

Policy and Inference: The Case of Product Labeling

Juanjuan Zhang*

September 4, 2013

Abstract

Consumers may infer a policy maker's private information about relevant economic variables from its policy choice. I study this inference effect and its policy implications in the context of mandatory product content disclosure. I find that mandatory disclosure does not always benefit consumers. For example, by observing the government's requirement to label genetically modified organisms (GMOs) in food products, consumers may infer that GMO safety is of enough concern, which leads to insufficient consumption. On the other hand, the lack of mandatory disclosure may send a falsely optimistic signal of GMO safety to consumers, which leads to excessive consumption. Optimal disclosure policy should take into account both the transparency of product content and consumer inference of content quality. Interestingly, higher disclosure implementation cost may broaden the scope of mandatory disclosure and may even increase social welfare. An experiment shows that emphasizing GMO disclosure does worsen consumers' perception of GMO safety.

Key words: observational learning, signaling, mandatory disclosure, product labeling, public policy

*Massachusetts Institute of Technology, Cambridge, MA 02139, jjzhang@mit.edu. The author gratefully acknowledges helpful comments from Alessandro Bonatti, Anthony Dukes, Dmitri Kuksov, Jeanine Miklós-Thal, David Tannenbaum, and J. Miguel Villas-Boas; from seminar participants at Cheung Kong Graduate School of Business, Emory University, Tsinghua University, and the University of California, Los Angeles; and from attendees of the 2012 MIT Applied Theory Summer Camp, the 2012 Northeast Marketing Conference, and the 2013 Frank M. Bass UTD-FORMS Conference.

1 Introduction

Government policy intervention can be an effective way to influence public choice and improve social welfare. This article argues that the very enactment of a policy may signal the policy maker's private information about relevant economic variables. The making of public policies often requires extensive research, whereas the public normally does not have this level of expertise. What the public can do, however, is infer the state of the economy from the observed policy. Policy making should take this inference effect into account.

For example, consider the mandatory labeling of genetically modified food, a policy proposal that has generated a heated debate in the United States (e.g., Harmon and Pollack 2012).¹ The use of genetically modified organisms (GMOs) in agricultural products is prevalent in the United States, although its health consequences have drawn much questioning (National Research Council 2004). Proponents of mandatory disclosure emphasize "transparency," arguing that consumers have the right to know the content of their food to make a well-informed purchase decision.² Opponents highlight two points. First, the implementation of mandatory labeling is costly (e.g., Ramessar et al. 2010). Second, although the risk assessment of GMOs is not conclusive, studies to date have not found significant evidence of health hazards (e.g., OECD 2000; Harmon and Pollack 2012).

The last point is rather delicate. If the justification for not mandating labeling, in view of its implementation cost, is that genetically modified food is sufficiently safe, then mandatory labeling itself could be interpreted as a signal that genetically modified food is hazardous. Indeed, concerns have been expressed that mandatory labeling would "create a falsely negative impression about the safety of genetically modified foods" and would "scare consumers away from genetically engineered products that science has found to be perfectly safe."³

The impact of mandatory labeling is complex because there are two sources of information asymmetry. First, consumers typically do not know whether a product is genetically modified without the benefit of labeling (Golan et al. 2001). Second, consumers tend to be less informed

¹As of the time of this article, labeling of genetically modified foods is not mandated by the Food and Drug Administration (FDA). At California's state-wide election in November 2012, residents voted on Proposition 37, which mandates the disclosure of a broad range of foods containing genetically modified ingredients. The Proposition was rejected.

²See <http://justlabelit.org> and <http://www.carighttoknow.org> for examples of campaigns that advocate consumers' entitlement to genetically modified product information.

³Source: <http://www.noprop37.com/news/editorial-more-food-labels-not-this-not-now>.

about the quality of genetically modified products than the policy maker. Since the policy maker's choice of whether to address the first information asymmetry depends on its private information about genetically modified products' quality, the policy decision itself influences the extent of the second information asymmetry.

To formalize this intuition and derive the policy implications, I study a game between a policy maker and consumers who choose whether to adopt a product. Product content is either "innocuous" (e.g., derived from natural methods) or "risky" (e.g., derived from genetic modification). Consumers do not observe product content unless the policy maker mandates disclosure. The policy maker observes the quality of the risky product, and chooses whether to mandate the disclosure of product content. The policy maker's objective is to maximize social welfare, which in this context is equivalent to maximizing consumer welfare net of any implementation cost of disclosure. The disclosure policy generates two informational effects: a *transparency effect* that resolves consumers' information asymmetry regarding product content, and an *inference effect* that influences consumers' beliefs about the risky product's quality.

To disentangle these two effects, I start with a benchmark scenario in which consumers observe the risky product's quality. As expected, mandatory disclosure benefits consumers: it helps those concerned about product quality avoid a risky product out of precaution and adopt an innocuous product with confidence. Indeed, unless the risky product is equivalent to the innocuous product, disclosure always improves consumer welfare, and the magnitude of this improvement decreases with the risky product's quality. It follows that, if implementation is costless, the policy maker should always mandate disclosure. If implementation is costly but not prohibitive, the policy maker should mandate disclosure if and only if the risky product's quality is sufficiently low. These results parallel the arguments for or against genetic modification labeling, namely, consumers' "right to know," implementation cost, and GMO safety.

The key assumption underlying this article, however, is that consumers are less informed than the policy maker about the risky product's quality. There are several ways to understand this assumption. Consumers normally do not have access to the level of scientific product testing that the policy maker is able to organize. Even if the policy maker releases the test results, many consumers do not have the expertise to interpret them correctly (Golan et al. 2001). Moreover, consumers may be overloaded with information. A Google search of the

phrase “genetically modified food safety” yields 2,710,000 results (as of November 2012) that represent divergent views. It is thus plausible that the disclosure policy itself serves as a relatively salient and unambiguous cue that consumers can rely on to infer content quality.

When this inference effect is taken into consideration, some of the conventional wisdom regarding disclosure needs to be revised. In particular, the effect of mandatory disclosure on consumer welfare is not always positive. Even if disclosure is costless to implement, there exists no perfect sequential equilibrium in which the policy maker always mandates disclosure. To see the reason, suppose the policy maker always mandates disclosure. Consumers learn nothing about the risky product’s quality from the policy. If the risky product is actually of high quality, some consumers, unaware of its true quality, will falsely reject the product. The policy maker thus has the incentive to convey high quality to consumers by deviating to nondisclosure. Consumers also have reason to believe that such deviation signifies high product quality because, intuitively, disclosure becomes unimportant if the risky product is reasonably safe.

Similar to the benchmark case, the equilibrium disclosure policy can be characterized by a threshold rule: the policy maker will mandate disclosure if and only if the risky product’s quality is sufficiently low. The quality threshold, however, is in general different from that in the benchmark case. Moreover, the quality threshold, and thus the “scope” of risky products subject to mandatory disclosure, can increase with disclosure implementation cost. This counterintuitive result occurs when the inference effect of disclosure dominates the transparency effect – as the quality threshold increases, the policy maker’s choice to mandate disclosure becomes a weaker sign of low quality, and its choice to not mandate disclosure becomes a stronger sign of high quality. This mitigates consumers’ Type I error in consumption (e.g., insufficient product adoption) following disclosure, and exacerbates their Type II error (e.g., excessive product adoption) following nondisclosure. Both effects increase the benefit of disclosure to consumer welfare, which serves to justify a higher disclosure implementation cost.

Finally, the effect of disclosure implementation cost on social welfare is subtle because of the inference effect. In the benchmark case, more costly implementation always (weakly) decreases social welfare. When consumers infer the risky product’s quality from the disclosure policy, however, more costly implementation can increase social welfare by influencing the signal strength of the disclosure policy. For example, suppose that the quality threshold decreases with disclosure implementation cost and that the policy maker chooses nondisclo-

sure. As implementation becomes more costly, consumers know that the policy maker is more reluctant to mandate disclosure, which means nondisclosure is less indicative of high product quality. This effect mitigates consumers' Type II error, which helps to improve social welfare.

The premise of this article is that consumers draw inferences from an observed policy. An experiment lends empirical support to this hypothesis. I find that emphasizing mandatory disclosure of GMOs in food products lowers consumers' perceived GMO safety compared to emphasizing nondisclosure. This result is robust across students versus non-students and across consumers with likely different degrees of attention to food safety.

The findings of this article suggest that the design of public policy should take into account the message it sends to the public regarding the underlying economic variable. There is a general belief that more information is socially beneficial (e.g., Milgrom and Weber 1982). Mandatory disclosure policies, in particular, are widely proposed to improve transparency in the marketplace (e.g., Coffee 1984; Easterbrook and Fischel 1984).⁴ This article argues that, although mandatory disclosure improves transparency, it might worsen other forms of information asymmetry as consumers second-guess the policy maker's motivation behind mandating disclosure. This possibility is illustrated in the context of product labeling, but the idea can be applied to other scenarios.

This article is also related to the literature on observational learning. The seminal works of Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992) study the process by which individuals learn about the value of an object (e.g., a product) through observing others' choice behaviors. Observational learning may lead to *ex post* inefficient decisions because there is information attrition in the learning process – others' choices may not perfectly reveal their private information about the object. The same reasoning extends beyond the setting of social learning. It applies broadly when decision-makers retrieve information from its coarse representation (see also Meyer 1991). It applies, for example, when discrete public policies (e.g., mandate disclosure or not) coarsely represent the finer-grained underlying economic variable (e.g., a risky product's quality). Policy design should thus consider the need to manage the observational learning process of the public who witness the policy but not its detailed rationale.

⁴This article focuses on mandatory disclosure. There is a large body of literature on voluntary disclosure of verifiable information (e.g., Grossman 1981; Milgrom 1981; Okuno-Fujiwara et al. 1990). The current article treats product suppliers' content disclosure as being exogenously mandated. The disclosure rule, however, is endogenously determined by the policy maker.

The rest of the article is organized as follows. Section 2 introduces the model setup. Section 3 derives the optimal disclosure policy in the benchmark case in which consumers observe the risky product’s quality. Section 4 presents the main analysis, whereby consumers do not observe the risky product’s quality. Section 5 explores the effect of disclosure implementation cost on social welfare. Section 6 reports an experiment, which shows that emphasizing GMO disclosure lowers consumers’ perceived GMO safety. Section 7 concludes with discussions and suggestions for future research. Proofs are presented in the Appendix.

2 Model Setup

Consider a one-shot game between a policy maker and a unit mass of consumers, each of whom chooses whether to adopt a product. Product content is either innocuous (denoted by $r = 0$) or risky (denoted by $r = 1$). By adopting the product ($a = 1$), a consumer derives a normalized utility of 1 if the product is innocuous, and s if the product is risky, where $s \in [0, 1]$ measures the inherent quality of the risky product. By not adopting the product ($a = 0$), the consumer enjoys her reservation utility v . The consumer privately observes v , but the distribution of v across all consumers, characterized by the continuous probability distribution function $f(v)$ over $[0, 1]$ with positive support everywhere, is common knowledge. The consumer seeks to maximize her expected utility.

The policy maker decides whether to mandate the disclosure of product content. If it mandates disclosure ($m = 1$), the consumer will be able to observe product content r . Meanwhile, the policy maker incurs an implementation cost $k \geq 0$, which includes, for instance, the cost of verifying the product label (Golan et al. 2001). If the policy maker does not mandate disclosure ($m = 0$), the consumer does not observe product content, a case denoted as $r = \emptyset$.⁵ The disclosure policy is publicized to the consumer.

⁵In practice, product labeling is carried out by the producer. The policy maker can enforce mandatory disclosure by verifying product labels and imposing a sufficient penalty of breach, which gives the producer the incentive to truthfully label its product. To focus on the mechanism of interest, I make the simplifying assumption that a product label is uninformative without mandatory disclosure. This could happen if the policy maker allows voluntary disclosure but does not verify the label. For example, at the time of this article, the FDA leaves the definition of “natural” in food products undefined (Killian 2012). A profit-maximizing producer will always claim that its product is natural, which makes such product labels uninformative. Nevertheless, the idea of the article applies as long as mandating disclosure provides the consumer with more accurate information about product content than not mandating disclosure.

The consumer and the policy maker share the common information that product is risky with probability $\lambda \in (0, 1)$. For example, the percentage of various foods using GMOs is public information (e.g., Ronald and McWilliams 2010). Moreover, the policy maker privately knows the risky product's quality s . The consumer does not have precise knowledge of s ; she holds the commonly known prior belief that s follows a continuous probability distribution function $p_0(s)$ over $[0, 1]$ with positive support everywhere.

The timing of the game is summarized as follows.

1. Nature draws the risky product's quality, and reveals it to the policy maker.
2. The policy maker publicly chooses whether to mandate disclosure of product content.
3. The consumer decides whether to adopt the product.

I derive the perfect sequential equilibrium of the game (Grossman and Perry 1986). This equilibrium concept has three elements. First, the consumer's adoption decision is the best response to every posterior belief of the risky product's quality, denoted as p , every disclosure policy m , and every product label r that follows the disclosure policy. Formally,

$$a(p, m, r) = \operatorname{argmax}_{a \in \{0, 1\}} a u(p, m, r) + (1 - a)v, \quad (1)$$

where $u(p, m, r)$ is the consumer's expected consumption utility of the product given her posterior belief p , the disclosure policy m , and product label r . Sections 3 and 4 will present the derivation of expected consumption utility.

Second, the policy maker chooses a disclosure policy to maximize social welfare, anticipating the consumer's best response induced by this policy. In this modeling context, the policy maker's objective is equivalent to maximizing consumer welfare net of any disclosure implementation cost, where *consumer welfare* equals expected consumer utility integrated over the consumer reservation utility distribution $f(v)$. I denote social welfare as $w(s, m, a)$. The policy maker's strategy is formalized as

$$m(s) = \operatorname{argmax}_{m \in \{0, 1\}} w(s, m, a). \quad (2)$$

Third, the consumer forms her posterior belief p in a *consistent* manner. It is given by Bayes' rule on the equilibrium path, and imposes the following restrictions on off-equilibrium beliefs. If the consumer observes an off-equilibrium policy m , she should try to rationalize it

by looking for a set $S \subseteq [0, 1]$ such that for every quality level $s \in S$ the policy maker is weakly better off choosing m (and inducing the consumer’s best response given beliefs concentrated on S) than in the candidate equilibrium, and that for every quality level $s \in [0, 1] - S$ the policy maker is weakly worse off choosing m than in equilibrium. If there exists such a nonempty set S , then the posterior belief for $s \in S$ is given by the conditional distribution: $p(s) = p_0(s) / \int_{s \in S} p_0(s) ds$.⁶

The following section proceeds with the benchmark case in which consumers observe the risky product’s quality s .

3 Consumers Observe the Risky Product’s Quality

In this benchmark case, consumers and the policy maker share the common understanding that the risky product offers quality s . The disclosure policy affects a consumer’s expected consumption utility only through the product label: $r \in \{1, 0\}$ if $m = 1$ and $r = \emptyset$ if $m = 0$. The consumer’s expected consumption utility is thus sufficiently captured by $u(s, r)$. If the product is labeled as innocuous under mandatory disclosure, the consumer knows that consumption utility is

$$u(s, r = 0) = 1. \tag{3}$$

A second possibility is that the product is labeled as risky under mandatory disclosure. To simplify exposition for subsequent analysis, I denote the corresponding consumption utility as $u^d(s)$, where d stands for “disclosure.”

$$u^d(s) \equiv u(s, r = 1) = s. \tag{4}$$

Finally, if the policy maker does not mandate disclosure, product label \emptyset is uninformative, and the consumer only knows that the product is risky with probability λ . Expected consumption

⁶The commonly used equilibrium refinement concepts, such as the Intuitive Criterion (Cho and Kreps 1987) and the D1 and D2 Criteria (Banks and Sobel 1987), cannot rule out any pooling equilibria in this context. Meanwhile, the use of the perfect sequential equilibrium concept is particular suitable in this setting – because the policy maker has only two actions to choose from, the perfect sequential equilibrium concept is exempt from the critique of Mailath, Okuno-Fujiwara and Postlewaite (1993).

utility, denoted as $u^{nd}(s)$ where nd means “nondisclosure,” is given by

$$u^{nd}(s) \equiv u(s, r = \emptyset) = \lambda s + 1 - \lambda. \quad (5)$$

Both $u^d(s)$ and $u^{nd}(s)$ increase with s ; the higher the quality of the risky product, the higher the consumer’s expected consumption utility (unless the consumer knows that the product is innocuous for certain). In addition, it is easy to see that

$$u^d(s) \leq u^{nd}(s) \leq 1, \quad \forall s \in [0, 1] \quad (6)$$

where the inequality holds strictly unless $s = 1$. That is, mandatory disclosure provides the consumer with greater transparency by revealing product content information and by “polarizing” the consumer’s valuations of the product based on its content.

The consumer’s product adoption strategy is given by $a(u) = \mathbf{1}(u \geq v)$ where $\mathbf{1}(\cdot)$ is the indicator function.⁷ Depending on the consumer’s reservation utility v , there are three possibilities: $0 \leq v \leq u^d(s)$ so that the consumer is always willing to adopt; $u^d(s) < v \leq u^{nd}(s)$ so that the consumer is willing to adopt if there is no disclosure but unwilling if the product is labeled as risky; $u^{nd}(s) < v \leq 1$ so that the consumer is willing to adopt only if the product is labeled as innocuous.

The consumer’s adoption strategy determines her expected utility under either disclosure policy. First consider the case of mandatory disclosure. When $0 \leq v \leq u^d(s)$, the consumer’s reservation utility is so low that she is always willing to adopt the product. It follows that the consumer’s expected utility is $\lambda s + 1 - \lambda$. When $u^d(s) < v \leq 1$, the consumer will only adopt if the product is labeled as innocuous. Her expected utility is thus $\lambda v + 1 - \lambda$. From the policy maker’s perspective, social welfare, integrated over all possible values of v and denoted by $w^d(s)$ for simplicity, equals

$$w^d(s) = \int_0^{u^d(s)} (\lambda s + 1 - \lambda) f(v) dv + \int_{u^d(s)}^1 (\lambda v + 1 - \lambda) f(v) dv - k. \quad (7)$$

In the case of nondisclosure, the cutoff reservation utility becomes $u^{nd}(s)$. When $0 \leq$

⁷As a tie-breaking rule, I assume that the consumer will choose to adopt the product when indifferent. This assumption simplifies analysis because it retains the focus on the consumer’s best response in pure strategies.

$v \leq u^{nd}(s)$, the consumer will again always adopt the product. When $u^{nd}(s) < v \leq 1$, the consumer will never adopt; without disclosure, the product's expected quality is insufficient to overcome the consumer's high reservation value. From the policy maker's perspective, social welfare as denoted by $w^{nd}(s)$ equals

$$w^{nd}(s) = \int_0^{u^{nd}(s)} (\lambda s + 1 - \lambda) f(v) dv + \int_{u^{nd}(s)}^1 v f(v) dv. \quad (8)$$

Let $\Delta(s) = w^d(s) + k - w^{nd}(s)$ denote the change in consumer welfare by mandating disclosure. It follows from Equations (7) and (8) that

$$\Delta(s) = \lambda \int_{u^d(s)}^{u^{nd}(s)} (v - s) f(v) dv + (1 - \lambda) \int_{u^{nd}(s)}^1 (1 - v) f(v) dv. \quad (9)$$

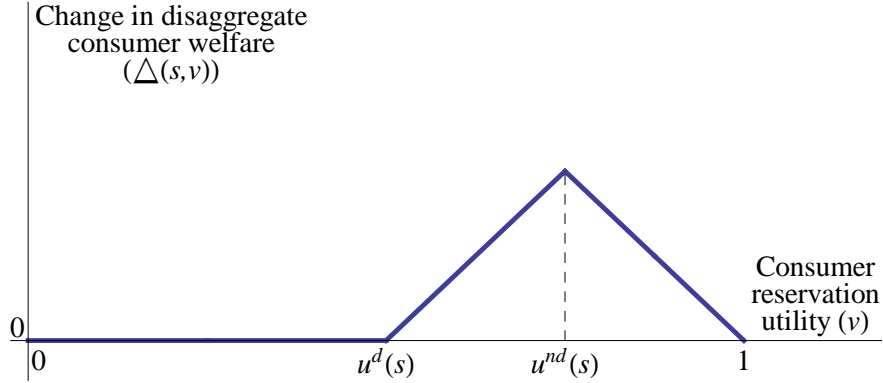
Figure 1 illustrates the effect of mandatory disclosure on consumer welfare. The parameters are set to be $\lambda = 0.5$ and $s = 0.5$, and $f(v)$ is assumed to be the uniform distribution over $[0, 1]$. The bold curve plots $\Delta(s, v)$, the disaggregate welfare change that mandatory disclosure brings to the consumer whose reservation utility is v . The area under the curve, $\Delta(s) = \int_0^1 \Delta(s, v) f(v) dv$, represents the change in aggregate consumer welfare from the policy maker's perspective.

If the consumer's reservation utility is below $u^d(s)$, she is always willing to adopt the product, hence the disclosure policy has no effect on consumer welfare. If the consumer's reservation utility is between $u^d(s)$ and $u^{nd}(s)$, mandatory disclosure helps her identify and avoid a risky product. If the consumer's reservation utility is between $u^{nd}(s)$ and 1, disclosure helps her recognize an innocuous product, which she would not have adopted were there no product label. In these two cases, mandatory disclosure improves transparency and allows the consumer to make "better informed" decisions. The following proposition holds in regard to the value of mandatory disclosure (see the Appendix for proof).

Proposition 1. *If consumers observe the risky product's quality (s), mandatory disclosure always improves consumer welfare unless the risky product is equivalent to the innocuous product ($s = 1$), and the magnitude of this improvement decreases with the risky product's quality. Formally:*

$$\Delta(s) \geq 0 \text{ and } \frac{\partial \Delta(s)}{\partial s} < 0, \forall s \in [0, 1], \quad (10)$$

Figure 1: Effect of Mandatory Disclosure on Consumer Welfare (Consumers Observe the Risky Product's Quality)



Notes. This figure assumes that $\lambda = 0.5$, $s = 0.5$, and $f(v) \sim U[0, 1]$. The bold curve plots $\Delta(s, v)$, the change in disaggregate consumer welfare brought forth by mandatory disclosure. The area under the curve, $\Delta(s) = \int_0^1 \Delta(s, v) f(v) dv$, represents the change in aggregate consumer welfare.

where the inequality holds strictly unless $s = 1$.

The proposition echoes the conventional belief that mandatory disclosure benefits consumers by improving information transparency. Rather intuitively, transparency is more important to consumers when the risky product's quality is worse. This result implies that, unless mandatory disclosure is too costly to implement, the optimal disclosure policy follows a threshold rule: the policy maker will mandate disclosure if and only if the risky product's quality is sufficiently low.⁸ The rest of the article will therefore focus attention on the following parameter space:

$$k \in [0, \bar{k}], \text{ where } \bar{k} = \Delta(0). \quad (11)$$

The corollary below formalizes the policy maker's equilibrium strategy in this benchmark case (see the Appendix for proof).

Corollary 1. *If consumers observe the risky product's quality (s), for any disclosure implementation cost $k \in [0, \bar{k}]$, there exists a quality threshold $s^*(k) = \Delta^{-1}(k) \in [0, 1]$ such that the policy maker mandates disclosure if and only if $s \leq s^*(k)$. The quality threshold $s^*(k)$ decreases with k .*

⁸As a tie-breaking rule, I assume that the policy maker will mandate disclosure when indifferent. This assumption simplifies analysis and retains the focus on the policy maker's pure strategies.

The above threshold rule accords well with intuition: mandating disclosure is worth its implementation cost only if the risky product’s quality poses a sufficient concern. The higher the implementation cost, the smaller the scope of risky products (in terms of their quality levels) that are subject to mandatory disclosure. The next section explores how these results change when the consumer does not observe the risky product’s quality.

4 Consumers Do Not Observe the Risky Product’s Quality

The previous section shows that mandatory disclosure always benefits consumers if consumers are aware of the risky product’s quality. This section presents the main argument of this article, that mandatory disclosure can generate perverse effects if consumers do not know the risky product’s quality but infer it from the disclosure policy.

When a consumer does not observe the risky product’s quality, the disclosure policy m affects expected consumption utility not only through the product label r , but also through the consumer’s posterior belief about the risky product’s quality $p(s|m)$. If the product is labeled as innocuous under mandatory disclosure, the consumer again knows that consumption utility equals 1. If the product is labeled as risky under mandatory disclosure, expected consumption utility is

$$\tilde{u}^d \equiv u[p(s|m = 1), r = 1] = \mathbb{E}(s|m = 1), \quad (12)$$

where $\mathbb{E}(s|m)$ denotes the consumer’s expectation of s conditional on disclosure policy m , and $\tilde{\cdot}$ differentiates values from their counterparts in the benchmark case of Section 3. Similarly, if disclosure is not mandated, the consumer’s expected consumption utility is

$$\tilde{u}^{nd} \equiv u[p(s|m = 0), r = \emptyset] = \lambda\mathbb{E}(s|m = 0) + 1 - \lambda. \quad (13)$$

The consumer’s expected consumption utilities \tilde{u}^d and \tilde{u}^{nd} depends on her beliefs about the risky product’s quality. However, the extent to which the consumer is making the right product choice depends on the actual quality s . Therefore, from the policy maker’s perspective, social welfare is a function of s . Specifically, by mandating disclosure, the policy maker

generates social welfare of

$$\tilde{w}^d(s) = \int_0^{\tilde{u}^d} (\lambda s + 1 - \lambda) f(v) dv + \int_{\tilde{u}^d}^1 (\lambda v + 1 - \lambda) f(v) dv - k. \quad (14)$$

By not mandating disclosure, the policy maker yields social welfare of

$$\tilde{w}^{nd}(s) = \int_0^{\tilde{u}^{nd}} (\lambda s + 1 - \lambda) f(v) dv + \int_{\tilde{u}^{nd}}^1 v f(v) dv. \quad (15)$$

Let $\tilde{\Delta}(s) = \tilde{w}^d(s) + k - \tilde{w}^{nd}(s)$ denote the improvement in consumer welfare caused by mandatory disclosure. Rearranging terms yields

$$\tilde{\Delta}(s) = \lambda \int_{\tilde{u}^d}^{\tilde{u}^{nd}} (v - s) f(v) dv + (1 - \lambda) \int_{\tilde{u}^{nd}}^1 (1 - v) f(v) dv. \quad (16)$$

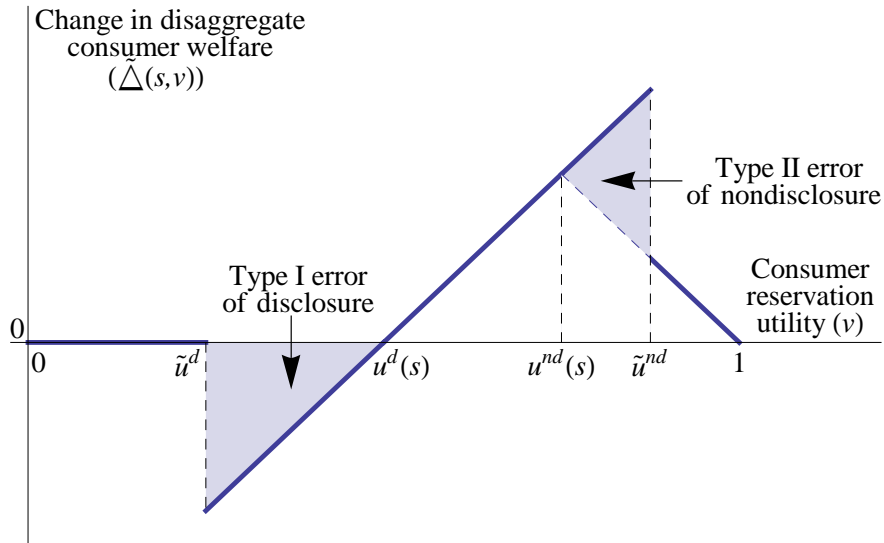
The policy maker should mandate disclosure if and only if $\tilde{\Delta}(s) \geq k$. How does the equilibrium disclosure policy compare with its counterpart in the benchmark case?

To gain more intuition, suppose the policy maker mistakenly assumes that consumers observe the risky product's quality. The policy maker will thus follow the optimal strategy in the benchmark case, which is to mandate disclosure if and only if $s \leq s^*(k)$. For illustrative purposes, suppose $0 < s^*(k) < 1$. It follows that $\mathbb{E}(s|m=1) < s^*(k) < \mathbb{E}(s|m=0)$, which in turn implies that $\tilde{u}^d < u^d(s^*(k))$ and that $u^{nd}(s^*(k)) < \tilde{u}^{nd}$.

In addition, let us assume that the risky product's quality s equals the quality threshold $s^*(k)$. Thinking that consumers observe s , the policy maker will be indifferent about whether to mandate disclosure – the transparency effect of mandatory disclosure on consumer welfare $\Delta(s^*(k))$ equals the implementation cost k . However, when consumers do not observe the risky product's quality in reality, the disclosure policy may generate two perverse effects. First, if $\tilde{u}^d < v \leq u^d(s^*(k))$, then disclosure will cause a *Type I error* in consumer choice. The consumer would have adopted the product had she observed the risky product's quality, but will not adopt because she underestimates quality by inferring it from mandatory disclosure. Second, if $u^{nd}(s^*(k)) < v \leq \tilde{u}^{nd}$, then nondisclosure will cause a *Type II error* in consumer choice. The consumer would have rejected the product had she observed the risky product's quality, but will adopt it because she overestimates quality by inferring it from nondisclosure.

Figure 2 illustrates these two errors using the same example from Figure 1, assuming that $s = s^*(k) = 0.5$ and that $p_0(s)$ is the uniform distribution over $[0, 1]$. The bold curve plots the disaggregate welfare change that mandatory disclosure brings to the consumer of reservation utility v . This curve falls below zero for v between \tilde{u}^d and $u^d(s^*(k))$, which means that the consumer is worse off with disclosure than without. The curve is above its benchmark level for v between $u^{nd}(s^*(k))$ and \tilde{u}^{nd} , which means that disclosure brings greater improvement in consumer welfare than in the benchmark scenario. The net effect of disclosure on aggregate consumer welfare, measured as the net area under the curve, is in general different than that in the benchmark case – it is lower than its benchmark counterpart for the example presented in this figure. Recall that $\Delta(s^*(k)) = k$. If $\tilde{\Delta}(s^*(k)) \neq k$, then the optimal disclosure policy in the benchmark is no longer optimal when consumers do not observe the risky product’s quality.

Figure 2: Effect of Mandatory Disclosure on Consumer Welfare (Consumers Do Not Observe the Risky Product’s Quality)



Notes. This figure assumes that $\lambda = 0.5$, $s = s^*(k) = 0.5$, $f(v) \sim U[0, 1]$, and $p_0(s) \sim U[0, 1]$.

If the Type I error of mandatory disclosure is sufficiently large, it might overwhelm the transparency effect. The net effect of mandatory disclosure on consumer welfare, $\tilde{\Delta}(s)$, could thus be negative, contrary to Proposition 1. In this case, the policy maker might not always want to mandate disclosure even if implementation is free ($k = 0$). If the Type II error of nondisclosure is sufficiently large, it might dominate the savings in disclosure implementation

cost. The net effect of mandatory disclosure on consumer welfare could be greater than \bar{k} , the highest possible consumer welfare improvement in the benchmark case. It follows that the policy maker might want to mandate disclosure even if the implementation cost is prohibitive ($k = \bar{k}$). Indeed, the following results hold true (see the Appendix for proof).

Proposition 2. *If consumers do not observe the risky product's quality, even if disclosure is costless to implement ($k = 0$), there exists no perfect sequential equilibrium in which the policy maker always mandates disclosure; even if disclosure is prohibitive to implement ($k = \bar{k}$), there exists no perfect sequential equilibrium in which the policy maker never mandates disclosure.*

Recall that in the benchmark case the policy maker will always mandate disclosure if $k = 0$ and will never mandate disclosure if $k = \bar{k}$ (except for being indifferent when $s = 1$). These results change when consumers do not observe the risky product's quality but infer it from the disclosure policy. The intuition behind Proposition 2 is as follows.

If the policy maker always mandates disclosure, a consumer learns nothing about the risky product's quality from such a policy beyond her prior belief. Depending on the consumer's off-equilibrium beliefs and the actual quality level, the policy maker may want to deviate to nondisclosure if doing so better communicates its private information about quality to the consumer. At the extreme, if by observing nondisclosure the consumer holds the off-equilibrium belief that $s = 1$, then the policy maker will indeed deviate to nondisclosure when $s = 1$. Doing so convinces the consumer of virtually innocuous product quality and eliminates the Type I error in consumer choice.

Following the same logic, if the consumer holds the off-equilibrium belief that deviation to nondisclosure indicates relatively high quality, then the policy maker will indeed deviate when quality is relatively high. The proof in the Appendix shows that there exists a quality threshold $\hat{s} \in (0, 1)$ such that, if the consumer believes that deviation to nondisclosure indicates $s \in [\hat{s}, 1]$, then the policy maker will fare better under nondisclosure if $s > \hat{s}$ and fare worse if $s < \hat{s}$. This result rules out the candidate pooling equilibrium of disclosure at $k = 0$.

By the same intuition, even if $k = \bar{k}$, there exists no perfect sequential equilibrium in which the policy maker always chooses nondisclosure. When the risky product's quality is sufficiently low, the policy maker would always want to deviate and mandate disclosure. In this case, the total information value of mandatory disclosure, which captures both the transparency effect and the inference effect, more than offsets the implementation cost.

The results so far emphasize the fact that the policy maker has different incentives to mandate disclosure given different quality levels of the risky product. A natural question to explore next is whether there exists a semi-separating equilibrium in which disclosure policies are governed by a threshold strategy and, if so, how does this threshold strategy compare with its counterpart in the benchmark case. The following proposition states the answer (see the Appendix for proof).⁹

Proposition 3. *If consumers do not observe the risky product’s quality (s), then for any disclosure implementation cost $k \in [0, \bar{k}]$ there exists a quality threshold $\tilde{s}^*(k) \in (0, 1)$ and a semi-separating equilibrium in which the policy maker mandates disclosure if and only if $s \leq \tilde{s}^*(k)$. The quality threshold $\tilde{s}^*(k)$ can increase with k .*

It is worth noting that the quality threshold $\tilde{s}^*(k)$ lies strictly between 0 and 1 even if k equals 0 or \bar{k} , a result that echoes Proposition 2. Moreover, the quality threshold $\tilde{s}^*(k)$ can increase with k , in contrast to the benchmark case where $s^*(k)$ always decreases with k (Corollary 1). To understand this result, recall that $\tilde{s}^*(k)$ is derived from $\tilde{\Delta}(\tilde{s}^*) = k$. The question then becomes: why can a higher \tilde{s}^* be associated with greater consumer welfare $\tilde{\Delta}(\tilde{s}^*)$ that justifies the higher implementation cost of disclosure? A closer look at the relationship between \tilde{s}^* and $\tilde{\Delta}(\tilde{s}^*)$ helps to reveal the intuition. It follows from Equation (16) that the total derivative of $\tilde{\Delta}(\tilde{s}^*)$ with respect to \tilde{s}^* can be decomposed into three components:

$$\frac{d\tilde{\Delta}(\tilde{s}^*)}{d\tilde{s}^*} = \underbrace{\int_{\tilde{u}^d}^{\tilde{u}^{nd}} (-\lambda)f(v)dv}_{\text{Decrease in transparency effect}} - \underbrace{\frac{\partial \tilde{u}^d}{\partial \tilde{s}^*} \lambda(\tilde{u}^d - \tilde{s}^*)f(\tilde{u}^d)}_{\text{Decrease in Type I error of disclosure}} + \underbrace{\frac{\partial \tilde{u}^{nd}}{\partial \tilde{s}^*} (\tilde{u}^{nd} - \lambda\tilde{s}^* - 1 + \lambda)f(\tilde{u}^{nd})}_{\text{Increase in Type II error of nondisclosure}} \quad (17)$$

where $\tilde{u}^d = \mathbb{E}(s|s \leq \tilde{s}^*)$ and $\tilde{u}^{nd} = \lambda\mathbb{E}(s|s > \tilde{s}^*) + 1 - \lambda$.

The first term on the right-hand side of Equation (17) reflects the change in the transparency effect of mandatory disclosure. This term is negative because, naturally, knowing

⁹When λ is sufficiently large, there also exists a semi-separating equilibrium in which the policy maker mandates disclosure if and only if the risky product’s quality is weakly higher than a threshold. However, following the reasoning of Schelling (1960), it is reasonable to assume that the equilibrium described in Proposition 3 serves as the “focal” equilibrium that the policy maker and the consumer can coordinate on because it accords better with popular beliefs regarding product labeling. For example, the common argument against labeling GMOs is that “genetic modification does not materially change the food” (Harmon and Pollack 2012). Moreover, the experiment in Section 6 shows that emphasizing GMO disclosure does worsen consumers’ perception of GMO safety.

whether a product is risky becomes less important when the risky product is of higher quality.

The second term captures the decrease in the Type I error of disclosure following an increase in \tilde{s}^* . This term is positive for two reasons. First, $\partial\tilde{u}^d/\partial\tilde{s}^* > 0$ by definition; knowing that the policy maker will mandate disclosure even when the risky product's quality is high, the consumer will not perceive mandatory disclosure as a strong indication of low quality. Second, $\tilde{u}^d < \tilde{s}^*$; when the policy maker is indifferent about whether to mandate disclosure ($s = \tilde{s}^*$) but is unable to communicate this true quality to the consumer, the consumer will naturally hold more pessimistic beliefs and will incur the Type I error in product adoption. A higher disclosure threshold mitigates this distortion by improving the consumer's quality beliefs upon observing mandatory disclosure.

The third term represents the change in the Type II error of nondisclosure following an increase in \tilde{s}^* . This term is also positive for two reasons. First, $\partial\tilde{u}^{nd}/\partial\tilde{s}^* > 0$; exemption from a strict mandatory disclosure policy means product quality is reassuringly high. Second, $\tilde{u}^{nd} > \lambda\tilde{s}^* + 1 - \lambda = u^{nd}(\tilde{s}^*)$; when actual quality equals the threshold, the consumer will infer higher quality from nondisclosure and make a Type II error of adoption. Raising the disclosure threshold thus exacerbates this distortion of nondisclosure, which makes disclosure a more desirable policy choice.

In summary, the welfare improvement brought forth by mandatory disclosure, $\tilde{\Delta}(\tilde{s}^*)$, decreases with \tilde{s}^* because of the transparency effect and increases with it because of the inference effect (the second and third terms of Equation (17)). When the inference effect dominates the transparency effect, the total effect $d\tilde{\Delta}(\tilde{s}^*)/d\tilde{s}^*$ is positive, which in turn means that the quality threshold $\tilde{s}^*(k)$ increases with disclosure implementation cost. In comparison, only the transparency effect is present in the benchmark case, so that the quality threshold can only decrease with disclosure implementation cost.

5 Effect of Disclosure Implementation Cost on Social Welfare

The above observations lead to another question: how does the implementation cost k affect social welfare? The answer is nuanced. On the one hand, a higher implementation cost

directly lowers social welfare when the policy maker chooses to mandate disclosure. On the other hand, the implementation cost changes the quality threshold for disclosure, which in turn affects the inference that consumers draw from the disclosure policy, and hence affects their product adoption decisions.

To disentangle the welfare consequences of the disclosure implementation cost, it is helpful to start with the benchmark case in which the inference effect is absent. Given the risky product's quality level s , social welfare (with the parameter k explicitly spelled out) equals

$$w(s; k) = \mathbf{1}[s \leq s^*(k)]w^d(s; k) + \mathbf{1}[s > s^*(k)]w^{nd}(s; k), \quad (18)$$

where $w^d(s; k)$ and $w^{nd}(s; k)$ are given by Equations (7) and (8), respectively. It is straightforward to see that $\partial w^d(s; k)/\partial k = -1$ and that $\partial w^{nd}(s; k)/\partial k = 0$. Meanwhile, the value of k determines the choice of the disclosure policy given s . Nevertheless, a higher disclosure implementation cost can never improve social welfare.

When the consumer does not observe the risky product's quality, for any given level of s , social welfare becomes

$$\tilde{w}(s; k) = \mathbf{1}[s \leq \tilde{s}^*(k)]\tilde{w}^d(s; k) + \mathbf{1}[s > \tilde{s}^*(k)]\tilde{w}^{nd}(s; k), \quad (19)$$

where $\tilde{w}^d(s; k)$ and $\tilde{w}^{nd}(s; k)$ are given by Equations (14) and (15), respectively. It follows that

$$\frac{d\tilde{w}(s; k)}{dk} = \mathbf{1}[s \leq \tilde{s}^*(k)] \left\{ \frac{\partial \tilde{w}^d}{\partial k} \lambda [u^d(s) - \tilde{u}^d] f(\tilde{u}^d) - 1 \right\} + \mathbf{1}[s > \tilde{s}^*(k)] \frac{\partial \tilde{w}^{nd}}{\partial k} [u^{nd}(s) - \tilde{u}^{nd}] f(\tilde{u}^{nd}), \quad (20)$$

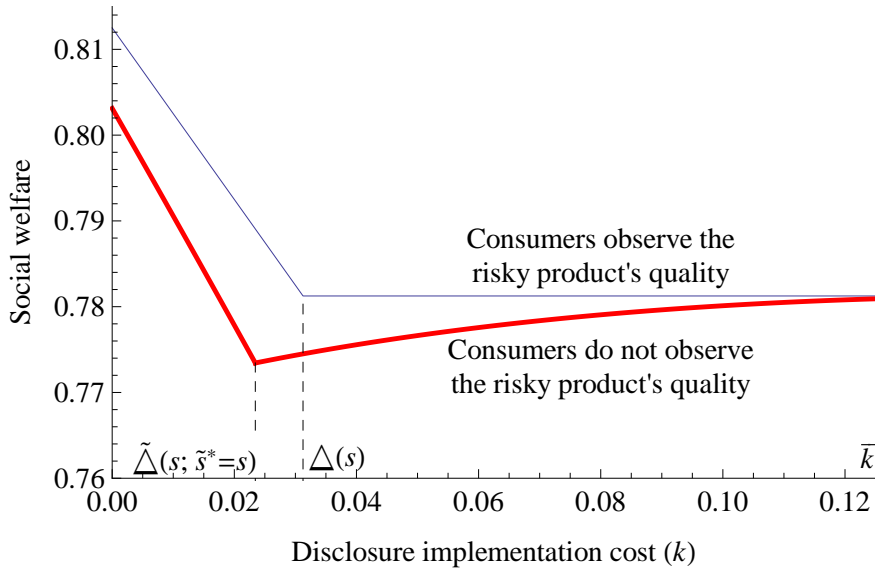
where $\tilde{u}^d = \mathbb{E}(s | s \leq \tilde{s}^*(k))$ and $\tilde{u}^{nd} = \lambda \mathbb{E}(s | s > \tilde{s}^*(k)) + 1 - \lambda$.

As in the benchmark case, disclosure implementation cost directly reduces social welfare in the case of mandatory disclosure, and affects the equilibrium disclosure policy by shifting the disclosure threshold $\tilde{s}^*(k)$. Different from the benchmark case, the shift in the disclosure threshold also changes consumer welfare by influencing consumer inference. It varies the consumer's expected utility from consumption (\tilde{u}^d or \tilde{u}^{nd} , depending on the disclosure policy), changes the adoption decision of the marginal consumer (whose reservation utility equals \tilde{u}^d or \tilde{u}^{nd}), and affects whether the marginal consumer is making a wise decision (how $u^d(s)$

compares with \tilde{u}^d , or how $u^{nd}(s)$ compares with \tilde{u}^{nd} .

To illustrate these effects, Figure 3 presents the example in which $f(v)$ and $p_0(s)$ are uniform distributions over $[0, 1]$, $\lambda = 0.5$, and s is fixed at 0.5. The thin curve plots social welfare when the consumer observes the risky product's quality. The policy maker switches from disclosure to nondisclosure when the implementation cost k reaches $\Delta(s)$. Social welfare first decreases with the implementation cost and then stays constant; it never increases.

Figure 3: Disclosure Implementation Cost and Social Welfare



Notes. This figure assumes that $s = 0.5$, $\lambda = 0.5$, $f(v) \sim U[0, 1]$, and $p_0(s) \sim U[0, 1]$.

When the consumer does not observe the risky product's quality, it can be shown that the quality threshold $\tilde{s}^*(k)$ decreases with k in this example. It follows that the policy maker will switch from disclosure to nondisclosure when the implementation cost increases to $\tilde{\Delta}(s; \tilde{s}^* = s)$, or equivalently, when $\tilde{s}^*(k)$ declines to s .¹⁰ It also follows that $\partial \tilde{u}^d / \partial k$ and $\partial \tilde{u}^{nd} / \partial k$ are both negative. Intuitively, mandating disclosure in spite of its high implementation cost suggests that the risky product's quality is really worrisome, whereas not mandating disclosure does not necessarily imply that the risky product is safe; it could simply mean that implementation is prohibitive.

¹⁰The notation $\tilde{\Delta}(s; \tilde{s}^* = s)$ emphasizes that the disclosure threshold \tilde{s}^* equals the actual quality level s . This is different from the notation $\tilde{\Delta}(s)$ used earlier, whereby the disclosure threshold and the actual quality level are allowed to diverge.

When the policy maker mandates disclosure, after a small increase in k and the ensuing small decrease in \tilde{u}^d , the consumer whose reservation utility equals \tilde{u}^d will change her mind and choose not to adopt the risky product. However, this decision is wrong whenever $s \geq \tilde{u}^d$, which always holds in this example since $s = 0.5 = \mathbb{E}(s)$ given the uniform prior $p_0(s)$. Therefore, an increase in k not only makes disclosure more costly to implement, but also exacerbates the consumer’s Type I error in this example. Because of the latter effect, the social welfare curve under disclosure declines faster in k than its counterpart in the benchmark case.

More interestingly, social welfare increases with disclosure implementation cost when the policy maker chooses nondisclosure. This result again traces back to the inference effect of the disclosure policy. After observing an increase in the implementation cost, the consumer knows that the disclosure threshold will decrease (for this example), and will infer a lower utility of consumption \tilde{u}^{nd} . The marginal consumer who holds a reservation utility of \tilde{u}^{nd} will thus choose not to adopt. This is the right decision because $u^{nd}(s) < \tilde{u}^{nd}$; that is, given the true quality level $s = \mathbb{E}(s)$, the consumer is prone to a Type II error following the nondisclosure policy. A higher implementation cost serves to mitigate this Type II error.

Depending on parameter values and the functional forms of $f(v)$ and $p_0(s)$, social welfare can also increase with the implementation cost under disclosure and decrease with it under nondisclosure. Nevertheless, the example presented above provides a “proof of existence” for the counterintuitive result that higher disclosure implementation cost can increase social welfare. The following proposition summarizes this result (the proof holds by construction).

Proposition 4. *If consumers do not observe the risky product’s quality (s), social welfare can increase with the disclosure implementation cost (k).*

6 A GMO Safety Perception Experiment

The premise underlying the analysis so far is that consumers draw inferences based on an observed public policy. This section presents an empirical evaluation of this mechanism in the context of mandatory GMO labeling. The focal empirical question is whether different GMO labeling policies affect consumers’ perception of GMO safety.

To answer this question, I conduct an experiment that exogenously varies GMO labeling policy statements and elicits subject’s GMO safety perception. Each subject is given a

questionnaire, which defines GMOs, describes the current policy in the U.S. regarding GMO labeling, and asks the subject to indicate how safe he or she thinks GMOs are on a 1-to-5 scale, with 1 being “totally unsafe” and 5 being “totally safe.”

The experiment follows a between-subject design that features two conditions: disclosure and nondisclosure. Because subjects may have preexisting knowledge about GMO labeling policies, the questionnaires are worded to evoke the emphasis on disclosure or nondisclosure without misreporting the current GMO disclosure policy in the U.S. In the the disclosure condition, subjects are informed that:

“At the time of this survey, there have been policy proposals in the U.S. to mandate the use of GMO ingredients to be disclosed on food labels. California voted in November 2012 on mandatory labeling of GMOs in food products.”

In the nondisclosure condition, the above statement is replaced by:

“At the time of this survey, the U.S. Food and Drug Administration (FDA) does NOT mandate the use of GMO ingredients to be disclosed on food labels.”

Human subjects are recruited from four locations to allow potentially different consumer groups to participate in the experiment. The first location is the student dining hall of a major university in the U.S. The second location is a food truck at the same university positioned around serving organic food. The third location is Shaw’s, a grocery chain that includes genetically modified products in its offerings.¹¹ The fourth location is Trader Joe’s, a grocery chain that claims to have “NO genetically modified ingredients” in their products.¹² In relation to the model, consumers that are attracted to these different locations might hold different reservation utilities for food consumption.

A total of 200 subjects participated in the experiment. At each location, 25 subjects were randomly assigned into the disclosure condition and another 25 into the nondisclosure condition. Participation was voluntary and unpaid.

Table 1 reports the experiment results. Across the entire sample, the mean perceived GMO safety is 2.650 in the disclosure condition and 3.620 in the nondisclosure condition.

¹¹Source: http://www.just-food.com/news/national-anti-ge-campaign-starts-with-protest-at-shawsstar-markets_id72477.aspx.

¹²Source: <http://www.traderjoes.com/products.asp>.

The difference is statistically significant ($p = 0.000$; two-tailed test). The same comparison holds for each of the four locations, although the difference is statistically insignificant for the dining-hall sub-sample.

Table 1: Experiment Results

	Entire Sample	Dining Hall	Organic Truck	Shaw's	Trader Joe's
Mean perceived GMO safety under disclosure	2.650 (1.029)	3.440 (1.003)	2.280 (0.737)	2.280 (0.792)	2.600 (1.118)
Mean perceived GMO safety under nondisclosure	3.620 (0.919)	3.560 (0.917)	3.680 (0.945)	3.720 (0.891)	3.520 (0.963)
t -stat of difference in mean perceived GMO safety	-7.176	-0.377	-5.715	-7.493	-3.874
p -value (two-tail)	0.000	0.709	0.000	0.000	0.001
# observations	200	50	50	50	50

Notes: The dependent variable is perceived GMO safety on a 1-to-5 scale, with 1 being “totally unsafe” and 5 being “totally safe.” At each survey location, there are 25 subjects in the disclosure condition and another 25 subjects in the nondisclosure condition, respectively. The numbers in parentheses are standard deviations.

In summary, the experiment shows that emphasizing mandatory disclosure of GMOs in food products lowers consumers’ perceived GMO safety. This finding is robust; it applies to both students and non-students, and to consumers with plausibly different degrees of attention to food quality. These results provide empirical evidence for the main premise of this article – policy choices affect consumer inferences.

7 Concluding Remarks

Public policies often result from extensive research. Consumers typically do not access the details of the research, but can draw inference about it from the policy. If this inference effect is ignored, even benevolent policies can lead to unintended outcomes. This article studies the inference effect and its policy implications in the context of mandatory product content disclosure. Some of the conventional wisdom regarding mandatory disclosure needs to be revised in the presence of the inference effect. Table 2 summarizes the main findings.

This article focuses on the case of product labeling for two reasons. Theoretically, prod-

Table 2: Summary of Main Findings

	Consumers observe the risky product's quality	Consumers do not observe the risky product's quality
Effect of disclosure on consumer welfare	Positive	Can be negative
Always mandate disclosure if disclosure implementation is costless	Yes	No
Effect of disclosure implementation cost on the scope of mandatory disclosure	Negative	Can be positive
Effect of disclosure implementation cost on social welfare	Non-positive	Can be positive

uct information disclosure is often seen as improving consumer welfare. Choosing the well-intentioned disclosure policy as a starting point helps highlight the adverse effect of ignoring consumer inferences. Practically, there has been heated debate over mandatory food content disclosure, particularly over the labeling of genetically modified foods. California recently voted on Proposition 37 that mandates the use of GMOs in food products to be disclosed on product labels. This article hopefully contributes a timely and fresh perspective to this ongoing debate.

Nevertheless, the idea that consumers draw inference from public policies can be applied to other settings. For example, from the requirement for physicians to disclose their affiliation with pharmaceutical companies, patients may infer that corrupted recommendations must be prevalent enough to be worth addressing. From the tightening of security scrutiny at airports, travelers may infer that security threats have caused sufficient concern. From the increase of taxation on a natural resource, manufacturers may infer a future decline in supply. Further research can study the implications of public inferences in these contexts.

This article assumes that product design is exogenously determined. In the long term, mandatory product information disclosure can change the supply of certain types of products (Ippolito and Mathios 1990; Golan et al. 2001). Future research can model the policy effect on product reformulation and its interaction with consumer inference. This article also assumes that the policy maker aims to maximize social welfare. It will be interesting to incorporate

other considerations, such as influence activities from product suppliers, into policy making. The effect on consumer inference is subtle. For example, if the policy maker is known to be biased against GMO labeling, there will be less transparency about product content, but the public will also be more cautious about GMO quality. The overall welfare consequences will be worth studying.

References

- Banerjee, Abhijit V. (1992), "A Simple Model of Herd Behavior," *Quarterly Journal of Economics*, 107(3), 797-817.
- Banks, Jeffrey S., and Joel Sobel (1987), "Equilibrium Selection in Signaling Games," *Econometrica*, 55(3), 647-661.
- Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch (1992), "A Theory of Fads, Fashion, Custom, and Cultural Changes as Information Cascades," *Journal of Political Economy*, 100(5), 992-1026.
- Cho, In-Koo, and David M. Kreps (1987), "Signaling Games and Stable Equilibria," *Quarterly Journal of Economics*, 102(2), 179-221.
- Coffee, John C. (1984), "Market Failure and the Economic Case for a Mandatory Disclosure System," *Virginia Law Review*, 70(4), 717-753.
- Easterbrook, Frank H., and Daniel R. Fischel (1984), "Mandatory Disclosure and the Protection of Investors," *Virginia Law Review*, 70(4), 669-715.
- Golan, Elise, Fred Kuchler, Lorraine Mitchell, Cathy Greene, and Amber Jessup (2001), "Economics of Food Labeling," *Journal of Consumer Policy*, 24(2), 117-184.
- Grossman, Sanford J. (1981), "The Informational Role of Warranties and Private Disclosure about Product Quality," *Journal of Law and Economics*, 24(3), 461-483.
- Grossman, Sanford J., and Motty Perry (1986), "Perfect Sequential Equilibrium," *Journal of Economic Theory*, 39(1), 97-119.
- Harmon, Amy, and Andrew Pollack (2012), "Battle Brewing Over Labeling of Genetically Modified Food," *New York Times*, May 24.
- Ippolito, Pauline M., and Alan D. Mathios (1990), "Information, Advertising and Health Choices: A Study of the Cereal Market," *RAND Journal of Economics*, 21(3), 459-480.
- Killian, Anastasia (2012), "California's Proposition 37: An Impending 'Natural' Disaster?" *Forbes*, October 15.
- Mailath, George J., Masahiro Okuno-Fujiwara, and Andrew Postlewaite (1993), "Belief-Based

- Refinements in Signalling Games,” *Journal of Economic Theory*, 60(2), 241-276.
- Milgrom, Paul R. (1981), “Good News and Bad News: Representation Theorems and Applications,” *Bell Journal of Economics*, 12(2), 380-391.
- Milgrom, Paul R., and Robert J. Weber (1982), “A Theory of Auctions and Competitive Bidding,” *Econometrica*, 50(5), 1089-1122.
- Meyer, Margaret A. (1991), “Learning from Coarse Information: Biased Contests and Career Profiles,” *Review of Economic Studies*, 58(1), 15-41.
- National Research Council (2004), *Safety of Genetically Engineered Foods: Approaches to Assessing Unintended Health Effects*, National Academies Press, Washington.
- Okuno-Fujiwara, Masahiro, Andrew Postlewaite, and Kotaro Suzumura (1990), “Strategic Information Revelation,” *Review of Economic Studies*, 57(1), 25-47.
- Organization for Economic Co-operation and Development (2000), *Report of the Task Force for the Safety of Novel Foods and Feeds*, May 17.
- Ramessar, Koreen, Teresa Capell, Richard M. Twyman, and Paul Christou (2010), “Going to Ridiculous Lengths—European Coexistence Regulations for GM Crops,” *Nature Biotechnology*, 28(2), 133-136.
- Ronald, Pamela C., and James E. McWilliams (2010), Genetically Engineered Distortions, *New York Times*, May 14.
- Schelling, Thomas C. (1960), *The Strategy of Conflict*, Harvard University Press.

Appendix

A.1 Proof of Proposition 1

Let $t_1(v, s) \equiv \lambda(v - s)$ and $t_2(v) \equiv (1 - \lambda)(1 - v)$. Equation (9) is rewritten as

$$\Delta(s) = \int_{u^d(s)}^{u^{nd}(s)} t_1(v, s) f(v) dv + \int_{u^{nd}(s)}^1 t_2(v) f(v) dv.$$

Note that

$$\frac{\partial t_1(v, s)}{\partial v} > 0, \quad \frac{\partial t_2(v)}{\partial v} < 0, \quad t_1(v, s)|_{v=u^d(s)} = 0, \quad t_2(v)|_{v=1} = 0.$$

Recall that $u^d(s) \leq u^{nd}(s) \leq 1$ for all $s \in [0, 1]$, where the inequality holds strictly unless $s = 1$. Recall also that $f(v) > 0$ for all $v \in [0, 1]$. Therefore, $\Delta(s) \geq 0$, where the inequality holds strictly unless $s = 1$.

It remains to show that $\frac{\partial \Delta(s)}{\partial s} < 0$.

$$\begin{aligned} \frac{\partial \Delta(s)}{\partial s} &= \int_{u^d(s)}^{u^{nd}(s)} \frac{\partial t_1(v, s)}{\partial s} f(v) dv + \frac{\partial u^{nd}(s)}{\partial s} t_1(v, s) f(v)|_{v=u^{nd}(s)} - \frac{\partial u^d(s)}{\partial s} t_1(v, s) f(v)|_{v=u^d(s)} \\ &\quad - \frac{\partial u^{nd}(s)}{\partial s} t_2(v) f(v)|_{v=u^{nd}(s)} \\ &= \int_{u^d(s)}^{u^{nd}(s)} \frac{\partial t_1(v, s)}{\partial s} f(v) dv + \frac{\partial u^{nd}(s)}{\partial s} [t_1(v, s) - t_2(v)] f(v)|_{v=u^{nd}(s)}. \end{aligned}$$

Note that

$$\frac{\partial t_1(v, s)}{\partial s} < 0, \quad [t_1(v, s) - t_2(v)]|_{v=u^{nd}(s)} = 0.$$

It follows that

$$\frac{\partial \Delta(s)}{\partial s} = \int_{u^d(s)}^{u^{nd}(s)} \frac{\partial t_1(v, s)}{\partial s} f(v) dv < 0, \forall s \in [0, 1).$$

Finally, note that $\Delta(s) = 0$ when $s = 1$ and that $\Delta(s) > 0$ for all $s < 1$. Therefore,

$$\frac{\partial \Delta(s)}{\partial s} < 0, \forall s \in [0, 1]. \quad \blacksquare$$

A.2 Proof of Corollary 1

Equations (4) and (5) imply that

$$u^d(0) = 0, \quad u^{nd}(0) = 1 - \lambda, \quad u^d(1) = 1, \quad u^{nd}(1) = 1.$$

It follows that

$$\Delta(0) = \lambda \int_0^{1-\lambda} v f(v) dv + (1 - \lambda) \int_{1-\lambda}^1 (1 - v) f(v) dv, \quad (\text{A1})$$

and that

$$\Delta(1) = 0. \quad (\text{A2})$$

Because $\Delta(s)$ is continuous and decreasing in s for all $s \in [0, 1)$, for any $k \in [0, \Delta(0)]$ there exists $s^*(k) = \Delta^{-1}(k) \in [0, 1]$ such that $\Delta(s) \geq k$ if and only if $s \leq s^*(k)$.

The quality threshold $s^*(k)$ is derived from

$$\zeta = \Delta(s^*) - k = 0.$$

By the implicit function theorem,

$$\frac{\partial s^*(k)}{\partial k} = -\frac{\partial \zeta / \partial k}{\partial \zeta / \partial s^*} = \frac{1}{\partial \Delta(s^*) / \partial s^*}.$$

Because $\partial \Delta(s) / \partial s < 0$ for any s , we have

$$\frac{\partial s^*(k)}{\partial k} < 0. \quad \blacksquare$$

A.3 Proof of Proposition 2

The first part of the proposition states that even if $k = 0$ there exists no perfect sequential equilibrium in which $m(s) = 1, \forall s \in [0, 1]$. Suppose such an equilibrium does exist. In this equilibrium, $\tilde{u}^d = \mathbb{E}(s) \in (0, 1)$, which is the unconditional mean of s based on its prior distribution $p_0(s)$. The value of \tilde{u}^{nd} depends on the consumer's off-equilibrium belief after

observing $m = 0$. Consider the following off-equilibrium belief: $s \in S = [\hat{s}, 1]$. It follows that $\tilde{u}^{nd} = \lambda \mathbb{E}(s|s \geq \hat{s}) + 1 - \lambda$.

Equation (16) implies that

$$\text{sign} \left[\frac{\partial \tilde{\Delta}(s)}{\partial s} \right] = \text{sign} (\tilde{u}^d - \tilde{u}^{nd}). \quad (\text{A3})$$

Because $\mathbb{E}(s|s \geq \hat{s}) \geq \mathbb{E}(s)$ by definition, we have $\tilde{u}^{nd} > \tilde{u}^d$. Therefore, $\partial \tilde{\Delta}(s)/\partial s < 0$, which establishes the following monotonicity property: if the policy maker is indifferent between $m = 0$ and $m = 1$ when $s = \hat{s}$, then it will prefer $m = 0$ for all $s > \hat{s}$, and will prefer $m = 1$ for all $s < \hat{s}$. The quality threshold \hat{s} , if it exists, is given by the following equation:

$$\phi(\hat{s}) = \lambda \int_{\mathbb{E}(s)}^{\lambda \mathbb{E}(s|s \geq \hat{s}) + 1 - \lambda} (v - \hat{s}) f(v) dv + (1 - \lambda) \int_{\lambda \mathbb{E}(s|s \geq \hat{s}) + 1 - \lambda}^1 (1 - v) f(v) dv = 0.$$

Note that

$$\phi(0) = \lambda \int_{\mathbb{E}(s)}^{\lambda \mathbb{E}(s) + 1 - \lambda} v f(v) dv + (1 - \lambda) \int_{\lambda \mathbb{E}(s) + 1 - \lambda}^1 (1 - v) f(v) dv > 0,$$

and that

$$\phi(1) = \lambda \int_{\mathbb{E}(s)}^1 (v - 1) f(v) dv + (1 - \lambda) \int_1^1 (1 - v) f(v) dv < 0.$$

Because $\phi(\hat{s})$ is continuous in \hat{s} , there exists $\hat{s} \in (0, 1)$ such that $\phi(\hat{s}) = 0$. Therefore, given the off-equilibrium belief $s \in S = [\hat{s}, 1]$, the policy maker will indeed want to deviate to $m = 0$ when $s > \hat{s}$ and will not want to deviate when $s < \hat{s}$. The existence of such a non-empty set S rules out the candidate pooling equilibrium in which $m(s) = 1, \forall s \in [0, 1]$.

The second part of the proposition states that even if $k = \bar{k}$ there exists no perfect sequential equilibrium in which $m(s) = 0, \forall s \in [0, 1]$. Suppose such an equilibrium does exist. In this equilibrium, $\tilde{u}^{nd} = \lambda \mathbb{E}(s) + 1 - \lambda$. The value of \tilde{u}^d depends on the consumer's off-equilibrium belief after observing $m = 1$. Consider the following off-equilibrium belief: $s \in S = [0, \check{s}]$. It follows that $\tilde{u}^d = \mathbb{E}(s|s \leq \check{s})$.

Because $\mathbb{E}(s|s \leq \check{s}) \leq \mathbb{E}(s)$ by definition, we have $\tilde{u}^{nd} > \tilde{u}^d$. By Equation (A3), $\partial \tilde{\Delta}(s)/\partial s < 0$, which establishes the following monotonicity property: if the policy maker is indifferent between $m = 0$ and $m = 1$ when $s = \check{s}$, then it will prefer $m = 0$ for all $s > \check{s}$, and will prefer

$m = 1$ for all $s < \check{s}$. The quality threshold \check{s} , if it exists, is given by the following equation:

$$\varphi(\check{s}) = \lambda \int_{\mathbb{E}(s|s \leq \check{s})}^{\lambda \mathbb{E}(s)+1-\lambda} (v - \check{s})f(v)dv + (1 - \lambda) \int_{\lambda \mathbb{E}(s)+1-\lambda}^1 (1 - v)f(v)dv = \bar{k}.$$

Note that

$$\varphi(0) = \lambda \int_0^{\lambda \mathbb{E}(s)+1-\lambda} vf(v)dv + (1 - \lambda) \int_{\lambda \mathbb{E}(s)+1-\lambda}^1 (1 - v)f(v)dv. \quad (\text{A4})$$

Recall that $\bar{k} = \Delta(0)$ is given by Equation (A1). Rearranging terms yields

$$\varphi(0) - \bar{k} = \int_{1-\lambda}^{\lambda \mathbb{E}(s)+1-\lambda} [\lambda v - (1 - \lambda)(1 - v)]f(v)dv.$$

Because $\lambda v - (1 - \lambda)(1 - v)$ increases with v and equals 0 at $v = 1 - \lambda$, we have

$$\varphi(0) - \bar{k} > 0. \quad (\text{A5})$$

Note also that

$$\begin{aligned} \varphi(1) &= \lambda \int_{\mathbb{E}(s)}^{\lambda \mathbb{E}(s)+1-\lambda} (v - 1)f(v)dv + (1 - \lambda) \int_{\lambda \mathbb{E}(s)+1-\lambda}^1 (1 - v)f(v)dv \\ &< (1 - \lambda) \int_{\lambda \mathbb{E}(s)+1-\lambda}^1 (1 - v)f(v)dv \\ &< (1 - \lambda) \int_{1-\lambda}^1 (1 - v)f(v)dv \\ &< \bar{k}. \end{aligned}$$

Because $\varphi(\hat{s})$ is continuous in \check{s} , there exists $\check{s} \in (0, 1)$ such that $\varphi(\check{s}) = \bar{k}$. Therefore, given the off-equilibrium belief $s \in S = [0, \check{s}]$, the policy maker will indeed want to deviate to $m = 1$ when $s < \check{s}$ and will not want to deviate when $s > \check{s}$. The existence of such a non-empty set S rules out the candidate pooling equilibrium in which $m(s) = 0, \forall s \in [0, 1]$. ■

A.4 Proof of Proposition 3

Suppose such an equilibrium does exist. Given the consumer's equilibrium beliefs, her expected utility from adopting the product is $\tilde{u}^d = \mathbb{E}(s|s \leq \tilde{s}^*(k))$ after observing disclosure

($m = 1$), and is $\tilde{u}^{nd} = \lambda \mathbb{E}(s|s > \tilde{s}^*(k)) + 1 - \lambda$ after observing nondisclosure ($m = 0$). The consumer's best response again is to adopt the product if and only if the expected utility is greater than or equal to her reservation utility v .

It follows that the policy maker's best response is $m = 1$ if and only if $\tilde{\Delta}(s) \geq k$, where $\tilde{\Delta}(s)$ is given by Equation (16). Because $\mathbb{E}(s|s \leq \tilde{s}^*(k)) < \mathbb{E}(s|s > \tilde{s}^*(k))$ by definition, we have $\tilde{u}^d < \tilde{u}^{nd}$. By Equation (A3), $\partial \tilde{\Delta}(s)/\partial s < 0$, which establishes the following monotonicity property: if the policy maker is indifferent between $m = 0$ and $m = 1$ when $s = \tilde{s}^*(k)$, then it will prefer $m = 0$ for all $s > \tilde{s}^*(k)$, and will prefer $m = 1$ for all $s < \tilde{s}^*(k)$. The quality threshold $\tilde{s}^*(k)$, if it exists, is given by the following equation:

$$\begin{aligned} \chi(\tilde{s}^*, k) &= \tilde{\Delta}(\tilde{s}^*) - k \\ &= \lambda \int_{\mathbb{E}(s|s \leq \tilde{s}^*)}^{\lambda \mathbb{E}(s|s > \tilde{s}^*) + 1 - \lambda} (v - \tilde{s}^*) f(v) dv + (1 - \lambda) \int_{\lambda \mathbb{E}(s|s > \tilde{s}^*) + 1 - \lambda}^1 (1 - v) f(v) dv - k \\ &= 0. \end{aligned}$$

First note that

$$\begin{aligned} \chi(0, k) &= \lambda \int_0^{\lambda \mathbb{E}(s) + 1 - \lambda} v f(v) dv + (1 - \lambda) \int_{\lambda \mathbb{E}(s) + 1 - \lambda}^1 (1 - v) f(v) dv - k \\ &= \varphi(0) - k, \end{aligned}$$

where $\varphi(0)$ is given by Equation (A4). Inequality (A5) then implies that

$$\chi(0, k) > 0, \quad \forall k \in [0, \bar{k}].$$

Meanwhile,

$$\chi(1, k) = \lambda \int_{\mathbb{E}(s)}^1 (v - 1) f(v) dv + (1 - \lambda) \int_1^1 (1 - v) f(v) dv - k < 0.$$

Because $\chi(\tilde{s}^*, k)$ is continuous in \tilde{s}^* , for any $k \in [0, \bar{k}]$ there exists $\tilde{s}^*(k) \in (0, 1)$ such that $\chi(\tilde{s}^*(k), k) = 0$. Hence the candidate equilibrium does exist.

It remains to prove that $\tilde{s}^*(k)$ can increase with k . By the implicit function theorem,

$$\frac{\partial \tilde{s}^*(k)}{\partial k} = -\frac{\partial \chi(\tilde{s}^*, k)/\partial k}{\partial \chi(\tilde{s}^*, k)/\partial \tilde{s}^*} = \frac{1}{d\tilde{\Delta}(\tilde{s}^*)/d\tilde{s}^*}.$$

Therefore, the goal is to prove that $d\tilde{\Delta}(\tilde{s}^*)/d\tilde{s}^*$ can be positive. To construct an example in which this is true, consider $p_0(s) \sim U[0, 1]$. It follows that

$$\tilde{u}^d = \frac{\tilde{s}^*}{2}, \quad \tilde{u}^{nd} = \lambda \frac{1 + \tilde{s}^*}{2} + 1 - \lambda.$$

Equation (17) is rewritten as

$$\frac{d\tilde{\Delta}(\tilde{s}^*)}{d\tilde{s}^*} = \lambda \left[-\int_{\frac{\tilde{s}^*}{2}}^{\lambda \frac{1+\tilde{s}^*}{2} + 1 - \lambda} f(v) dv + \frac{\tilde{s}^*}{4} f\left(\frac{\tilde{s}^*}{2}\right) + \frac{\lambda(1 - \tilde{s}^*)}{4} f\left(\lambda \frac{1 + \tilde{s}^*}{2} + 1 - \lambda\right) \right]$$

Assume further that $f(v) = 11(v - 1/2)^2 + 1/12$. It can be verified that $d\tilde{\Delta}(\tilde{s}^*)/d\tilde{s}^* > 0$ for a non-empty set of parameters – for example, when $k = 0.0165$ (such that $\tilde{s}^* = 0.5$) and $\lambda = 0.8$. ■