RULE IMPACT ANALYSIS, *supra* note **Error! Bookmark not defined.**, at 64. As before, however, the willingness-to-pay criterion may run into normative objections, even from the standpoint of welfare. *See generally* JOHN BRONSTEEN ET AL., HAPPINESS AND THE LAW (2015) (raising questions about willingness-to-pay in view of people's occasional failure to know what will promote their welfare).

[19] See Jonathan Gruber and Sendhil Mullainathan, Do Cigarette Taxes Make Smokers Happier? (2002), available at http://www.nber.org/papers/w8872

20 *See* Graphic Warnings for Cigarettes Rule, *supra* note, at 36,719 (noting the longer lifespans, fewer cancers and diseases, as well as increased property and monetary values of non-smokers); Improve Tracking of Workplace Injuries and Illnesses, 81 Fed. Reg. 29,624, 29,628 (proposed May 12, 2016) (to be codified at 29 C.F.R. pts. 1904, 1902) [hereinafter OSHA Reporting Requirement Rule] (requiring that employees have access to OSHA logs); Fuel Economy Labels Rule, *supra* note 14, at 39,517 ("The agencies believe that informed choice is an end in itself, even if it is hard to quantify; the agencies also believe that the new labels will provide significant benefits for consumers, including economic benefits, though these benefits cannot be quantified at this time."); *see also* CALORIE LABEL RULE IMPACT ANALYSIS, *supra* note, at 11 ("The final rule may also assist consumers by making the long-term health consequences of consumer food choices more salient and by providing contextual cues of food consumption.").

benefits would have to be, in order to justify the costs, and then do what they can to generate a reasonable lower bound. Suppose, for example, that an energy-efficiency label for refrigerators would cost $10 million annually and that eight million refrigerators are sold in the United States every year. Even if the average consumer saves only $0.50 annually as a result of the label, the cost will be made up in just three years. Breakeven analysis can be crude, but in some cases, it will suggest that the argument for labels is either very strong or very weak.

Third Parties—and Morality

Some actual or imaginable labels are meant to protect third parties, not consumers as such. Suppose that some or many consumers are concerned about the use of certain minerals to finance mass atrocities, and they favor labeling, or some kind of disclosure requirement, so that consumers can decline to purchase products that contain such minerals. Or suppose that consumers care about where goods were made, perhaps because they want to purchase products from their own nation or perhaps because they do not want to purchase products from nations that do not respect human rights. They might seek "country of origin" labels for that reason. Or suppose that some or many consumers care about the welfare of animals in general or certain animals in particular; because they do, they seek labels to reflect how animals were (mis)treated.

In some of these cases, the third-party effects are not obscure, and the real challenge is how to quantify them. As before, it is necessary to begin by making some projections about consumer behavior. To what extent would consumers change their purchasing habits in response? Even if that question can be answered, it would be necessary to tie any such changes to reduced harm or increased benefit for third parties. And even if that problem can be resolved, it would be necessary to quantify and monetize the resulting effects. It is no wonder that in the context of conflict minerals, the agency concluded that quantification was not possible.[21] Perhaps it should have engaged in some form of breakeven analysis, explaining that the requirement was likely to survive cost–benefit analysis even if its effect were modest. But perhaps it lacked the information that would have allowed it to make that analysis plausible.

Some disclosure requirements, including mandatory labels, are not simple to defend within a standard cost–benefit framework, not for the reasons I have been sketching, but because considerations

[21] *See* Nat'l Ass'n of Mfrs. v. SEC, 800 F.3d at 547 ("The Commission was 'unable to readily quantify' the 'compelling social benefits' the rule was supposed to achieve: reducing violence and promoting peace and stability in the Congo." (quoting Conflict Materials, 77 Fed. Reg. 56,274, 56,350 (Sept. 12, 2012) (to be codified at 17 C.F.R. pts. 240 & 249b))).

of equity, distributional effects, or human dignity are involved. When values of this kind are involved, it is perfectly legitimate for agencies to consider them. It might well be sufficient for agencies simply to point to such considerations and not to fold them into a cost–benefit analysis. Agencies are authorized to give independent consideration to equity and human dignity. If the statutory goal is to achieve distributional goals, by transferring resources from some people to others, then cost–benefit balancing is not the rule of decision, and it is not all that matters. A rule might have costs in excess of benefits, in the sense that the losers lose more than the winners gain, but perhaps the winners are poor or otherwise deprived, and perhaps have a special claim to attention under the relevant law or as a matter of principle.

I have suggested that if quantification of the benefits of labels is required, the question might be: how much would (informed) consumers be willing to pay for such labels? Within a certain framework, that question is the right one. But even if we put the foregoing difficulties to one side, it is not at all clear that the framework is the right one. If the issue involves human dignity, equity, or distributional considerations—or any kind of harm to third parties—why should the proper analysis depend on how much people are willing to pay for it? It seems senseless to say that labels motivated by distributive goals should be imposed to the extent that people are willing to pay for them.

To say this is not to say that consequentialist considerations do not matter at all. Insofar as harms to third parties are involved, cost–benefit analysis can be used, acknowledging the empirical problems sketched above. Insofar as the issue involves equity or dignity, breakeven analysis might be useful. To the extent that distributive goals are involved, a key question is whether such goals would, in fact, be promoted by labels or disclosure. That question would seem relevant to the “conflict minerals” problem. Some kind of means–ends analysis, explaining how the means are connected to the ends, would seem indispensable to an evaluation of labels that are designed to promote distributive goals (or for that matter equity or human dignity). Agencies should be expected to undertake that analysis—or to explain why they cannot.

Taking Stock

In numerous contexts, Congress has required or authorized federal agencies to impose disclosure requirements. In all those contexts, executive agencies are required, by executive order, to catalogue the benefits and costs of disclosure requirements, and to demonstrate that the benefits justify the costs. Such agencies face persistent challenges in projecting benefits, and they use four different approaches, including a refusal to do so on the ground that quantification is not feasible; breakeven analysis; projection of end-states, such as economic savings or health outcomes; and estimates of willingness-to-pay for the relevant information.

Each of these approaches raises serious questions and runs into strong objections. In principle, the right question generally involves willingness-to-pay. But in practice, people often lack enough information to give a sensible answer to the question how much they would be willing to pay for (more) information. (How much would you be willing to pay for information about the presence of chemical XYZ in your favorite food, when you know little or nothing about chemical XYZ or its effects?)

We have also seen that when preferences are labile or endogenous, even a sensible answer may fail to capture the welfare consequences, because people may develop new tastes and values. In these circumstances, a breakeven analysis is the very least that should be required, and it is sometimes the most that agencies can do. If it is accompanied by some account of potential outcomes, acknowledging uncertainties, a breakeven analysis will often show that mandatory disclosure is justified on welfare grounds—and often that it is not.