

Drawing false inferences from mandated disclosures

OREN BAR-GILL*

Harvard Law School, Cambridge, MA, USA

DAVID SCHKADE

University of California, San Diego, Rady School of Management, La Jolla, CA, USA

CASS R. SUNSTEIN

Harvard Law School, Cambridge, MA, USA

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.

Submitted 6 November 2017; accepted 14 November 2017

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.

* Correspondence to: Oren Bar-Gill, Harvard Law School, Cambridge, MA, USA. Email: bargill@law.harvard.edu

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 (RtK) 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) 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.

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 RtK 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. 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 RtK.

In this paper, we analyze the factors that influence the inferences that consumers draw from a disclosure mandate, both theoretically and empirically. 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 RtK what they are buying. Puzzlingly,

but importantly, consumers who attribute to the government a RtK motive 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 genetically modified organism (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 prior, pre-disclosure beliefs 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.¹

The general arguments about false inferences from disclosure have implications for the intense and continuing debate about the labeling of genetically modified (GM) foods. In Europe, and increasingly in the United States, there is considerable public concern about GMOs and about food that contains them (GM food) (see Weirich, 2007). 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 (see Pub. L. No. 114–216, § 1, 130 Stat. 834 (2016) (codified at 7 U.S.C. § 1621 et seq.)).

This paper adds to the growing literature on the benefits and costs of mandated disclosure. For a skeptical view on the overall merits of mandated disclosure as a regulatory technique, see Ben-Shahar and Schneider (2014). For a less skeptical view, see Loewenstein *et al.* (2014) and Bar-Gill (2012). More specifically, this paper formalizes and provides empirical support for a claim made by opponents of GM food labels – that GM labels might

1 We note, however, that, given the prevalence of pre-disclosure and pre-inference misperceptions, false inferences are not necessarily welfare reducing. For example, if, pre-disclosure, the relevant risk was underestimated, then a disclosure mandate can efficiently reduce the underestimation.

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 (see, e.g., Pinholster, 2012). Most closely related to our work is the unpublished manuscript by Zhang (2014). 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 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 a made-up ingredient (Z25).

The remainder of this paper is organized as follows. The next section 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. (The more technical analysis is relegated to the Online Appendix.) The subsequent 'evidence' section describes results from our survey study that confirm the theoretical predictions. The final section offers some concluding remarks.

Drawing false inferences: theory

We now present the False Inference theory. 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' estimates 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." In particular, we show that false beliefs about the motivation behind the government's decision to mandate disclosure often result in false inference.

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.²

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 .
- (2) The government would mandate disclosure regardless of its beliefs about the harmfulness of Z25. 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).³

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, arriving at a 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 .

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, knowing that the government decided to mandate a Z25 disclosure increases the consumer's (posterior) probability of harm. Several forces 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. More interesting, the degree of pre-disclosure certainty, as reflected in the prior, affects the inferences drawn from mandated disclosure. In the extreme, 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 and the posterior is equal to the prior. There are two extreme cases. The first occurs when, before learning whether or not the government mandates disclosure,

² We can extend this framework to allow for a continuous outcome variable measuring the probability that Z25 is harmful.

³ This assumes that the two reasons are mutually exclusive. We can relax this assumption.

the consumer was already certain that Z25 is harmful. In this case, $p_1 = p_0 = 1$. 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. In this case, $p_1 = p_0 = 0$. Beyond these extreme cases, we find that as the pre-disclosure uncertainty increases, namely, as p_0 moves away from $p_0 = 1$ or $p_0 = 0$, the effect of mandated disclosure on consumer beliefs increases.

The second force that affects the consumer's posterior is the accuracy of the government's information about the harm from Z25. The more accurate the information, 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.) In fact, the posterior is influenced not by the actual accuracy of the government's information, but by the perceived accuracy of the information. In particular, if consumers overestimate the accuracy of the government's information, the level of updating will be higher.

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. 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, and 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. Also, as in the simpler case, the level of updating increases in the perceived accuracy of the government's information about the harm from Z25.

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 information 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). 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. 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, which is formalized in the Online Appendix, is summarized in the following proposition.

Proposition: The consumer's posterior probability, p_1 , is determined by the following factors:

- 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. As the pre-disclosure uncertainty increases, the effect of the government's decision to mandate disclosure increases.
- 2) Accuracy of the Government's Information: The posterior, p_1 , is increasing with the perceived accuracy of the government's information about the harm from Z25.
- 3) Government Motives: The posterior, p_1 , is increasing with the perceived likelihood that the disclosure mandate was motivated by the government's belief that Z25 is harmful.

Drawing false inferences: evidence

The most interesting, and the most important, theoretical predictions from the previous section 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. (We also test predictions about the effect of consumers' prior beliefs.)

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: 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 this ingredient 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 RtK 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 asked 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 seven-point “Extremely Conservative” to “Extremely Liberal” scale). In the GMO conditions, respondents also indicated their level of knowledge about GMOs on a five-point 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. [Table 1](#) lists the four studies that we conducted and the numbers of participants in each study.

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

Participants

We recruited 1675 volunteers on two consecutive Wednesdays through Amazon Mechanical Turk (MTurk; www.mturk.com) to participate in an online study.⁴ The study took an average of 3.8 minutes to complete and participants were paid a typical MTurk rate for participating. Participants were all US residents, aged 18+ (50% of sample in 30–49 years of age category) and 47% female; political views covered the full spectrum from extremely conservative to extremely liberal, and were slightly more liberal than the national average (as is typical of MTurk samples).

⁴ Respondents on MTurk, though not a nationally representative sample, have been shown to be similar to respondents on most other survey platforms (Huff & Tingley, 2015).

Table 1. Summary of studies.

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

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 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.

For GMOs, we obtain similar results when consumers believe that the disclosure mandate (or warning) was motivated by new research or by RtK. In




Figure 1. Effect of Action and perceived Motives on Posteriors for Z25.
RTK = right to know.

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.

More generally, we find that perceived Motive affects belief updating (Posterior – Prior Risk), as predicted by our theory. When consumers infer




Figure 2. Effect of Action and perceived Motives on Posteriors for genetically modified organisms.

RTK = right to know.

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](#).

There is no doubt that the inferences that consumers draw about product risk affect their decisions on 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




Figure 3. Effect of Action and inferred Motives on updating.
RTK = right to know.

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 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.) 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 decision-makers 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 RtK when there is good reason to know; namely, when there is evidence that the product is harmful.




Figure 4. Purchase likelihood by Action and Motive.
RTK = right to know.

False beliefs about government motives

The results reported in the prior subsection 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 US 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 (see Degnan, 2007; US Food and Drug Administration, 2015). 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 a RtK motive with a new research motive, as 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.

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 assessments (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




Figure 5. Which Priors have larger updates?

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 depicts the average amount of updating as a function of the consumer's Prior.

Concluding remarks

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 (see Comstock, 2014). For instance, many consumers do not want to eat tuna that actually contains

dolphin meat. Many other consumers want to purchase products that are “made in America,” for social or moral reasons. 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.⁵ It is of course an empirical question whether many consumers 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.

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.

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

⁵ 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 is multidimensional. For this group of consumers, we face a trade-off – 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).

reassured 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 (see Beales *et al.*, 1981). 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.

Supplementary Material

To view supplementary material for this article, please visit <https://doi.org/10.1017/bpp.2017.12>.

References

- Bar-Gill, Oren (2012), *Seduction by Contract: Law, Economics and Psychology in Consumer Markets*, Oxford, UK: Oxford University Press.
- Beales, Howard, Richard Craswell, and Steven C. Salop (1981), ‘The Efficient Regulation of Consumer Information’, *Journal of Law and Economics*, 24: 491–539.
- Ben-Shahar, Omri and Carl E. Schneider (2014), *More Than You Wanted to Know: The Failure of Mandated Disclosure*, Princeton, NJ: Princeton University Press.
- Comstock, Gary (2014), ‘Ethics and Genetically Modified Foods’, in Ronald L. Sandler (ed.) *Ethics and Emerging Technologies*, UK: Palgrave MacMillan.
- Degnan, Fred H. (2007), ‘Biotechnology and the Food Label’, in Weirich, Paul (ed.), *Labeling Genetically Modified Food: The Philosophical and Legal Debate*, Oxford, UK: Oxford University Press.

- Huff, Connor and Dustin Tingley (2015), “‘Who Are These People?’ Evaluating the Demographic Characteristics and Political Preferences of MTurk Survey Respondents’, *Research & Politics*, 2015: 1–12.
- Loewenstein, George, Cass R. Sunstein, and Russell Golman (2014), ‘Disclosure: Psychology Changes Everything’ *Annual Review of Economics*, 6: 391–419.
- Pinholster, Ginger (2012), ‘AAAS Board of Directors: Legally Mandating GM Food Labels Could “Mislead and Falsely Alarm Consumers”’ (available at <https://www.aaas.org/news/aaas-board-directors-legally-mandating-gm-food-labels-could-“mislead-and-falsely-alarm”>).
- US Food and Drug Administration (2015), *Guidance for Industry: Voluntary Labeling Indicating Whether Foods Have or Have Not Been Derived from Genetically Engineered Plants*.
- Weirich, Paul (ed.) (2007), *Labeling Genetically Modified Food: The Philosophical and Legal Debate*, Oxford, UK: Oxford University Press.
- Zhang, Juanjuan (2014), ‘Policy and Inference: The Case of Product Labeling. Unpublished manuscript’, http://jjzhang.scripts.mit.edu/docs/Zhang_2014_GMO.pdf.