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

HARVARD

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

DRAWING FALSE INFERENCES FROM MANDATED DISCLOSURES

Oren Bar-Gill
David Schkade
Cass R. Sunstein

Discussion Paper No. 912

06/2017

Harvard Law School
Cambridge, MA 02138

Drawing False Inferences from Mandated Disclosures

Oren Bar-Gill,[*] David Schkade,[**] and Cass R. Sunstein[***]

Abstract

Disclosure mandates are pervasive. Though designed to inform consumers, such mandates may lead consumers to draw false inferences – for example, that a product is harmful when it is not. When deciding to require disclosure of an ingredient in or characteristic of a product, regulators may be motivated by evidence that the ingredient or characteristic is harmful to consumers. But they may also be motivated by a belief that consumers have a right to know what they are buying or by interest-group pressure. Consumers who misperceive the regulator's true motive, or mix of motives, will draw false inferences from the mandated disclosure. If consumers think that the disclosure is motivated by evidence of harm, when in fact it is motivated by a belief in a right-to-know or by interest-group pressure, then they will be inefficiently deterred from purchasing the product. We analyze this general concern about disclosure mandates. We also offer survey evidence demonstrating that the risk of false inferences is serious and real. Our framework has implications for the ongoing debate over the labeling of food with genetically modified organisms (GMOs); it suggests that the relevant labels might prove misleading to some or many consumers, producing a potentially serious welfare loss. Under prevailing executive orders, regulators must consider that loss and if feasible, quantify it.

[*] William J. Friedman and Alicia Townsend Friedman Professor of Law and Economics, Harvard Law School.

[**] Jerome Katzin Professor of Management, Rady School of Management, University of California San Diego.

[***] Robert Walmsley University Professor, Harvard Law School. We are grateful to the Program on Behavioral Economics and Public Policy for support. For helpful comments and suggestions, we thank Lucian Bebchuk, John Coates, Einer Elhauge, and workshop participants at Harvard Law School.

I. Introduction

Red Auerbach, the late, great coach of the Boston Celtics, liked to say, "It's not what you say; it's what they hear." What do consumers "hear" when the government mandates the disclosure of a certain ingredient or characteristic of a product? Our argument, in brief, is that consumers often hear something very different from what the government intends to convey. The result can be a serious welfare loss, with harms to producers and consumers alike.

In many cases, consumers hear "DANGER! DON'T BUY!!" That may be precisely what the government wants consumers to hear. In such cases, the government concluded, on the basis of scientific evidence, that the relevant ingredient or characteristic is harmful to consumers, and it is using the disclosure mandate to convey this information and reduce demand for the harmful product. Think cigarette labels.

In other cases, however, the government does not want to send a "DANGER!" signal. There may be no scientific basis for concluding that the ingredient or characteristic is harmful. The disclosure mandate may be motivated by a belief that consumers have a right to know what they are buying, whether or not the ingredient or characteristic is harmful. Or it may be motivated by interest-group pressures. Or, perhaps, there is some preliminary evidence of possible harm, but far from enough to merit a "DANGER! DON'T BUY!" warning; only, maybe, a much weaker message: "Some Preliminary, Inconclusive Cause for Concern. Not Sure If You Should Buy or Not." Or government may be recognizing some kind of social value (say, on behalf of products bought in the country in which they are sold, or products with certain national origins[1]) or moral commitment (say, on behalf of animal welfare or natural products), which has nothing to do with health risks.

The problem is that, in these cases, consumers may hear "DANGER!" even though the government does not mean to issue a "DANGER!" warning at all. The concern is that the mandatory label would mislead consumers, thus producing a welfare loss. We study the inference problem that consumers face when the government decides to mandate the disclosure of an ingredient or characteristic of a product. Our analysis establishes that consumer's post-disclosure beliefs about the product are influenced by (1) the consumer's pre-disclosure beliefs, (2) the consumer's estimate of the accuracy of the government's information, and (3) the consumer's beliefs about the government's motives.[2]

---

[1] For a discussion of "country of origin labeling," see *Country of Origin Labelling (COOL)*, U.S. DEP'T OF AGRIC., https://www.ams.usda.gov/rules-regulations/cool

[2] A similar inference problem is studied in Juanjuan Zhang, Policy and Inference: The Case of Product Labeling (Sept. 23, 2014) (unpublished manuscript), http://jjzhang.scripts.mit.edu/docs/Zhang_2014_GMO.pdf. Zhang reports an interesting initial empirical study showing that people perceive higher risk following a GMO disclosure mandate (compared to no action). Our study includes (and replicates) this

The consumer's pre-disclosure beliefs play a critical role. Suppose that before learning of the government's decision to mandate disclosure, the consumer is fairly certain that the ingredient or characteristic is harmful. If so, the disclosure mandate will have a minimal effect on the consumer's post-disclosure beliefs (and perhaps none at all). Similarly, if, pre-disclosure, the consumer is fairly certain that the ingredient or characteristic is harmless, then again the disclosure mandate will have a minimal effect on the consumer's post-disclosure beliefs (and perhaps none at all). In essence, when consumers are already well-informed, or think that they are well-informed, the additional signal derived from the government's decision to mandate disclosure carries little weight.

In contrast, when, pre-disclosure, consumers are uncertain about whether the ingredient or characteristic is harmful or not, the government's decision to mandate disclosure will carry more weight. This means that we should be most worried about potentially misleading decisions to mandate disclosure when many consumers are uncertain about whether the ingredient or characteristic is harmful. In many areas, consumers, or a large number of them, are indeed uncertain, because the underlying questions are technical, complex, or subject to competing (but apparently plausible) interpretations.

The perceived quality or accuracy of the government's evidence about whether the ingredient or characteristic is harmful also affects the consumer's post-disclosure beliefs. When the government is thought to have superior information, the decision to mandate disclosure will naturally carry more weight. It follows that the perceived professional expertise of the government agency that decides to mandate the disclosure will affect the inferences that consumers draw from any such mandate. And this is all as it should be: consumers should give more weight to the government's decision to mandate disclosure when they believe that the government has better information and greater expertise. The concern that a disclosure mandate will mislead consumers arises when consumers over- (or under-) estimate the quality of the government's information or its level of professional expertise.

Finally, and perhaps most interestingly, the government's perceived motivation for mandating disclosure will critically influence the inferences that consumers draw from a decision to mandate disclosure. If consumers think that the government requires disclosure because it found that the product is harmful, then they will be more likely to revise their beliefs about the product's harmfulness. If, by contrast, consumers think that the government requires disclosure because it believes in a "right to know" or because it succumbed to interest-group pressure, then they will be less likely to revise their beliefs about the product's harmfulness. Again, this is all as it should be.

---

effect, also adding warning as a possible action. In addition, we examine the effect of the government's motive (stated and perceived), how this interacts with government action, how prior assessments are updated after learning the action and motive, effects on purchase intentions, and a comparison of the pattern of effects for GMOs to those for an unknown new product (Z25).

The concern, and our central focus here, is that a decision to mandate disclosure will mislead consumers. This concern arises *when consumers misperceive the government's motives* -- for example, if they think that the government decided to mandate disclosure because it concluded that the product is harmful, when in fact the disclosure mandate was motivated by a belief in a right to know.

In this Article, we analyze the factors that influence the inferences that consumers draw from a disclosure mandate, both theoretically and empirically.[3] In particular, we measure the effect of inferred motives on the inferences that consumers draw from mandated disclosures. Empirically, we confirm that consumers' beliefs about product risk increase when they think that the disclosure mandate was motivated by new research, but not when they think that the mandate was driven by political pressure. We obtain more subtle empirical results when consumers think that the government chose to mandate disclosure because it believes that consumers have a right to know (RtK) what they are buying. Puzzlingly but importantly, consumers who attribute to the government a RtK motive end up seem to perceive a higher level of risk. It appears that these consumers are incorrectly conflating a RtK motive with a new-evidence-of-risk motive.

We are especially concerned about updating that leads to false inferences about product risk. Such false inferences will occur when consumers attribute the wrong motive to the government's decision to mandate disclosure. In particular, consumers will wrongly increase their estimate of product risk, if they wrongly think that the disclosure mandate was motivated by new research finding that the product is harmful, when in fact the government's motives were very different. In the GMO context, where the actual disclosure mandate was not motivated by such new research, our survey results suggest that about 50% of consumers attribute a false motive and thus draw false inferences (the 50% figure includes consumers who attribute a RtK motive but think that a RtK is important because there is evidence of risk). We also confirm empirically that the magnitude of the false inference problem is inversely correlated with the strength of the consumer's priors about product risk.

What are the welfare costs of the false inferences that we identify? Quantification is challenging, but in qualitative terms, the answer is obvious: False inference leads to misperception of risk, and misperception of risk distorts consumers' purchase decisions. A deeper look reveals further insight, especially as we consider the task of regulators who are deciding whether to adopt a disclosure mandate and how to specify its costs and

[3] There is a great deal of work on the uses and limits of disclosure remedies in general. For a skeptical view, *see* OMRI BEN-SHAHAR & CARL E. SCHNEIDER, MORE THAN YOU WANTED TO KNOW: THE FAILURE OF MANDATED DISCLOSURE (2014); for a less skeptical view, *see* George Loewenstein et al., *Disclosure: Psychology Changes Everything*, 6 Annual Review of Economics 1 (2014), available at http://www.annualreviews.org/doi/full/10.1146/annurev-economics-080213-041341. To our knowledge, ours is the first systematic treatment of the false inference problem.

benefits (as the United States Department of Agriculture will soon be required to do[4]). If, pre-disclosure, the relevant risk was underestimated, then a disclosure mandate can efficiently reduce the underestimation. But if the effect of the disclosure-based inference is stronger, the disclosure mandate might substitute underestimation with overestimation, in which case the overall welfare assessment would be more difficult. Finally, if pre-disclosure, the relevant risk was overestimated, then a disclosure mandate will inefficiently increase the overestimation. In due course, we shall explore these possibilities in the context of cost-benefit analyses by regulators, which are mandatory at the federal level.[5]

The general arguments about false inferences from disclosure have implications for the intense and continuing debate about labeling of GM foods. In Europe, and increasingly in the United States, there is considerable public concern about genetically modified organisms (GMOs) and about food that contains them (GM food).[6] In response to this concern, governments have given serious consideration to the idea of requiring GM food labels, and some (including the United States) have already done so through legislation[7] (which may require implementing regulations, to which our analysis is relevant).

Opponents of GM food labels make a version of the false inference argument developed here: GM labels might affirmatively mislead some or many consumers, by leading them to believe, falsely, that GM foods pose risks to health or the environment, when in fact the scientific consensus is that no such risks exist.[8] Supporters of GM labels counter that consumers have a right to know what they are eating and that even a small risk of great harm justifies the disclosure mandate (since there is no conclusive scientific

[4] *See* Pub. L. No. 114-216, § 1, 130 Stat. 834 (2016) (codified at 7 U.S.C. § 1621 et seq.) (establishing the National Bioengineered Foot Disclosure Standard); under Executive Order 13563, the Department of Agriculture will have to catalogue the costs and benefits of the required regulations.

[5] *See* Exec. Order 13,563, 3 C.F.R. 215 (2012).

[6] *See* LABELING GENETICALLY MODIFIED FOOD: THE PHILOSOPHICAL AND LEGAL DEBATE (Paul Weirich ed., 2007).

[7] *See* Pub. L. No. 114-216, § 1, 130 Stat. 834 (2016) (codified at 7 U.S.C. § 1621 et seq.) (establishing the National Bioengineered Foot Disclosure Standard). Note that the requirement allows considerable flexibility for the regulated class, and the flexibility should significantly reduce compliance costs: Food producers can comply with a symbol or with text, but also with a barcode consumers can scan to obtain information on ingredients. Implementing regulations will be required to specify this requirement; for reasons to be elaborated, they should take account of the false inference problem both in the regulatory design and in the assessment of costs and benefits.

[8] See Ginger Pinholster, Food Labels Could "Mislead and Falsely Alarm Consumers" (2012), available at https://www.aaas.org/news/aaas-board-directors-legally-mandating-gm-food-labels-could-"mislead-and-falsely-alarm

proof that the risk of harm is zero).[9] By demonstrating the existence of the false inference problem and showing its adverse effects on consumer welfare, our analysis has implications for these arguments. As noted, the United States Department of Agriculture (USDA) is required to catalogue the costs and benefits of its regulations, including for the coming labels for GM foods. In doing so, the USDA must account for the false inference problem.

The remainder of this Article is as follows. Part II develops the general theoretical argument about false inferences and derives the conditions under which mandated disclosure is more or less likely to result in a false inference. Part III describes results from our survey study that confirm the theoretical predictions. Part IV considers the welfare implications of the false inference problem, with particular reference to the debate over GMO labeling. Part V is a brief conclusion.

II. Drawing False Inferences: Theory

We now present the False Inference theory. Though some of the discussion is a bit technical, the central intuitions and results are straightforward. Consumers hold some prior beliefs about the dangerousness of a product or a product feature (as we shall call it, for shorthand). Upon learning that the government decided to mandate (or not to mandate) disclosure of this feature, consumers update their beliefs. This updating, or inference, process can bring consumers' estimate of product risk closer to the actual, scientific risk measure. But under conditions that we specify, the updating process can drive the consumer's estimate further away from the objective truth.

This is what we call "false inference." Assume, for example, that the (fictional) Z25 ingredient increases the risk of colon cancer by 5% percentage points and that, pre-disclosure, consumers believe that Z25 increases the risk of colon cancer by only 1%. If the inference process, triggered by the disclosure mandate, increases consumers' estimate from the 1% prior to, say, 4%, then the inference is epistemologically desirable. On the other hand, if the inference process increases the consumer's estimate from 1% to 20%, or from (an accurate) 5% to 20%, we have a false inference problem. We show that false

[9] Nassim Nicholas Taleb et al., The Precautionary Principle (with Application to the Genetic Modification of Organisms) (Sept. 4, 2014) (unpublished manuscript), http://www.fooledbyrandomness.com/pp2.pdf. Consider in particular this suggestion: "More generally, engineered modifications to ecological systems (through GMOs) are categorically and statistically different from bottom up ones. Bottom-up modifications do not remove the crops from their long term evolutionary context, enabling the push and pull of the ecosystem to locally extinguish harmful mutations. Top-down modifications that bypass this evolutionary pathway unintentionally manipulate large sets of inter-dependent factors at the same time, with dramatic risks of unintended consequences. They thus result in fat- tailed distributions and place a huge risk on the food system as a whole." *Id.* at 10.

beliefs about the motivation behind the government's decision to mandate disclosure often result in false inference.

A. Framework of Analysis

Suppose that a consumer is choosing between two food products, A and B. Product A carries a government-mandated "Contains Z25" disclosure. Product B does not. The consumer wants to purchase healthy food products. But she is uncertain about the health effects of Z25. For expositional purposes, we assume that there is a particular health risk, *H* (measured in dollars), that is potentially associated with Z25. In this basic framework, the outcome is binary – either Z25 is harmful or not.[10]

The consumer knows that there are two possible reasons why the government would mandate a Z25 disclosure:

(1) The government would mandate disclosure because it believes that Z25 generates the risk *H*. Formally, the government receives one of two possible signals – either that Z25 is harmful or that Z25 is harmless. The accuracy of these signals is $p_G > \frac{1}{2}$. Namely, if Z25 is harmful, then there is a probability $p_G > \frac{1}{2}$ that the government gets a harmful signal and a probability $1 - p_G$ that the government gets a harmless signal. And if Z25 is harmless, then there is a probability $p_G > \frac{1}{2}$ that the government gets a harmless signal and a probability $1 - p_G$ that the government gets a harmful signal.

(2) The government would mandate disclosure regardless of its beliefs about the harmfulness of Z25 (indeed, the government would mandate disclosure even if it received no signal about whether Z25 is harmful or not). For example, the government believes that consumers should have as much information as possible about the ingredients in food products, regardless of any associated health risks. Or the government agency succumbs to interest-group pressure and mandates the disclosure.

The consumer attributes probability *q* to reason (1) and probability 1-*q* to reason (2).[11]

Before learning that the government mandates a Z25 disclosure, the consumer believed that Z25 generates the risk *H* with probability p_0. This is the consumer's prior. After learning that the government mandates a Z25 disclosure, the consumer updates her beliefs, according to a standard Bayesian updating process.[12] We next derive the

[10] We can extend this framework to allow for a continuous outcome variable measuring the probability that Z25 is harmful.

[11] This assumes that the two reasons are mutually exclusive. We can relax this assumption.

[12] For a summary of Bayesian updating, see C. R. GALLISTEL & ADAM PHILIP KING, MEMORY AND THE COMPUTATIONAL BRAIN: WHY COGNITIVE SCIENCE WILL TRANSFORM

consumer's updated, posterior probability that Z25 generates risk *H*. (The posterior probability is the consumer's final, post-updating probability estimate.) We denote this posterior probability p_1.

B. Motive for Disclosure is Harmfulness

It is instructive to begin with the special case where $q = 1$, namely, where the government would mandate disclosure if and only if it believes that Z25 is harmful. In this case, the Bayesian inference problem is depicted in the following tree diagram:

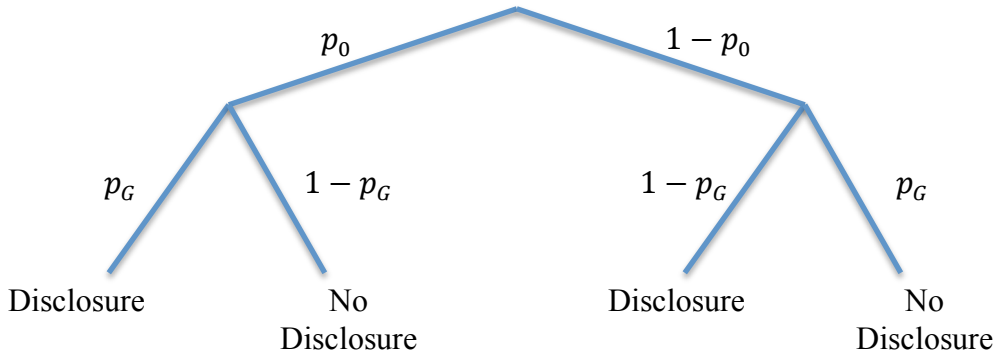

Figure 1: Bayesian Inference when Motive for Disclosure is Harmfulness

Pre-disclosure, the consumer believes that Z25 is harmful with probability p_0 (left side of the tree) and harmless with probability $1 - p_0$ (right side of the tree). If Z25 is harmful, then with probability $p_G > \frac{1}{2}$ the government will get a signal that Z25 is harmful and mandate disclosure; and with probability $1 - p_G$ the government will get a signal that Z25 is harmless and decline to mandate disclosure. If Z25 is harmless, then with probability $p_G > \frac{1}{2}$ the government will get a signal that Z25 is harmless and decline to mandate disclosure; and with probability $1 - p_G$ the government will get a signal that Z25 is harmful and mandate disclosure.

Knowing that the government decided to mandate a Z25 disclosure, the consumer would believe that Z25 is harmful with a (posterior) probability:

$$p_1 = \frac{p_0 \cdot p_G}{p_0 \cdot p_G + (1 - p_0) \cdot (1 - p_G)}$$

NEUROSCIENCE 27–42 (2010). Note that we are not claiming that people always act in a Bayesian manner. For an entertaining account of the origins of the now-widespread view that they do not, see MICHAEL LEWIS, THE UNDOING PROJECT (2016). We are so assuming for purposes of this simplified account.

In calculating her posterior, the consumer considers the likelihood that the government correctly mandates disclosure ($p_0 \cdot p_G$) and compares it to the overall likelihood that the government mandates disclosure – correctly or incorrectly ($p_0 \cdot p_G + (1 - p_0) \cdot (1 - p_G)$). The posterior is basically the share of correct disclosure mandates.

The mathematical formula for the posterior, p_1, captures several forces that intuitively affect the inferences that consumers draw from government-mandated disclosure: First, the consumer's prior has a strong effect on the posterior probability. A higher prior translates into a higher posterior:

$$\frac{dp_1}{dp_0} = \frac{p_G \cdot (1 - p_G)}{[p_0 \cdot p_G + (1 - p_0) \cdot (1 - p_G)]^2} > 0$$

In the extreme cases, where the consumer is certain about the health effects of Z25, the government's decision to mandate disclosure has no effect on the consumer's beliefs. The posterior is equal to the prior: $p_1 = p_0$. There are two extreme cases. The first occurs when, before learning whether or not the government mandates disclosure, the consumer was already certain that Z25 is harmful. Formally, this means that the consumer's prior was $p_0 = 1$. If the consumer was already 100% certain that Z25 is harmful, then any signal emanating from the government's decision to mandate disclosure will have no effect. Indeed, plugging $p_0 = 1$ into the posterior equation above, we get:

$$p_1 = \frac{1 \cdot p_G}{1 \cdot p_G + (1 - 1) \cdot (1 - p_G)} = 1$$

The second extreme case occurs when, before learning whether or not the government mandates disclosure, the consumer was already certain that Z25 is not harmful. Formally, this means that the consumer's prior was $p_0 = 0$. Again, the signal emanating from the government's decision to mandate disclosure will have no effect. Plugging $p_0 = 0$ into the posterior equation, we get:

$$p_1 = \frac{0 \cdot p_G}{0 \cdot p_G + (1 - 0) \cdot (1 - p_G)} = 0$$

The second force that affects the consumer's posterior is the accuracy of the government's signal, as measured by p_G. The more accurate the signal, the more upward updating would be expected – from p_0 to p_1. (We expect only upward updating, since updating is triggered by the government's decision to mandate disclosure – a decision that is motivated by a signal that Z25 is harmful.) Updating is captured by the difference

$$p_1 - p_0 = \frac{p_0 \cdot (1 - p_0) \cdot (2p_G - 1)}{p_0 \cdot p_G + (1 - p_0) \cdot (1 - p_G)}$$

(Note that, since $p_G > \frac{1}{2}$, we get $p_1 - p_0 > 0$, which implied upward updating.) We see that, as explained above, in the cases of pre-disclosure certainty, when $p_0 = 1$ or $p_0 = 0$, there is no updating, i.e., the government's signal has no effect on the consumer's posterior: $p_1 - p_0 = 0$, or $p_1 = p_0$. We also see that as the pre-disclosure uncertainty increases, namely, as p_0 moves away from $p_0 = 1$ or $p_0 = 0$, the effect of the government's signal increases.[13] Finally, we see that the level of updating increases in the accuracy of the government's signal:

$$\frac{d(p_1 - p_0)}{dp_G} = \frac{p_0 \cdot (1 - p_0)}{[p_0 \cdot p_G + (1 - p_0) \cdot (1 - p_G)]^2} > 0$$

In fact, the posterior is influenced not by the actual p_G, but by the perceived p_G. In particular, if consumers overestimate the accuracy of the government's signal, the level of updating will be higher.

C. Uncertainty About the Motive for Disclosure

We now reintroduce uncertainty about the government's motives, namely, with probability *q* the government mandates disclosure because it believes that Z25 is harmful and with probability 1-*q* the government mandates disclosure for other reasons that have nothing to do with the potential harmfulness of Z25. In this case, the Bayesian inference problem is depicted in the following tree diagram:

[13] The derivative of the difference $p_1 - p_0$ w.r.t. p_0 is:

$$\frac{d(p_1 - p_0)}{dp_0} = (2p_G - 1)\frac{(1 - p_0)^2 \cdot (1 - p_G) - {p_0}^2 \cdot p_G}{[p_0 \cdot p_G + (1 - p_0) \cdot (1 - p_G)]^2}$$

The derivative is increasing as we move upward from $p_0 = 0$ and as we move downward from $p_0 = 1$. (For subtle reasons, the difference $p_1 - p_0$ is not maximized at $p_0 = \frac{1}{2}$; for one, as p_0 increases there is less room for upward updating.)

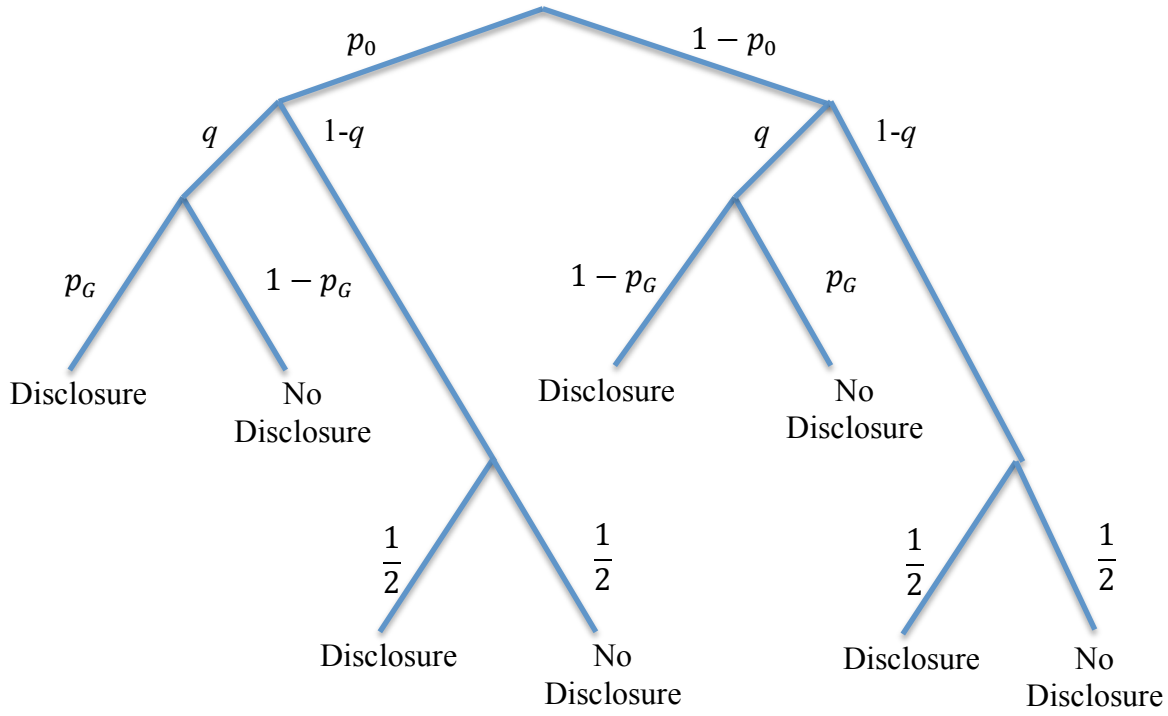

Figure 2: Bayesian Inference when Motive for Disclosure is Unclear

On the left side of the tree (where Z25 is harmful, according to the consumer's prior), with probability $p_G > \frac{1}{2}$ the government will get a signal that Z25 is harmful and with probability $1 - p_G$ the government will get a signal that Z25 is harmless. But these signals determine the government's decision whether to mandate disclosure only with probability q. With probability 1-q, the decision to mandate disclosure is completely uninformative, as captured by the 50%-50% probability distribution on the 1-q branch. On the right side of the tree (where Z25 is harmless), with probability $p_G > \frac{1}{2}$ the government will get a signal that Z25 is harmless and with probability $1 - p_G$ the government will get a signal that Z25 is harmful. But these signals determine the government's decision whether to mandate disclosure only with probability q.

Knowing that the government decided to mandate a Z25 disclosure, the consumer would believe that Z25 is harmful with a (posterior) probability:

$$p_1 = \frac{p_0 \cdot \left[qp_G + (1-q)\frac{1}{2}\right]}{p_0 \cdot \left[qp_G + (1-q)\frac{1}{2}\right] + (1-p_0) \cdot \left[q(1-p_G) + (1-q)\frac{1}{2}\right]}$$

Or:

$$p_1 = \frac{p_0 \cdot \left[qp_G + (1-q)\frac{1}{2}\right]}{q \cdot [p_0 \cdot p_G + (1-p_0) \cdot (1-p_G)] + (1-q) \cdot \frac{1}{2}}$$

In calculating her posterior, the consumer considers the likelihood that the government correctly mandates disclosure – based on an accurate signal that Z25 is harmful or randomly ($p_0 \cdot \left[qp_G + (1-q)\frac{1}{2}\right]$), and compares it to the overall likelihood that the government mandates disclosure – correctly or incorrectly ($q \cdot [p_0 \cdot p_G + (1-p_0) \cdot (1-p_G)] + (1-q) \cdot \frac{1}{2}$).

As in the simpler case, where the government's disclosure motives were clear, we find that a higher prior leads to a higher posterior. We also find that, in the extreme cases, when, pre-disclosure, the consumer is certain about the health effects of Z25, the government's decision to mandate disclosure has no effect on the consumer's beliefs. The posterior is equal to the prior: $p_1 = p_0$. As p_0 moves away from $p_0 = 1$ or $p_0 = 0$, the level of updating increases. Updating is captured by the difference

$$p_1 - p_0 = \frac{q \cdot p_0 \cdot (1-p_0) \cdot (2p_G - 1)}{q \cdot [p_0 \cdot p_G + (1-p_0) \cdot (1-p_G)] + (1-q) \cdot \frac{1}{2}}$$

Also, as in the simpler case, the level of updating increases in the accuracy of the government's signal:

$$\frac{d(p_1 - p_0)}{dp_G} = \frac{q \cdot p_0 \cdot (1-p_0)}{\left[q \cdot [p_0 \cdot p_G + (1-p_0) \cdot (1-p_G)] + (1-q) \cdot \frac{1}{2}\right]^2} > 0$$

And now that we have uncertainty about the motive for disclosure, we can also measure the effect of this uncertainty on the level of updating. As can be expected, the consumer will update more when disclosure is likely motivated by a signal that Z25 is harmful (i.e., when q is large) and the consumer will update less when disclosure is likely motivated by other reasons (i.e., when q is small). Formally, the level of updating is increasing in q:

$$\frac{d(p_1 - p_0)}{dq} = \frac{\frac{1}{2} \cdot p_0 \cdot (1-p_0) \cdot (2p_G - 1)}{\left[q \cdot [p_0 \cdot p_G + (1-p_0) \cdot (1-p_G)] + (1-q) \cdot \frac{1}{2}\right]^2} > 0$$

We can also identify two special cases. When $q = 0$, the disclosure is not informative and $p_1 = p_0$. When $q = 1$, we are back in the special case of a clear motive to mandate disclosure only if Z25 is harmful and the posterior is:

$$p_1 = \frac{p_0 \cdot p_G}{p_0 \cdot p_G + (1 - p_0) \cdot (1 - p_G)}$$

In fact, the posterior is influenced not by the actual q, but by the perceived q. In particular, if consumers overestimate the likelihood that the government's decision to mandate disclosure is motivated by a finding of harmfulness, the level of updating will be higher.

The preceding analysis is summarized in the following proposition.

Proposition:

1) Consumer's Prior

    a. The posterior, p_1, is increasing with the prior, p_0.

    b. In the extreme cases, where the consumer is certain about the health effects of Z25, the government's decision to mandate disclosure has no effect on the consumer's beliefs: (i) When $p_0 = 0$, the disclosure is irrelevant and $p_1 = p_0 = 0$; and (ii) When $p_0 = 1$, the disclosure is irrelevant and $p_1 = p_0 = 1$.

    c. As the pre-disclosure uncertainty increases, namely, as p_0 moves away from $p_0 = 1$ or $p_0 = 0$, the effect of the government's signal increases.

2) Accuracy of the Government's Signal

    a. The level of updating increases in the accuracy of the government's signal, p_G.

    b. Consumer overestimation of p_G results in excessive updating.

3) Government Motives

    a. The posterior, p_1, is increasing with q.

    b. When $q = 1$, we are back in the special case of a clear motive to mandate disclosure only if Z25 is harmful and

$$p_1 = \frac{p_0 \cdot p_G}{p_0 \cdot p_G + (1 - p_0) \cdot (1 - p_G)}$$

    c. When $q = 0$, the disclosure is not informative and $p_1 = p_0$.

    d. Consumer overestimation of q results in excessive updating.

III. Drawing False Inferences: Evidence

The most interesting, and the most important, theoretical predictions from Part II involve the effect of the government's motive, as perceived by consumers. We conducted a survey study to test these predictions. In particular, we set out to test how perceived motives affect the inferences that consumers draw from the government's decision to mandate disclosure. Our results confirm that when consumers believe that the government is motivated by new research, they draw stronger inferences from the government's decision to mandate disclosure. Our results also provide suggestive evidence that many consumers hold false beliefs about government motives – false beliefs that result in false inferences, producing welfare losses.

A. Methodology and Survey Design

The primary focus of this study is on the relationships between three variables: (1) a government regulatory Action, (2) the government's Motive for the Action and (3) consumer Risk perceptions about the subject of the Action. We examined two food ingredients as subjects of potential regulation: genetically modified organisms (GMOs) and Z25, a fictional synthetic preservative that is "sometimes added to make food stay fresher and last longer".

Government Action. In all cases, the government chose one of three actions: (1) make no requirement and let food producers decide whether to mention it on the label (No Action), (2) require that the label has a clear statement that the food contains this ingredient (Disclosure) or (3) require that the label includes a warning that the food contains this ingredient and it could pose at least some risk for some people (Warning). The distinction between disclosure and warning allows us to explore the implications of the disclosure's content and framing on perceived risk.

Government Motive. We examined three possible government Motives for the government's regulatory Action:

(1) People have a right to know what is in their food, and government regularly evaluates whether or not information should be added to food labels (RtK).

(2) Political pressure from food lobbying groups. Industry groups typically argue against adding information and consumer groups typically argue in favor.

(3) New research findings are published or reported about the safety or risk of an ingredient.

Some respondents were told the government's Motive immediately after learning of the Action (Manipulated Motive) while others were asked "Why do you think the

government decided to take this action?" and then to select one of the three Motives as the best reason (Perceived Motive). Comparing the effects of Manipulated and Perceived Motives allows for a direct test of the inference problem.

Perceived Risk. Respondents assessed risk on a 0 ("Definitely Won't Cause Harm") to 100 ("Definitely Will Cause Harm") scale. All respondents assessed Perceived Risk after learning about the Action and either learning about or assessing the Motive (Posterior Risk). Some respondents also assessed perceived risk before learning about the Action or Motive (Prior Risk). All respondents answered a purchase intent question after assessing their Posterior: "Given the action taken by the government, how does this affect your willingness to purchase foods that contain GMO (or Z25)?" on a five point scale from "Much Less Likely to Purchase" to "Much More Likely to Purchase".

Other Variables. Respondents also reported their age, gender, shopping frequency, and political views (on the widely used 7-pt "Extremely Conservative" to Extremely Liberal" scale). In the GMO conditions, respondents also indicated their level of knowledge about GMOs on a 5-pt scale from "No Knowledge" to "A great deal of knowledge".

Design. The GMO and Z25 studies were conducted on consecutive Wednesdays at the same time of day. The studies were parallel except that Prior Risk was assessed in the GMO study (since Z25 was fictional we did not assess priors). The additional heterogeneity due to previous GMO knowledge accounts for the larger samples in the GMO conditions. There were no significant differences between the samples on demographics, shopping frequency or political views. In each study, participants were randomly assigned to either the Manipulated or Perceived Motive condition. In the Perceived Motive conditions, the order in which Motives appeared was randomized across participants to remove any order effect.

	GMOs	**Z25**
Manipulated Motive	N = 418	N = 360
Perceived Motive	N = 534	N = 363

Table 1: Summary of studies

Prior and Posterior Risk measures were analyzed using double censored Tobit regression. Other variables were analyzed using a logistic regression.

Participants. We recruited 1,675 volunteers on two consecutive Wednesdays through Amazon Mechanical Turk (www.mturk.com) to participate in an online study.[14]

[14] Respondents on MTurk, though not a nationally representative sample, have been shown to be similar to respondents on most other survey platforms. Connor Huff & Dustin Tingley, *"Who Are These People?" Evaluating the Demographic Characteristics*

The study took an average of 3.8 minutes to complete and participants were paid a typical MTurk rate for participating. Participants were all U.S. residents, age 18+ (50% of sample in 30-49 years category), 47% female; political views covered the full spectrum from extremely conservative to extremely liberal, and on slightly were slightly more liberal than the national average (as is typical of MTurk samples).

B. Perceived Motives Affect Inferences from Government Actions

Consistent with the theoretical prediction, the effect of a disclosure mandate (or warning) depends on the particular Motive that consumers attribute to the government's decision. For Z25, we find that, when consumers believe that the disclosure mandate (or warning) was motivated by political pressure, the government's Action does not increase the Posterior risk. In contrast, when consumers believe that the disclosure mandate (or warning) was motivated by new research or by the government's belief that consumers have a right to know (RtK) what they are eating, the government's action results in a statistically significant increase in Posterior risk. The evident oddity, supporting the concern about false inferences, is that the Posterior risk *does not increase more when it is motivated by new research than when it is motivated by RtK*. These results are depicted in Figure 1 below.

and Political Preferences of MTurk Survey Respondents, RES. & POL., July–Sept. 2015, at 1.

Figure 1

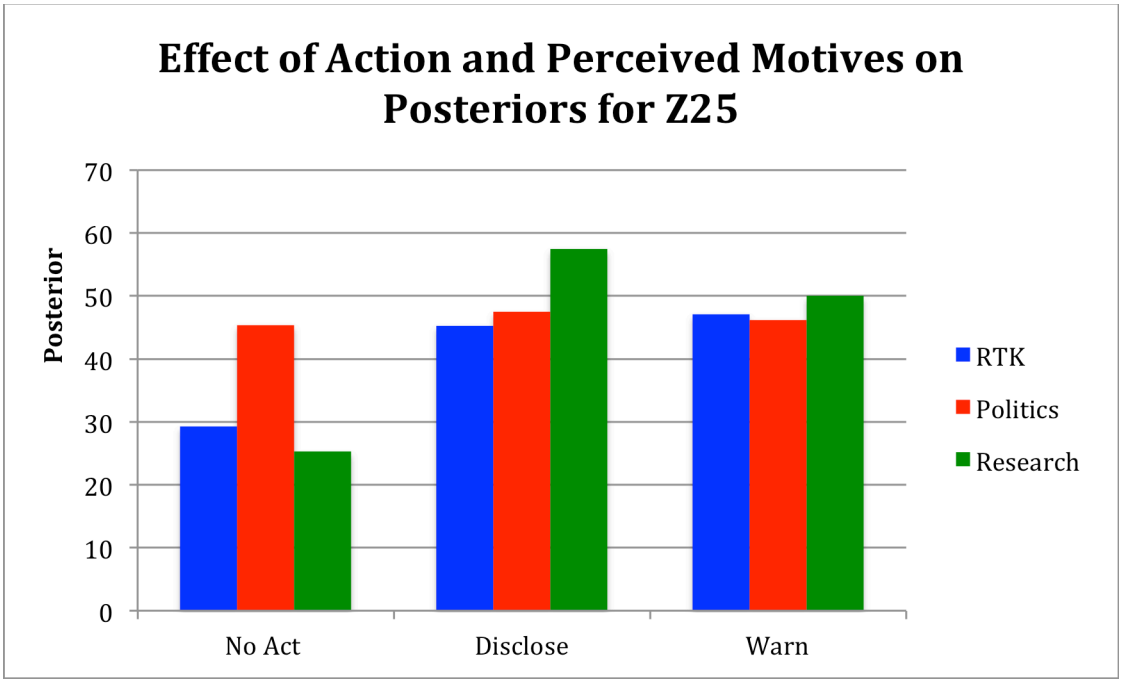

For GMOs, we obtain similar results when consumers believe that the disclosure mandate (or warning) was motivated by new research or by RtK. In these cases, the government's Action results in a statistically significant increase in Posterior risk. As expected, the strength of the government's action affects the Posterior risk. Specifically, the perceived risk is larger when the government mandates a strong warning, as compared to a weaker disclosure.

The results actually flip when consumers believe that the disclosure mandate (or warning) was motivated by political pressure. In these cases, the government's Action results in a small but statistically significant *decrease* in perceived risk. This latter result is surprising. A disclosure mandate (or warning) motivated by political pressure should not affect the Posterior risk.

Perhaps consumers are reasoning that if government acted in response to political pressure, the risk must be small; if it were large, political pressure would not be the reason for government's action. Or perhaps they are reasoning as follows: GM foods pose a greater competitive threat to non-GM foods, when GMOs are harmless. Therefore, the producers of non-GM foods exert more pressure on the government to act against GM foods, when they realize that GM foods are harmless. By this logic, consumers associate political pressure with evidence that GM foods are safe. In any event, consumers' risk perceptions are less affected by the government Action when they feel that the Motive is political.

The results are depicted in Figure 2 below.

Figure 2

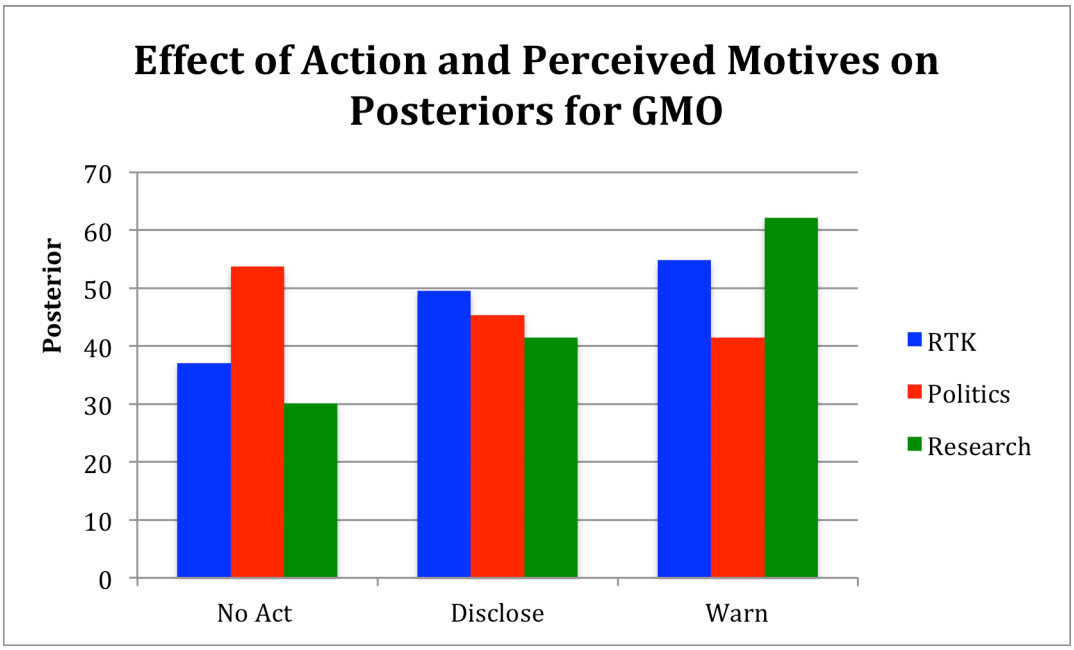

More generally, we find that Perceived Motive affects belief updating (Posterior – Prior Risk), as predicted by our theory. When consumers infer that the government's Action was motivated by political pressure, they do not update their beliefs about product risk (at least not on average; for GMOs they update downward). When consumers infer that the government's Action was motivated by new research, they update their beliefs about product risk upward when the government decides to mandate disclosure (or a warning) and downward when the government decides to take no action.

When consumers infer that the government's Action was motivated by RtK, the results are, as noted, more puzzling. They update their beliefs about product risk upward when the government decides to mandate a warning and downward when the government decides to take no action (there is no statistically significant updating when the government decides to mandate a weaker disclosure). These results are depicted in Figure 3 below.

Figure 3

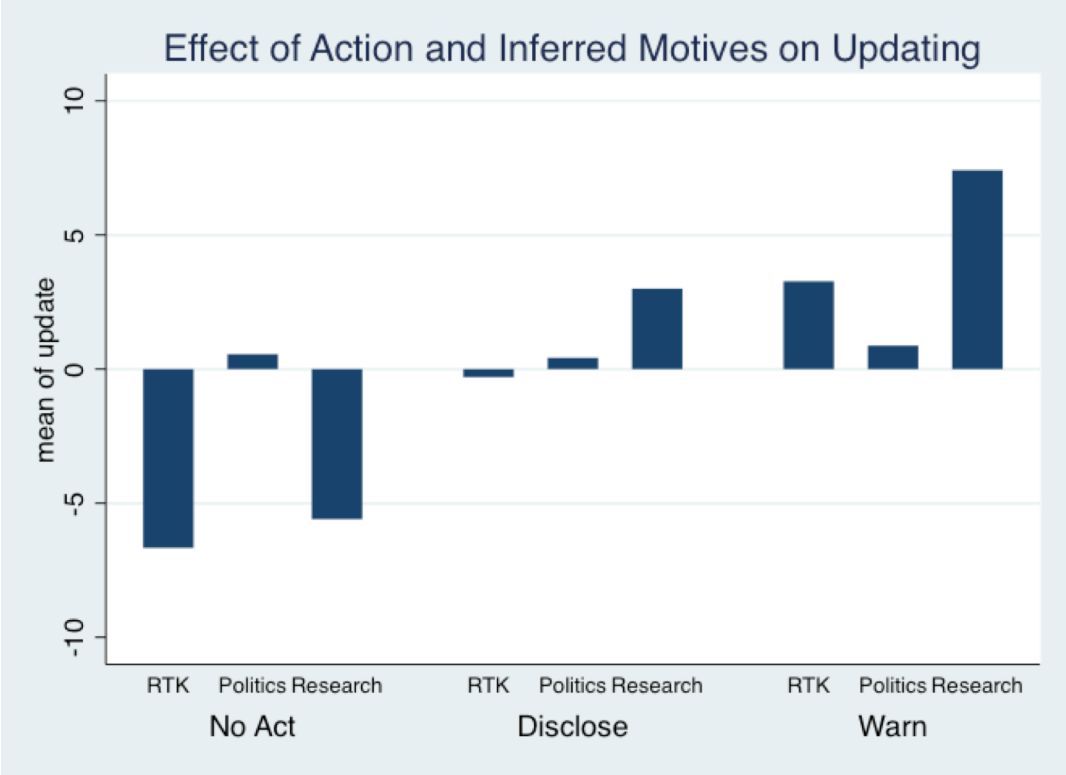

There is no doubt that the inferences that consumers draw about product risk affect their decisions whether to buy the product. Indeed, we are concerned about false inferences largely because they distort purchasing decisions. In our survey, we find that the government's Action to mandate disclosure (or a warning) reduces the reported likelihood of purchase – for both Z25 and GMOs. This effect is strongest when consumers infer that the government's Action was based on new research, weaker when consumers infer that it was based on RtK, and weakest (basically zero) when consumers infer that the government's decision was based on political pressure. These results are depicted in Figure 3 below.

Figure 4

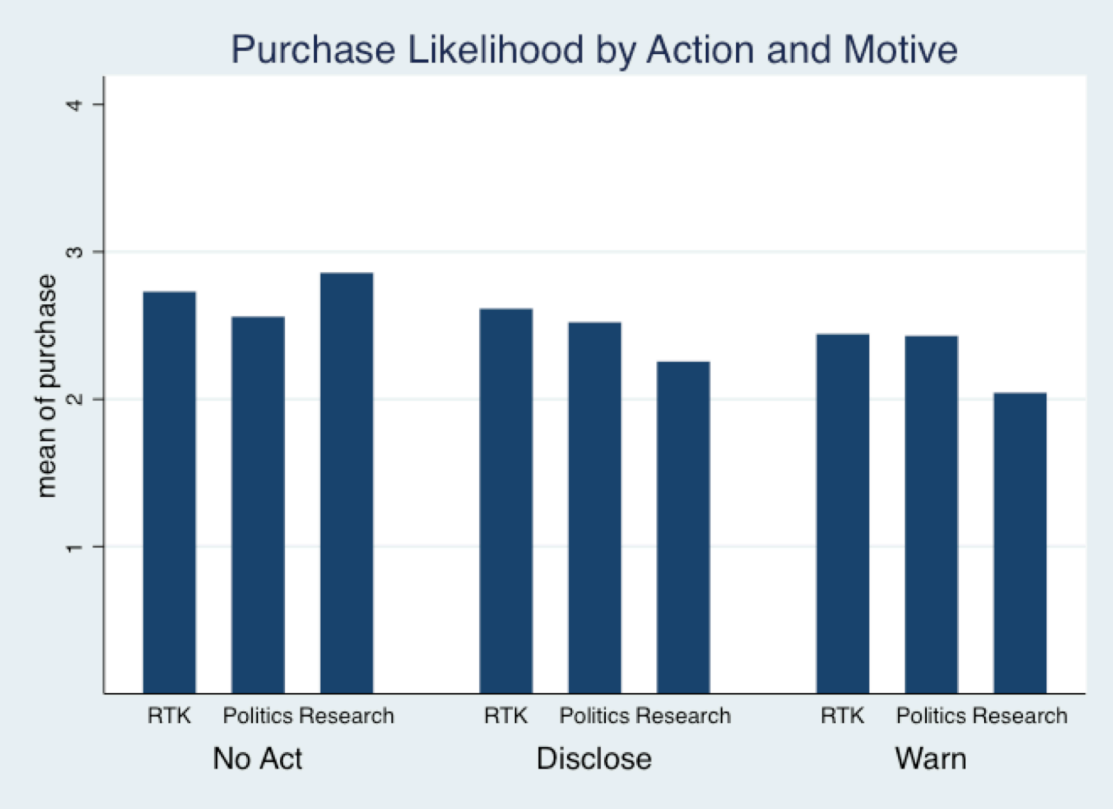

Consumer inferences clearly depend on their beliefs about the government's Motive. Faced with a mandated disclosure (or warning), consumers correctly increase their estimate of Posterior risk when they believe that the disclosure was based on new research. When they believe that the disclosure was driven by political pressure, consumers generally do not update their estimate of product risk – again a correct inference. (And for GMOs, they update downward.) But (and this is the central point for our purposes) when they believe that the disclosure was based on RtK, consumers seem to be making a false inference – they *increase* their estimate of product risk, whereas rational Bayesian decisionmakers would not update their estimate. It is reasonable to speculate that consumers do not accept a pure RtK motive; rather, they think that the government is motivated by a right to know when there is good reason to know, namely, when there is evidence that the product is harmful.

C. False Beliefs About Government Motives

The results reported in Section B confirm the theoretical prediction that consumer beliefs about the government's motive affect the inferences that consumers draw from the government's decision to mandate disclosure (or a warning). These results also have normative implications, which depend on the accuracy of consumers' beliefs about the government's motive. If these beliefs are accurate, then the inference that consumers draw from the disclosure mandate (or warning) will also be accurate. But if beliefs about the government's motive are inaccurate, then consumers will draw false inferences from the disclosure mandate (or warning).

The accuracy of consumer beliefs about the government's motives are hard to measure, largely because motives are themselves hard to measure. Still, our study provides suggestive evidence that a substantial group of consumers holds inaccurate beliefs. In particular, we know that the U.S. government, when mandating the GMO disclosure, was not motivated by new research about the risk of GMOs. Indeed, the research suggested that GM foods were harmless, at least in terms of human health.[15] Yet, when we asked our subjects about the government's motives, 16% answered that the GMO disclosure was motivated by new research on the harm that GM foods cause. To the extent that consumers conflate an RtK motive with a new research motive, as

[15] For a helpful discussion of scientific judgments in the United States, see Fred H. Degnan, *Biotechnology and the Food Label*, *in* LABELING GENETICALLY MODIFIED FOOD, *supra* note 4, at 17, 24–27. For a more recent overview, see U.S. FOOD & DRUG ADMIN., GUIDANCE FOR INDUSTRY: VOLUNTARY LABELING INDICATING WHETHER FOODS HAVE OR HAVE NOT BEEN DERIVED FROM GENETICALLY ENGINEERED PLANTS (2015), http://www.fda.gov/Food/GuidanceRegulation/GuidanceDocumentsRegulatoryInformation/LabelingNutrition/ucm059098.htm [hereinafter LABELING GUIDANCE]. As the FDA, notes in its guidance document:

> In the 1992 Policy, FDA stated that it was not aware of any information showing that bioengineered foods differ from other foods in any meaningful or uniform way, or that, as a class, foods developed by the new techniques present any different or greater safety concern than foods developed by traditional plant breeding (Ref. 5). Further, FDA concluded that the method of development of a new plant variety (including the use of new techniques such as rDNA technology) is generally not material information within the meaning of section 201(n) of the FD&C Act, and would not usually be required to be disclosed in the labeling for the food. This determination was reviewed and upheld by the court in *Alliance for Bio-Integrity v. Shalala*, 116 F. Supp. 2d 166, 178–79 (D.D.C. 2000) (finding that FDA's determination that genetic engineering, alone, is not a material fact that warrants food labeling was entitled to deference) (Ref. 10). Labeling provided by manufacturers on a wholly voluntary basis regarding whether a food was or was not bioengineered as described in this guidance is acceptable to FDA, provided that such labeling is truthful and not misleading. Some consumers are interested in the information provided in such labeling.
>
> *Id.*

suggested above, the 38% of subjects that chose RtK as the motive for the disclosure mandate were also drawing false inferences. In total, up to 54% of consumers are subject to the false inference problem.

D. The Effect of Prior Beliefs

Our theoretical model predicts that the false inference problem would be larger when consumers are uncertain about the relevant risks before encountering the mandated disclosure (or warning), and that the problem would be smaller when consumers start off with strong beliefs that the product is either safe or not. The strength of consumers' prior beliefs matters, because the extent of updating in response to the disclosure mandate (or warning) depends on these prior beliefs: weaker Prior Risk assessments (closer to 50%) result in more updating and stronger Prior Risk (closer to either 0% or 100%) result in less updating. Therefore, the false inferences problem is larger for consumers with weaker priors.

Our survey results confirm that stronger priors result in less updating. Specifically, 8.6% of respondents held very strong Priors – believing, with 100% certainty, that the product is harmful or believing, again with 100% certainty, that the product is safe. Of these respondents, only 13.1% updated their priors in response to the disclosure mandate (or warning). In contrast, 67.7% of the remaining respondents – those who were not so sure about the risk or safety of the product – updated their Prior Risk in response to the disclosure mandate (or warning). Figure 5 below depicts the average amount of updating as a function of the consumer's prior.

Figure 5

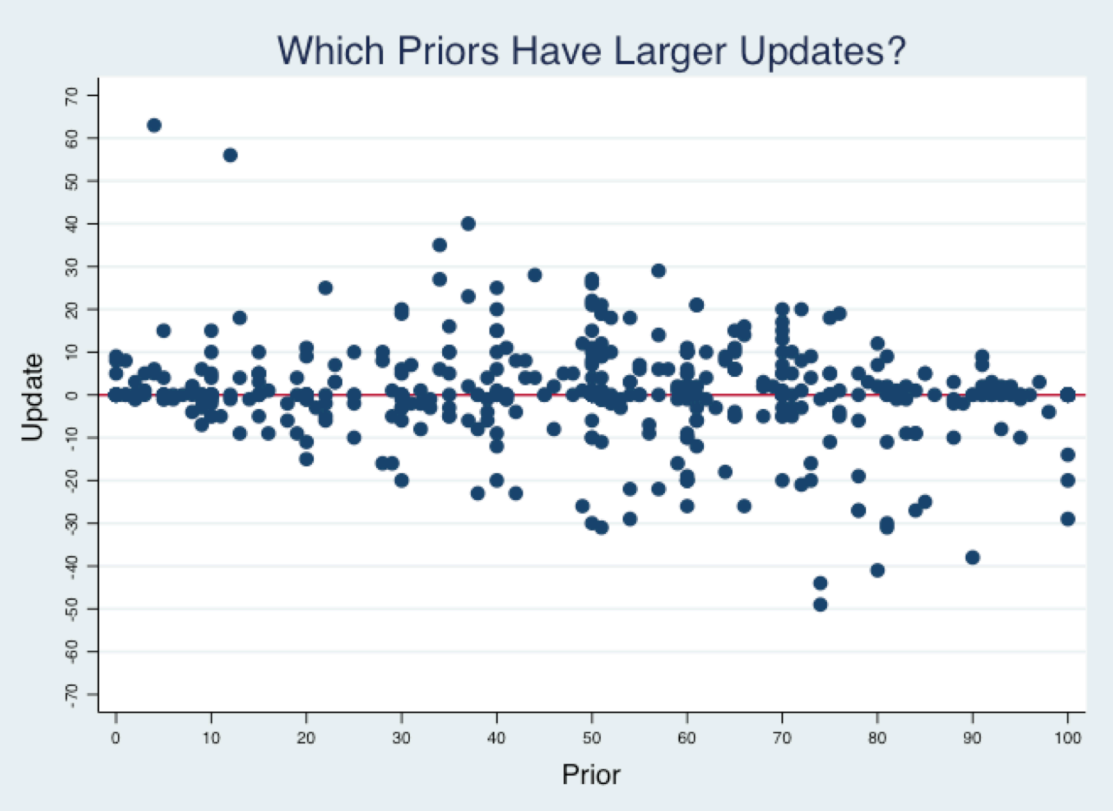

IV. Normative Implications

What are the welfare costs of false inferences? Putting quantitative assessments to one side, the qualitative answer is obvious. False inference leads to misperception of risk – consumers will either over- or underestimate the risk associated with an ingredient or characteristic of the product. Consumers who overestimate the risk might inefficiently decline to purchase the product. Instead these consumers will purchase an otherwise less attractive alternative (or decide not to purchase any product in this category), which would reduce the consumers' welfare. Consumers who underestimate the risk might inefficiently purchase the product, when in fact they should be purchasing a less risky alternative. Again, the result would be a reduction in the consumers' welfare.

When deciding whether to mandate disclosure, the government agency must compare the welfare cost from false inference to the welfare cost in the absence of the disclosure mandate. Under prevailing executive orders, agencies are required to catalogue

and to quantify (to the extent feasible) both the costs and the benefits of regulations,[16] and the costs of false inference must be noted and taken into account even if they are hard or impossible to quantify.

As we have seen, in the absence of a disclosure mandate consumers may be imperfectly informed, namely, they will suffer from under- or overestimation of risk. The question is whether the pre-disclosure misperception is better or worse than the post-disclosure misperception. To answer this question, it is helpful to distinguish between three cases, described in the following table.

Case	Pre-Disclosure Misperception	Post-Disclosure Misperception
1	Underestimation of risk	Less underestimation of risk
2	Underestimation of risk	Overestimation of risk
3	Overestimation of risk	More overestimation of risk

Table 2: Pre-disclosure vs. Post-disclosure misperceptions

In Case 1, pre-disclosure consumers suffer from underestimation of risk, and the disclosure mandate reduces the degree of underestimation. Consumers' risk estimate is now closer to the objectively correct estimate, and so their purchase decisions are more efficient and their welfare is higher. In Case 2, pre-disclosure consumers suffer from underestimation of risk, and post-disclosure they suffer from overestimation of risk. The purchase decisions are distorted in both cases – excessive purchase pre-disclosure and insufficient purchase post-disclosure. The effect on consumers' welfare is indeterminate in the abstract; empirical work would be needed to tell. In Case 3, pre-disclosure consumers suffer from overestimation of risk. and the disclosure exacerbates this bias. Consumers' risk estimate is now farther from the objectively correct estimate, and so their purchase decisions are less efficient and their welfare is lower.

Quantification is of course challenging. But at least in principle, the preceding analysis lends itself to direct implementation by regulators. Survey studies, like the one reported here, can provide information about the direction and even the magnitude of the misperception pre- and post-disclosure. On the basis of this information, the policy prescriptions in Case 1 and Case 3 are straightforward: Mandate disclosure in the former, but not in the latter. Case 2 poses a more difficult problem. In the absence of disclosure, underestimation of risk leads to overconsumption of the product, whereas disclosure results in overestimation of risk and thus underconsumption of the product.

It is important but insufficient to compare the magnitudes of the two misperceptions. Even if the underestimation is smaller than the overestimation, it may have a larger effect on consumption. Ideally, the regulator should assess the elasticity of demand with respect to risk perceptions (noting that this elasticity can be quite different for under- vs. overestimation of risk). As noted, surveys might provide relevant

[16] *See* Exec. Order 13,563, 3 C.F.R. 215 (2012).

information. If sufficient information cannot be obtained, regulators should, consistent with standard practice, candidly acknowledge uncertainties. When quantification is not possible and significant uncertainties remain, regulators have some helpful strategies, including the use of lower and upper bounds.[17] In some cases, it is imaginable that existing knowledge will make it difficult to decide whether the benefits of disclosure justify the costs – though as we will explain, corrective steps could reduce the risk of false inference and thus increase net benefits.

B. The Problem of Moral Preferences

The preceding analysis focuses on false inferences and the misperceptions of risk that they create. These misperceptions are troubling whenever (at least some) consumers care about the relevant risk and thus make purchasing decisions based on their potentially biased estimate of this risk. We acknowledge, however, that some consumers may have other, morally-laden preferences and that these preferences can affect the desirability of a disclosure mandate.

For example, a certain ingredient or characteristic of a product might be inherently objectionable to some consumers.[18] For example, many consumers do not want to eat tuna that actually contains dolphin meat.[19] Many other consumers want to purchase products that are "made in America," for social or moral reasons.[20] At least some consumers apparently find the idea of GM foods to be intrinsically objectionable, regardless of the absence of risk to health or to the environment. By their own lights, these consumers benefit from knowing whether a product contains GMOs, and they would be willing to pay something to obtain that knowledge. To that extent, they would benefit from a disclosure mandate, just as they would benefit from learning about a health risk. They would not be harmed from the false inference problem, simply because they do not much care about, and thus do not draw inferences about, the risk of harm to health or to the environment.[21] It is of course an empirical question whether many consumers

[17] See Cass R. Sunstein, *The Limits of Quantification*, 102 Cal L Rev 1369 (2014); Cass R. Sunstein, *On Mandatory Labelling, With Special Reference to Genetically Modified Food*, U. Pa. L. Rev. (forthcoming 2017).

[18] For a statement and critical evaluation, see Gary Comstock, *Ethics and Genetically Modified Foods*, *in* ETHICS AND EMERGING TECHNOLOGIES 473 (Ronald L. Sandler ed., 2014), http://hettingern.people.cofc.edu/150_Spring_2015/Comstock_Ethics_and_Genetically_Modified_Foods.pdf.

[19] *See Dolphin Safe Fishing*, INT'L MARINE MAMMAL PROJECT, http://savedolphins.eii.org/campaigns/dsf.

[20] For a recognition of this point, see Shaun Zinck, *What It Takes to Qualify for the "Made in America" Label*, QUALITY LOGO PRODUCTS, https://www.qualitylogoproducts.com/blog/the-made-in-america-label-rules/

[21] Another group of consumers might find GM foods inherently objectionable and also care about the risk of harm to health or to the environment. Their preference profile

actually have this preference; it is possible that in some cases, those who appear to hold moral preferences of this kind actually are concerned about health and environmental risks.

Another type of preference is the preference to know as much as possible about the products that you are purchasing.[22] Consumers with such a preference would also benefit from a disclosure mandate.[23] However, it is not clear whether many, or even any, consumers actually have that preference, which seems barely intelligible. In principle, there is a great deal to know about cell phones, hamburgers, television sets, and automobiles. Much of that information is irrelevant to what rational (or boundedly rational) consumers care about. Even if it were costless to supply that information, it is not costless to process it.

In any case, an important limit to the argument derives from the "as much as possible" qualifier. As consumers, we inevitably make decisions under conditions of imperfect information, and that is not a problem. In neoclassical terms, there is an optimal stopping point in our search for information – partly because the costs of search can be high and partly because the benefits of search can be low.[24] In behavioral terms, there is a limit on the amount of information that we can effectively absorb.[25] Therefore, learning about GMO contents might crowd out information about, say, sugar or salt. A general preference to know does not in and of itself justify a disclosure mandate. We suspect, though we do not know, that those who say that there is a RtK do not support that right in general, but only in particular contexts, in which they suspect that there is reason for concern from the standpoint of environment, health, or morality.

C. Counteracting False Inferences

From the point of view of regulators, it is important to ask whether false inferences might be combated with more disclosure or with improved framing. If so, the welfare costs would be reduced or avoided. One question is whether voluntary disclosure can be expected to provide a corrective. Another question is whether supplemental disclosure might be mandated. We briefly discuss both possibilities.

---

is multidimensional. For this group of consumers, we face a tradeoff — a disclosure mandate would provide a benefit (with respect to the dimension of their preferences that finds GM foods inherently objectionable), but might also impose a cost (with respect to. the dimension of their preferences that cares about the risk of harm to health or to the environment).

[22] For relevant discussion, see Edna Ullmann-Margalit, *On Not Wanting To Know*, *in* REASONING PRACTICALLY 72 (Edna Ullmann-Margalit ed., 2000).

[23] For a vivid statement to this effect, see *Why Label?*, JUST LABEL IT!, http://www.justlabelit.org/right-to-know-center/right-to-know/.

[24] George J. Stigler, *The Economics of Information*, 69 J. POL. ECON. 213 (1961).

[25] *See* OMRI BEN-SHAHAR & CARL E. SCHNEIDER, MORE THAN YOU WANTED TO KNOW: THE FAILURE OF MANDATED DISCLOSURE (2014).

Consider a mandate that requires all sellers who use Z25 in their products to include a Z25 label on their packaging, and assume that this disclosure mandate is not based on evidence that Z25 is harmful to consumers. Sellers of Z25 products would have a clear incentive to educate consumers and convince them that Z25 is harmless (or, at least, that there is no evidence to the contrary). The question may not be hypothetical. In the United States, sellers of GM food might want to engage in an advertising campaign or add a disclosure: "There is no evidence that GM food is hazardous to human health."

For two reasons, however, such voluntary disclosure might not always occur. First, it might be futile or even counterproductive. A statement that GM food has not been found to be hazardous to human health places "GM food" and "hazardous" in the same sentence. Many consumers might not be assured by that kind of proximity; their concern might even grow. Rational sellers would take that possibility into account. Second, the necessary information triggers a collective action problem: A single seller will be reluctant to invest millions of dollars in an advertising campaign to educate consumers about the safety of GM food, if all sellers of GM food would reap the benefits of such a campaign.[26] Perhaps an industry group could solve this collective action problem, or perhaps a simple label, including a corrective statement, would have benefits in excess of costs (assuming the proximity problem might be solved).

Should a federal agency mandate some kind of corrective disclosure, to combat the risk of false inferences? For instance, if there is concern that a GMO disclosure would lead to overestimation of risk, the government can mandate a supplemental disclosure: "The best scientific evidence suggests that GMOs carry no health risks." On plausible assumptions, such a mandate would make sense: It would reduce the welfare costs of false inferences without imposing costs on those who draw such inferences (assuming the costs of the disclosure are themselves modest). One question is whether the proximity problem just identified would mean that the mandate would be futile or counterproductive. Another question is the magnitude of the welfare loss from false inferences, and whether it can be reduced or eliminated through voluntary action. If the loss is large, if voluntary action is insufficient, and if the loss can be successfully combatted through a corrective mandate, such a mandate would deserve consideration.

V. Conclusion

Disclosure mandates are pervasive. In this Article, we have analyzed – theoretically and empirically – an unexplored and potentially significant cost of mandatory disclosure: the false inference problem. In some cases, many consumers will hear a loud signal of "DANGER!" even though the evidence justifies no such signal. The central implication is simple: Policymakers should incorporate an understanding of the false inference problem into their analysis of whether and how to mandate disclosure

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

(including their analysis of costs and benefits). When false inferences are made, a disclosure requirement might turn out to be affirmatively harmful.

In extreme cases, large numbers of consumers will make false inferences, and hence any such mandate may impose costs in excess of benefits. A disclosure requirement will make people less informed. In less extreme cases, the false inference problem is limited to a relatively small subset of consumers; the welfare cost decreases net benefits, and regulators should acknowledge that fact, but disclosure might nonetheless be a good idea on balance. In all cases, policymakers should consider whether to authorize or mandate additional disclosures so as to reduce the risk of false inferences.