Optimal consumer regulation when consumers make inferences from regulatory policy

BY ABRAHAM L. WICKELGREN

University of Texas at Austin

Abstract

Consumers often value products that are potentially harmful to themselves. In many cases, consumers do not know exactly how harmful the product is, but can engage in some research to learn more about it. Alternatively, the government can choose to ban products that are very harmful. Whether or not the government does so will affect both the consumers purchase and research decision. This paper analyses when it is optimal to ban such products when consumers can become informed at a cost. It also compares the regulatory policy of banning dangerous products to taxing them under the assumption that consumers can make better inferences from bans than taxes.

1 Introduction

When deciding to purchase a potentially dangerous product, a fully-informed consumer can simply trade-off the expected risk with their personal benefit from the product. In this simple setting, there is no role for government regulation of dangerous products. Prohibiting dangerous products simply prevents those consumers whose personal value from the product exceeds the expected harm from buying a product that increases their utility, while those who value the product less than the harm can simply refrain from purchasing it.

But, for many products, consumers do not know the everything that is known about the possible impacts of using the product, at least not without expending some effort researching those risks. For example, when a consumer performs a search or browses a website, that consumer is unlikely to know exactly what data the website is collecting, who it is being shared with, or exactly how that sharing of data will impact them in the future. But, by doing some costly research on website's privacy policies and the uses of consumer data in general, this consumer can learn about how any data collection might affect them.

In other contexts, a consumer considering purchasing a mosquito treatment for their yard is unlikely to know what the health risks of that treatment is without spending a significant amount of time researching it. Whether doing that research is worth it depends both on the value of the mosquito treatment and on the prior distribution of risk. Similarly, some consumers considering a complex financial product like an adjustable rate mortgate, will not have a great idea of how risky the product is without learning a lot about interest rate volatility, the likelihood of being able to refinance the mortgage, among other factors.

In these situations, the prior distriubtion of risk depends on the consumer's expectation of government regulation. If regulators have a relatively lasseiz-faire attitude, consumers would rightly believe that available websites or social media products or mosquito treatments could pose a non-trivial risk to their privacy or human health or that low initial payments on adjustable rate mortgates might often appear too good to be true. On the other hand, if regulators are much more stringent, consumers would correctly believe that the maximum risk associated with these products might be fairly small (otherwise, the products would be unlikely to allowed on the market).

How these beliefs affect a consumer's decision to learn the actual harm of a product will vary with the consumer's value of the product. For consumers with very high values, there may be no reason to learn the true harm under any regulatory policy since the probability they would not want to purchase the product based on that information may be very low. Similarly, consumers with very low value would also have very low value of information since they are almost never going to purchase the product.

For consumers with low, but somewhat higher values, they might only spend the effort to learn the true harm if they believe the regulator would have banned the product if the harm were very high. Otherwise, the probability of finding a low enough harm level to warrant purchase might be too small. On the other hand, for consumers with somewhat higher values, they might purchase the product without learning the true harm if they believe the regulator will ban very high harm levels, but they might be induced to learn the true harm level if the regulator is unlikely to ban the product when harm is very high.

This dyanmic creates the potential for multiple equilibria if the regulator cannot commit in advance to a regulatory policy. If consumers believe they are in a lasseiz-faire regime, then less stringent regulation is optimal because relatively low-valued consumers will not purchase the product anyway, those with somewhat higher values will learn the true harm and not purchase it if their value is below that harm level, so banning the product when harm is high will only prevent those with very high values from buying the product.

On the other hand, if consumers believe that they are in a stringent regulatory regime, that regime will be optimal as well. Fairly high-valued consumers will buy the product without researching it, while lower-valued consumers will research the product thinking it is fairly likely to have a low enough harm to warrant purchase. Both of these behaviors will be inefficient if the regulator does not ban the product when harm is high.

We also consider the option to tax the high risk product rather than ban it. Standard analysis would suggest that a tax must be weakly superior to a prohibition since one could think as a prohibition as akin to an infinite tax. In this model, however, this is not the

case because we assume that consumers can more easily make inferences from prohibitions than they can from taxes. If regulators ban products, this easily communicates to consumers that sufficiently risky products are not available for sale. On the other hand, governments tax products for a wide variety of reasons, and many products are subject to many different taxes from different levels of government. Thus, it is much more difficult for a consumer to correctly infer the product risk level from a tax. In the model, therefore, we allow for consumers to make inferences when the government policy is to ban products that are too risky, but we do not allow them to make inferences from the level of tax.

In our model, the product have have three levels of harm (high, low, and zero). We find that the optimal tax on a high-harm product is such that it induces consumers to purchase the product without acquiring information if and only if their value exceeds the marginal cost plus the high harm. This leads to efficient allocation in this case. But, it also leads to greater costs from information acquisition. Banning the high value product, by contrast, eliminates the gain from efficient purchases of the high-harm product, but it reduces information acquisition costs and leads to more efficient purchases when the product is harmless.

If consumers can become informed at very low cost, then the optimal tax is zero and banning the product is inefficient. As the cost of information increases, the optimal tax on the high-harm product becomes positive and banning the high-harm product can be optimal (at information costs either below or above the level at which a positive tax is optimal, depending on the values of o ther parameters). Banning the high-harm product is optimal for a greater range of parameter values as the magnitude of the surplus from buying the high-harm decreases (either through an increase in cost or high-harm). The ban is optimal for a smaller range of parameter values as the magnitude or probability of low-harm increases.

The analysis of optimal regualtion is further complicated when one realizes that the distribution of product risk is likely affected by the expectation of regulation. If the potential seller expects that high risk products are likely to be banned, it may have a greater incentive to invest in safety. On the other hand, it would also have less incentive to attempt to create the product at all. This part of the analysis is left for later drafts.

This analysis has a few important implications. First, and most obviously, is the importance of clearly articulating and sticking to a consistent regualtory policy so that consumers can determine how much research they need to do on their own about product risks. Second, while taxes can influence uninformed consumers' purchase decisions, if they do not send as clear sigals about a product's underlying risk, they may not be as

effective at influencing consumers' decisions to learn about product risks. As a result, in some cases, an outright ban may be superior to taxes because it signals to consumers that the products that are not banned are safer in a way that taxes (or their absence) may not.

This analysis is based on a few assumptions. First, we assume that the government cannot eliminate (or greatly reduce) the consumers' cost of learning the true harm through requiring warnings. While the government can require a simple product wanring such as "this product may cause cancer," for most consumers, this does not provide sufficient information about whether to purchase the product or not. The warning doesn't indicate what the risk of cancer is, how severe the cancer may be, nor how far in the future the risk is likely to occur. Moreover, the answers to all of these questions might depend on personal information that the government does not have. While the government could provide the same data that the consumer can learn on their own, it cannot be done with a simple warning and interpreting that data will still require a substantial level of effort. Thus, in many cases, it is reasonable to assume that while the government ban a product for free, it cannot costlessly inform consumers so that they can make their own, individually-optimal decisions, without further cost to themselves.

Second, we assume that the consumers cannot learn the risk by purchasing a small amount of the product (and, thereby, incurring a proportionately small amount of risk). For risks that are either low probability events or occur after a long period of time, which certainly describes the risks of a great many products this will be the case.

Third, and perhaps most controversially, we assume that consumers cannot make inferences from the price of the product. While in the model, cost of production and consumer demand are known, this is a modeling convenience. In the real world, almost no consumers know a firm's cost of production or the overall demand curve. As a result, consumers cannot observe a product's price and back out the underlying safety of the product. While one could imagine consumers getting a slightly more precise distribution of product risks if they had well-defined priors about demand and costs, the extra precision would be quite minor if these distribuitons were not very precise. An even more realistic model, would include the fact that making such an inference would require significant mental effort for most consumers, so that given the expected benefit, it would be very unlikely be worth it to make the effort.

This paper is related to the literature consumer information acquisition and inferences from disclosures. Chan and Leland (1982) first analyzed markets in which consumers can aquire information about quality at a cost. Milgrom (2008) reviews the literature on persuasion games and what it suggests about mandated disclosure. Bar-Gill et al. (2019) show that consumers often draw the wrong inferences from mandated disclosures,

which highlights another reason why warnings may not be effective. Armstrong (2015) discusses how differently informed consumers interact in the market, though his paper reviews models in which this difference in exogenous, not endogenously chosen.

Zhang (2014) considers the effect of consumer inference on government disclosure policy. In that paper, the consumer doesn't know whether a product has a risky ingredient in it or not. The government's policy choice is whether or not to require firms to disclose whether the product contains this ingerdient or not (at some cost). If the government knows the magnitude of the risk from this ingredient, then the consumer makes an inference about this risk from whether or not the government mandates disclosure. Because there is no option to ban the product, the decision to mandate disclosure signals that the ingredient is more harmful, while the decision not to signals the ingredient is less harmful.

While the Zhang paper shares the insight that consumers make inferences from government regulatory decisions, it differs from this paper in some important ways. First, as mentioned above, the most interventionist policy choice the government can make is to require disclosure of a given ingredient. This necessarily makes the inference from disclosure different than when the government also has the option to ban products entirely. That paper also does not consider taxes as an alternative policy instrument. Second, the consumer decision in binary, to buy or not to buy. They do not have the option to find out the harm at some cost. Thus, it does not examine how the regulation influences the decision to acquire this information.

The next section outlines the model. Section 3 analyzes the optimal tax regime. Section 4 analyzes the ban regime. Section 5 compares the two. Section 6 concludes.

2 Model

A monopoly firm produces a product at constant marginal cost c that causes harm to consumers of either $\{h, l, 0\}$, $1 > h > l > 0$. The fixed cost of production is $f \geq 0$. The firm knows the harm level at the time of sale (and so does a government regulator). Consumers do not know the harm level unless they spend k to learn the harm. Consumers value for one unit of the product (gross of harm and price) is given by $v \in [0, 1]$ with probability density function g and associated distribution function G.

In period -1, the government regulator may or may not be able to set a regulatory policy in advance. We will consider two possible policy choices. First, the regulator must decice whether or not to allow the product with harm h or l for sale in period 1. Second,

the regulator can choose a tax t_h (t_l) to impose on the product that causes harm $h(l)$. If the regulator can set these policies in advance, we assume that the consumers can observe regulations limiting sale but cannot observe taxes. The rationale for this latter assumption is that governments tax products for many reasons, so it is much harder for consumers to make inferences based on taxes than based on prohibitions.

In period 0, the firm chooses an amount s to invest in safety. An investment of s changes the ex ante probability of harm h from q_h to $q_h - s$, the probability of harm l from q_l to $q_l - (1-z)s$, and the probability of harm 0 from $1-q_h-q_l$ to $1-q_h-q_l+$zs. $z>0$ but z could be greater or less than one. The cost of an investment of s is $C(s); C', C'' > 0$.

In period 1, the government and firm observe the harm. The description of the harm is too complex to explain to consumers in an easy to understand warning. If the government was able to commit to a policy in period -1, then it imposes that policy (allowing, banning, or taxing the sale according to the harm). If not, then the government decides whether to allow the firm to sell the product and whether to tax it. The firm then decides whether to pay f to offer the product for sale, and, if so, it chooses its price p. After observing p and (trivially) whether the product is available, consumers decide whether or not to spend k to learn the true harm and then make purchase decisions (fully informed, if they spent k, otherwise not). If the government was not able to commit to a policy in period -1, then I assume the consumers to not directly observe the policy before deciding on whether to become informed and/or buy the product–but, of course, consumers can base these decisions about conjectures about the government policy.

As we will see below, the firm will often have a different optimal price depending on the level of harm if consumers do not make any inferences regarding harm from the price. For simplicity, we will disallow such inferences. As discussed above, the justification, other than simplicity, for this is that in practice, consumers do not know enough about the overall demand curve or the marginal costs to make any precise inference from prices. Note, making inferences from prices is much easier in a non-monopoly model where consumers can make such inferences from differences in prices between firms.

3 Period 1–no ban

In this section, we assume the government has not banned the product from sale and consumers correctly believe this to be the policy. We assume the firm pays any tax, so consumer decisions are simply based on the price p, the cost of information k, the value for the product v, and their beliefs about the probability distribution of harm (r_h, $r_l, 1-r_h-r_l$ for harm levels $h, l, 0$).

3.1 Information acquisition decision

We first analyze the decision to become informed. High-valued consumers are choosing between either buying the product without information or obtaining information If $v > p + h$, information has no value. For consumers with values just below h, the value of information is $r_h(p + h - v)$–these consumers only refrain from purchasing the product if harm is high, and doing so saves them the loss of $p + h - v$. Thus, for these consumers, obtaining information is worth it if and only if $v < p + h - k/r_h$. Note that if $k > r_h((1 - r_h)h - r_l l)$, then the cutoff for obtianing information is below the level at which consumers will purchase the product without information.

All consumers with $v > p + r_h h + r_l l$, those who purchase the product if uninformed, will obtain information if $v < p + h - k/r_h$. But, for those consumers with even lower values who will not purchase the product if uninformed, the value of information now is that if they learn the harm is low enough, they will purchase the product. If $v < p + l$ (the consumer would only purchase the good if it were harmless) then the value of information is $(1 - r_h - r_l)(v - p)$. These consumers only obtain information if $v > p + k/(1 - r_h - r_l)$. But, if $k/(1 - r_h - r_l) > l$, this means the marginal consumers would purchase the product if harm were l, so the value of information is $(1 - r_h - r_l)(v - p) + r_l(v - p - l)$. In that case, consumers obtain information if only if $v > p + (k + r_l l)/(1 - r_h)$. Note that if $k > r_h((1 - r_h)h - r_l l)$, then the cutoff for obtianing information is above the level at which consumers will not purchase the product without information. Thus, for information costs above this level, consumers never obtain information.

Thus, we have the following lemma.

Lemma 1. *If the product is never banned from sale,*

(a) Consumers with $v > p + h - k/r_h$ purchase the product without information.

(b) If $k \leq (1 - r_h - r_l)l$, then all consumers with $v \in [p + k/(1 - r_h - r_l),\ p + h - k/r_h]$ obtain information, if the product is harmless, then all of these consumers buy the product. If harm is *l, then these consumers purchase the product if and only if $v > p + l$, and if the harm is h, then none of these consumers purchase the product. All consumers with $v < p + k/(1 - r_h - r_l)$ do not obtain information or buy the product.*

(c) If $k \in ((1 - r_h - r_l)l, r_h((1 - r_h)h - r_l l)]$, then all consumers with $v \in [p + (k + r_l l)/(1 - r_h),\ p + h - k/r_h]$ obtain information. If harm is *l or 0, then these consumers purchase the product, if the harm is h, then none of these consumers purchase the product. All consumers with $v < p + (k + r_l l)/(1 - r_h)$* do not obtain information *or* buy the product.

(d) If $k > r_h((1-r_h)h - r_l l)$, *no consumers obtain information and consumers purchase the product if and only if* $v > p + r_h h + r_l l$.

Because this problem is interesting only if consumers sometimes acquire information, from here on, we will assume $k \leq r_h((1-r_h)h - r_l l)$.

3.2 Pricing decision

We can use Lemma 1 to derive the demand curve for the product based on harm level. If harm is high, then the only consumers who purchase the product are those who purchase the product without information. Thus, demand is given by $1 - G(p + h - k/r_h)$.

If harm is low and $k \leq (1 - r_h - r_l)l$, then consumers purchase the product whenever $v > p + l$, so demand is $1 - G(p + l)$. If $k > (1 - r_h - r_l)l$, then all consumers who obtain information have a high enough value to purchase when harm is low, so demand is $1 - G(p + (k + r_l l)/(1 - r_h))$.

If harm is zero and $k \leq (1 - r_h - r_l)l$, then all consumers who obtain information or purchase the product without information buy the product. So demand is $1 - G(p + k/(1 - r_h - r_l))$. If $k > (1 - r_h - r_l)l$, again all consumers who obrain information buy the product, so demand is $1 - G(p + (k + r_l l)/(1 - r_h))$.

The following lemma gives profit-maximizing prices when consumer values are uniformly distributed.

Lemma 2. *If* $g \sim U[0, 1]$, *then profit-maximing prices, given taxes of* t_h *and* t_l, *for each harm level are given by:*

(a) $p_h = \frac{1+c+t_h-h}{2} + \frac{k}{2r_h}$; *(b)* $p_l = \frac{1+c+t_l-l}{2}$ *if* $k \leq (1 - r_h - r_l)l$ *and* $p_l = \frac{1+c+t_l}{2} - \frac{k+r_l l}{2(1-r_h)}$ *if* $k > (1 - r_h - r_l)l$; *(c)* $p_0 = \frac{1+c}{2} - \frac{k}{2(1-r_h-r_l)}$ *if* $k \leq (1 - r_h - r_l)l$ *and* $p_0 = \frac{1+c}{2} - \frac{k+r_l l}{2(1-r_h)}$ *if* $k > (1 - r_h - r_l)l$.

Proof. These follow directly from maximizing $(p - t - c)(1 - \alpha(p))$, where $\alpha(p)$ is the minimum consumer value which purchases the product, derived above for each harm level and cost of information. □

Notice that the price of the high harm product is increasing in the cost of information, while the price of the harmless product is decreasing in the cost of information. Both of these result from the fact that increasing information costs increases demand for the high harm product and decreases demand for the harmless product. The price of the low harm product doesn't depend on the cost of information if that cost is small, but once that cost gets high enough, the low harm product price is decreasing in the cost of information.

3.3 Ex post welfare effect of taxes

Now that we have derived prices in the uniform case, we can analyze ex post welfare and determine the effect of taxes. If harm is high and information costs are low, ex post welfare is given by:

$$\int_{p_h+h-k/r_h}^{1}(v-c-h)\mathrm{d}v - k[G(p_h+h-k/r_h)-G(p_h+k/(1-r_h-r_l))] \tag{1}$$

If G is uniform, then this simplifies to:

$$\{r_h^2(1-r_h-r_l)(3(1-c-h)+t)(1-c-h-t)+2kr_h(1-r_h-r_l)(1-c-h-4hr_h+t)- k^2(1-r_l-r_h(9-8r_l))\}/8r_h^2(1-r_h-r_l) \tag{2}$$

This is increasing in t if and only if $k>r_h(1-c-h+t)$. If the expected surplus from the highest valued consumer purchasing the product when the harm is high is smaller than the cost of aquiring information, then it is socially optimal to tax the high harm product at least a little.

If $k>(1-r_h-r_l)l$, then expected welfare when the product harm is high is:

$$\int_{p_h+h-k/r_h}^{1}(v-c-h)\mathrm{d}v - k[G(p_h+h-k/r_h)-G(p_h+(k+r_l l)/(1-r_h))] \tag{3}$$

If G is uniform, then this simplifies to:

$$\{r_h^2(1-r_h)(3(1-c-h)+t)(1-c-h-t)+2kr_h[](1-r_h)(1-c-h-4hr_h+t)-r_l l- k^2(1-9r_{h\,l})\}/8r_h^2(1-r_h-r_l) \tag{4}$$

This is also increasing in t if and only if $k>r_h(1-c-h+t)$.

If harm is low and $k\leq(1-r_h-r_l)l$, then welfare is:

$$\int_{p_l+l}^{1}(v-c-l)\mathrm{d}v - k[G(p_l+h-k/r_h)-G(p_l+k/(1-r_h-r_l))] \tag{5}$$

If G is uniform, then this simplifies to:

$$\{(3(1-c-l)+t)(1-c-l-t)-8k(H-k(1-r_l)/(1-r_h-r_l))/8 \tag{6}$$

The derivative of this with respect to t is -2(1-c-l+t)<0, so it is never optimal to tax the product if harm is low if the cost of information is low.

If harm is low and $k>(1-r_h-r_l)l$, then welfare is:

$$\int_{p_l+\frac{k+r_l l}{(1-r_h)}}^{1}(v-c-l)\mathrm{d}v - k\left[G(p_l+h-k/r_h)-G\left(p_l+\frac{k+r_l l}{(1-r_h)}\right)\right] \tag{7}$$

If G is uniform, the derivative of this with respect to t is $-\left(1-c+t+\frac{k-l(2-2r_h-r_l)}{1-r_h}\right)/4$. This is decreasing in k, and at $k=(1-r_h-r_l)l$, it is -(1-c-l+t)/4<0, so it is never optimal to tax the product if the harm is low. We summarize these in the following lemma.

Lemma 3. *If the product can be sold at any harm level, then it is never optimal to tax the low harm product or the no harm product. It is optimal to tax the high harm product if and only if $k>r_h(1-c-h)$, in which case the optimal tax is $t_h=k/r_h-(1-c-h)$*

If k is small, it is not optimal to tax the high harm product because the monopoly price is already exceeding its social cost. For larger k, it is optimal to tax the high harm product because the monopolist would otherwise set the price such that consumers with a value lower than the social cost, $c+h$, would buy without learning the true harm. The optimal tax then just induces a consumer with value equal to the social cost indifferent to purchasing without acquiring information–will make the socially optimal choice.

4 Period 1–high harm ban

Now let's examine what happens if consumers know the regulator will ban the product if the harm is high. The analysis will be similar except that now beliefs about the probability distribution of harm $(0, r_l^b, 1-r_l^b$ for harm levels $h,l,0)$. If we were to assume that the ex ante probabilities were fixed, this would make $r_l^b=r_l/(1-r_h)$. But, because we want to allow for the ban to affect the ex ante investment decision, this may not be the case.

4.1 Information acquisition decision

If the product is for sale, consumers with $v\geq p+r_l^b l$ choose between buying the product without acquiring information and acquiring information. For these consumers, the value of information is $r_l^b(p+l-v)$ whenever $l>v$. Thus, these consumers acquire information if and only if $v<p+l-k/r_l^b$. For consumers with smaller v, they choose between acquiring information or not purchasing at all, thus their value of information is $(1-r_l^b)(v-p)$. These consumers acquire information if and only if $v>p+k/(1-r_l^b)$. Thus, we have the following lemma.

Lemma 4. *If the regulator will ban the product when the harm is high:*

(a) Consumers with $v\geq p+l-k/r_l^b$ buy the product without acquiring information.

(b) Consumers with $v\in[p+k/(1-r_l^b), p+l-k/r_l^b)$ spend k to learn the true harm and buy the product if and only if it is harmless.

(c) Consumers with $v\leq p+k/(1-r_l^b)$ do not buy the product.

Notice that consumers only ever acquire information if $k < r_l^b(1-r_l^b)l$.

4.2 Pricing decision

If harm is low, then demand only comes from consumers who do not acquire information, so is $1-G(p+l-k/r_l^b)$. If the product is harmless, then all consumers who acquire information also buy the product, so demand is $1-G(p+k/(1-r_l^b))$.

The following lemma gives profit-maximizing prices when consumer values are uniformly distributed.

Lemma 5. *If $g \sim U[0,1]$, then profit-maximing prices if the regulator bans the high harm product, given taxes of t_l^b, for the low harm product are given by:*

(a) $p_l = \frac{1+c+t_l^b-l}{2} + \frac{k}{2r_l^b}$ *(b)* $p_0 = \frac{1+c}{2} - \frac{k}{2(1-r_l^b)}$.

Proof. These follow directly from maximizing $(p-t-c)(1-\alpha(p))$, where $\alpha(p)$ is the minimum consumer value which purchases the product, derived above for each harm level and cost of information. □

The low harm product's price is now increasing in the cost of information while the harmless product's price is decreasing in the cost of information.

4.3 Ex post welfare effect of taxes

When the product harm is low, welfare is given by:

$$\int_{p+l-k/r_l^b}^{1} (v-c-l)g(v)\mathrm{d}v - k[G(p+l-k/r_l^b) - G(p+k/(1-r_l^b))] \tag{8}$$

If G is uniform, the derivative of this with respect to t_l^b is $(k-r_l^b(1-c-l+t_l^b))/4r_l^b$. Thus, some taxation of the low harm product is optimal if and only if $k > r_l^b(1-c-l)$. This condition is similar to the condition for optimal tax of the high harm product. The fact that harm is smaller in this case makes the right hand side larger, so makes some tax less likely to be optimal. If $r_l^b > r_h$, then this conclusion is reinforced. But, if $r_l^b < r_h$, then it's possible that it is optimal to tax the low harm product when the high harm product is banned even if taxing the high harm product is not optimal.

5 Welfare comparison: ban v. taxes

In this section, we compare the alternative of taxing the product when the harm is high versus banning it. Note that we assume that consumers cannot make inferences about the harm level from the tax, but they can if the government announces a clear ban. Thus, it

is not the case that a sufficiently large tax is equivalent to a ban. Under this assumption, the advantage of a ban is that it provides more information to consumers when the product is not banned. On the other hand, the the optimal tax ensures that only those consumers who value the product above $c+h$ purchase it, leading to efficient allocation when harm is high (even though there is still inefficient information gathering among lowered value consumers who will not purchase the product).

To analyze this formally, note that expected welfare under the optimal tax is:

$$r_h \int_{p_h+h-k/r_h}^{1} (v-h-c)\mathrm{d}v + r_l \int_{p_l+l}^{1} (v-l-c)\mathrm{d}v + (1-r_h-r_l) \int_{p_0+k/(1-r_h-r_l)}^{1} (v-c)\mathrm{d}v - k\left(h - \frac{k(1-r_l)}{r_h(1-r_h-r_l)}\right) \tag{9}$$

The prices are given in Lemma 2 and the optimal tax in Lemma 3. The integral terms represent the welfare from consumption for each possible harm level. The last term is the welfare loss from the cost of acquiring information.

Expected welfare under the policy of banning the product when harm is high is given by:

$$r_l^b(1-r_h^b) \int_{p_l+l-k/r_l^b}^{1} (v-l-c)\mathrm{d}v + (1-r_l^b)(1-r_h^b) \int_{p_0+k/(1-r_l^b)}^{1} (v-c)\mathrm{d}v - k\left(l - \frac{k}{r_l^b(1-r_l^b)}\right) \tag{10}$$

Notice that because we have defined r_l^b in the ban case as conditional on the product being sold, for welfare analysis, we have to convert them to unconditional probabilities for welfare analysis. Abusing notation somewhat, we define r_h^b as the unconditional probability of the high harm state when the high harm product is banned.

For the remaninder of this section, we will confine our analysis to the case of exogenous probabilities (i.e., we will will not allow for the regulation to affect the type of product developed), in which case we have $r_h^b = r_h$ and $r_l^b = r_l/(1-r_h)$. Prices are given by Lemma 5. We will assume that $k \leqslant r_l^b(1-c-l)$, so taxes are not optimal when the high harm product is banned.

If $k \leq (1-r_h-r_l)l$, then the welfare difference between banning when harm is high and allowing it with the optimal tax is:

$$\left\{-4(1-c-h)^2 r_h + 2k(1-c-h+5(h-l(1-r_h))) - k^2 \frac{8(1-r_l)r_l - 7(1-r_h)^3 r_h - r_h r_l}{r_h r_l (1-r_h-r_l)}\right\} / 8 \tag{11}$$

This expression assumes that at least some consumers obtain information in each case. For this to be case for all $k \leq (1-r_h-r_l)l$, we much have $r_l \geq (1-r_h)^2$ (ensures that some consumers obtain information in the ban case) and $r_l \geq (l-r_h h)/l$ (ensures that some consumers obtain information in the no ban case). The first condition is binding if and only if $l < \frac{h}{2-r_h}$. It also assumes that it is optimal to tax the high harm product when it is not banned, so $k \geq r_h(1-c-h)$. If this is not the case, then the welfare difference (with no tax on the high harm product) is:

$$\left\{-3(1-c-h)^2 r_h + 10k(h-l(1-r_h)) - 7k^2 \frac{(1-r_l)r_l - (1-r_h)^3 r_h}{r_h r_l (1-r_h-r_l)}\right\}/8 \tag{12}$$

If k is sufficiently small, it is clearly optimal to allow the sale of the product for any harm and potentially tax it when harm is high. On the other hand, if the surplus from purchasing the product for the highest value consumer, $1-c-h$, is sufficiently small, then banning the product generates more welfare for sufficiently small k. The following proposition shows how the various parameters affect the optimality of the ban versus the allowing the product with the optimal tax (which could be zero).

Proposition 6. *Banning the product is relatively more desirable for larger h or c (if $k < \text{Max}[4r_h(1-c-h), (1-r_h-r_l)l]$), smaller l or r_l (if $k < (1-r_h-r_l)l$).*

Proof. For $k \leq (1-r_h-r_l)l$, inspection reveals that both 11 and 12 are increasing in h and c and decreasing in l. Differentiating both 11 and 12 with respect to r_l gives

$$-\frac{7k^2[(1-r_h)^4 + r_l(r_l - 2(1-r_h)^3)]}{8r_l^2(1-r_h-r_l)^2} < 0$$

Now consider the case of $k > (1-r_h-r_l)l$, so that all consumers who obtain information purchase the product as long as harm is not high. The welfare difference between banning when harm is high and allowing it with the optimal tax is now:

$$\left\{\frac{3r_l(1-r_h-r_l)l^2}{1-r_h} - 4r_h(1-c-h)^2 + \frac{2k}{1-r_h}[(1-r_h)(1-c-h) + 5((1-r_h)h - ((1-r_h)^2 + r_l)l)] - k^2\left[\frac{8-r_h}{(1-r_h)r_h} - \frac{7(1-r_h)^2}{r_l(1-r_h-r_l)}\right]\right\}/8$$

Differentiating this with respect to h gives a squared denominator and a numerator whose sign depends on $r_h(1-c-h)+k>0$. Differentiating this with respect to c again yields a squared denominator and a numerator whose sign depends on $4r_h(1-c-h)-k>0$. The derivative with respect to l has the sign of $3lr_l(1-r_h-r_l)-5k((1-r_h)^2+r_l))<0$ whenever $k>(1-r_h-r_l)l$.

□

The advantage of taxing the high harm product is that it enables efficient allocation when harm is high. The advantage of banning the high-harm product is that then if the product is for sale, consumers know the harm either low or zero. This improves both their information acquisition decision in this case and their purchase decision. When there is a lot of surplus from buying the high harm product, the tax option tends to be superior. On the other hand, if knowing the product isn't high harm makes purchase optimal for a large number of consumers (l or r_l small), then providing that information is very valuable for consumers. The effect of the cost of information can be non-monotonic because when costs are very low there is no need for the regulator to ban the product because it is easy for consumers to find out if purchase is optimal or not. But, if costs are very high, then consumers may not change their information acquisition decision much if they learn harms are not high, reducing the benefit from the ban.

6 Conclusion

As consumers, we often purchase products without precise knowledge of any possible adverse effects they may have on us. These may be health effects associated with food, pesticides, or many other products. These could be financial risks associated with complex financial products. These could be accident risks associated with appliances we might purchase.

While in principle, product liability might could make us indifferent to our lack of knowledge, in practice that is rarely the case. Suits are expensive, causation is difficult to prove, and for harms that come regardless of a design defect, recovery may not be available at all, especially if there was some type of warning. If product liability were perfect, of course, there would be no reason to regulate dangerous or hazardous products.

Because liability and warnings are imperfect, consumers are often left with a choice not only whether to purchase a product or not but also whether to gather more information about the risk. Gathering such information is costly. In this paper, we show that this

cost can affect both the optimal tax for dangerous products and whether or not it is optimal to ban them outright. When government can commit to banning dangerous products, this provides consumers with valuable information about the hazards of products that are still on the market. Because making similar inferences from taxes is much harder, we show that sometimes committing to ban products with only small amounts of possible social surplus can be optimal ex ante, even though it foregoes possible ex post social surplus associated with an optimal tax regime.

Although not yet analyzed in this version, another possible benefit to committing to banning dangerous products is that it can possibly increase the marginal private benefit to investing in making products safer. On the other hand, it also might make the expected private benefit of developing a new product less than cost, further complicating the optimal regulatory policy.

Lastly, this analysis has some implications for changing regulatory policy. When regulatory preferences or information changes, optimal regulatory policy might remain sticky. For example, say new regulators believe the fraction of very high-valued consumers (who buy without research) is somewhat lower than previous regulators did. This might mean that more restrictive regulation would have been optimal initially, but it might not warrant banning exisitng products if many consumers have already incurred the research costs of learning those risks.

Because optimal regulation depends both on the risks of the product and the distribution of benefits, changing regulation because of new information about one of these factors also runs the risk of sending the wrong message to consumers. For example, if new regulators place more weight on the importance of enabling purchases by high-valued consumers, and thereby allow the sale of previously banned products, some consumers might wrongly infer that the change was due to new information about the harmfulness of the product. While it is possible to mitigate this by accurately conveying the reasons for the change, it is unlikely to fully resolve the problem. This suggests another reason why optimal regulation might be sticky, at least with respect to changes in other paramters aside from harmfulness.

Bibliography

Armstrong, Mark, "Search and Rip-off Externalities," 47 Review of Industrial Organization 273-302 (2015).

Bar-Gill, Oren, David Schkade, and Cass R. Sunstein, "Drawing false inferences from mandated disclosures," 3 Behavioral Public Policy 209-227 (2019).

Chan, Yuk-Shee and Hayne Leland, "Prices and Qualities in Markets with Costly Information," 49 Review of Economic Studies 499-516 (1982).

Milgrom, Paul, "What the Seller Won't Tell You: Persuasion and Disclosure in Markets," 22 Journal of Economic Perspectives 115-131 (2008).

Zhang, Juanjuan, "Policy and Inference: The case of product labeling," Working paper available at https://www8.gsb.columbia.edu/programs/sites/programs/files/finance/COI.pdf (2014).