False Advertising

Rhodes, Andrew and Wilson, Chris M

Toulouse School of Economics, School of Business and Economics, Loughborough University

18 July 2016

False Advertising

Andrew Rhodes and Chris M. Wilson*

July 18, 2016

Abstract

There is widespread evidence that some firms use false advertising to overstate the value of their products. We consider a model in which a policymaker is able to punish such false claims. We characterize an equilibrium where false advertising actively influences rational buyers, and analyze the effects of policy under different welfare objectives. We establish precise conditions where policy optimally permits a positive level of false advertising, and show how these conditions vary intuitively with demand and market parameters. We also consider the implications for product investment and industry self-regulation, and connect our results to the literature on demand curvature.

Keywords: Misleading Advertising; Product Quality; Pass-through; Self-Regulation

JEL codes: M37; L15; D83

*Rhodes: Toulouse School of Economics, France; andrew.rhodes@tse-fr.eu. Wilson: School of Business and Economics, Loughborough University, UK; c.m.wilson@lboro.ac.uk. We would like to thank Alexei Alexandrov, Mark Armstrong, Mike Baye, Daniel Garcia, Matthew Gentzkow, Thomas Jeitschko, Justin Johnson, Bruno Jullien, Martin Peitz, Michael Powell, Robert Ritz, Jean Tirole, Tianle Zhang, and various audiences including those at the North American Econometric Society (Philadelphia), CEPR (Zurich), ESSET (Gerzensee), IIOC (Boston), EEA (Mannheim), CREST, HECER, THEMA, Toulouse, the Berlin IO Day, the 7th Workshop on the Economics of Advertising and Marketing (Vienna), EARIE (Munich) and the NIE Workshop on Advertising (Manchester). We also thank Kamya Buch for her research assistance.

1 Introduction

Buyers are often reliant on firms to obtain information about product characteristics. To exploit this, some firms deliberately engage in what we call false advertising - the use of incorrect or exaggerated product claims. They do this in a range of different contexts and despite potential legal penalties[1]. Recent policy cases include Dannon which paid \$21m to 39 US states after it misled consumers about the health benefits of its Activia yogurt products, Skechers which paid \$40m after falsely stating that its toning shoes helped with weight loss, and the manufacturer of Nurofen which was fined 1.7m Australian dollars after making false claims about the effectiveness of its painkillers. Similarly, in some related examples, car manufacturers such as Volkswagen and Mitsubishi are facing multi-billion dollar penalties after cheating tests in order to make false claims about their emission levels or fuel efficiency[2]. Additional evidence also comes from academic research which carefully documents the existence of false advertising and its ability to increase demand[3].

However, despite the existence of false advertising, the theoretical literature has largely restricted attention to truthful advertising. In this paper, we develop an equilibrium model of false advertising where false advertising can actively influence rational buyers. Tougher legal penalties reduce the frequency of false adverts, but also increase their credibility. As a result of the latter effect, we show when and how stronger penalties can reduce buyer and social welfare. In particular, by using some results on demand curvature, the paper derives precise conditions on demand and market parameters such that a policymaker optimally uses a low penalty to permit a positive level of false advertising. We then consider several wider issues including investment incentives, and the potential optimality of industry self-regulation.

In more detail, Section 2 introduces our main model where a monopolist is privately informed about its product quality. While we later extend the results in a number of ways, we initially focus on the case where quality is either 'high' or 'low' and where the two types have symmetric marginal costs. The policymaker first commits to a penalty for false advertising. Then having learned its type, the firm chooses a price and makes a (possibly false) claim about

[1] In the US, most federal-level regulation is conducted by the FTC which punishes offenses with various public measures, including possible monetary penalties. In Europe, most countries employ varying levels of industry self-regulation alongside statutory regulations. For instance, in the UK, most regulation is conducted by the industry-led Advertising Standards Authority. It is endorsed by various governmental bodies, which have the power to issue fines.

[2] For further details, see http://www.ftc.gov/news-events/press-releases/2010/12/dannon-agrees-drop-exaggerated-health-claims-activia-yogurt, http://www.ftc.gov/news-events/press-releases/2012/05/skechers-will-pay-40-million-settle-ftc-charges-it-deceived, http://www.bbc.co.uk/news/business-36167011, http://www.bbc.co.uk/news/business-36651853, and http://www.bbc.co.uk/news/business-36137719. Accessed 07/18/16.

[3] E.g. Zinman and Zitzewitz (2016), Rao and Wang (2015), and Cawley *et al* (2013). See also Mayzlin *et al* (2014) for false advertising in the form of fake user reviews.

its quality. Buyers subsequently update their beliefs and make their purchase decisions, before the policymaker instigates any penalties. We believe this set-up closely approximates many important markets where buyers are unable to verify claims, or can only do so after a long time, and where policy plays a key role in regulating advertising.

Section 3 characterizes an equilibrium where the high type advertises truthfully and where the low type may engage in false advertising. This equilibrium does not resemble a standard price signaling equilibrium due to the assumption of symmetric costs. Instead, it smoothly unifies several otherwise separate cases depending on the level of penalty. Firstly when the policymaker's penalty is large, there is no false advertising. Here, as in the standard disclosure literature, quality claims are fully verifiable and product quality is perfectly revealed within a full separating equilibrium. Secondly when the penalty is small, the low type always conducts false advertising within a full pooling equilibrium. In this case, advertising is cheap talk and so buyers simply maintain their prior beliefs even after observing a high claim. Finally, when the penalty is moderate, our equilibrium involves a novel form of partially verifiable advertising. Here, the low type engages in false advertising probabilistically by mixing between i) pooling with the high type, with a false advert and a relatively high price, and ii) advertising truthfully with a relatively low price. Therefore when buyers observe a high claim, they positively update their belief that quality is high. Hence, in contrast to full separation, false advertising does arise in equilibrium but unlike in full pooling, advertised claims do still provide buyers with some new information about the firm's quality.

Section 4 then analyzes how marginal changes in the level of penalty affect a variety of welfare measures. We first consider buyer surplus. Here, a reduction in the penalty increases the probability of false advertising and generates two opposing effects. The first 'persuasion' effect harms buyers by prompting them to buy too many units of a product at an inflated price. The second effect derives from the impact of false advertising on damaging the credibility of claims. Understanding the impact on credibility goes back to at least Nelson (1974) and is well-documented empirically (e.g. Darke and Richie 2007). However, instead of viewing this impact as detrimental, we document a beneficial 'price' effect whereby false advertising counteracts monopoly power by lowering buyers' quality expectations and prompting lower prices.

To compare these two effects under a relatively general form of demand, we then utilize some recent results on demand curvature and cost pass-through (Weyl and Fabinger 2013). Such results are being used in an increasing number of applications, such as price discrimination (e.g. Chen and Schwartz 2015, Cowan 2012). However, rather than focusing on cost changes, we analyze the impact of changes in quality on price, which we term as 'quality pass-through'. In many cases, we show how the persuasion effect dominates such that buyer surplus is maximized by eliminating false advertising. However, we also formalize a set of parameter conditions where the price effect dominates. Here, the optimal penalty is softer so as to induce a positive level of false advertising. Such parameter cases include those where product quality levels are high or where the probability of high quality is large.

Next, we turn to the effect on profits. Unsurprisingly the low (high) type always prefers smaller (larger) penalties. Interestingly however, ex ante, the monopolist weakly prefers strong penalties to eradicate false advertising so that it can price more effectively under high quality. Hence, if the monopolist could commit (perhaps via some third-party), its choice of penalty would coincide with that preferred by buyers in many circumstances. This offers potential support for Europe's use of self-regulation. However, for other circumstances, in contrast to the view that self-regulation may be too soft, the monopolist's preferred penalty is too strong relative to buyers' preferences.

Lastly, we consider total welfare. Here, an increase in the probability of false advertising leads to two different effects. On the one hand, false advertising lowers the credibility of any high claim and so prompts any type with such a claim to further reduce its output below the socially desirable level. However, on the other hand, false advertising also allows the low type to expand its output. We then characterize a set of parameter conditions where this latter output expansion is beneficial and dominates the former effect, such that a positive level of false advertising is welfare optimal.

Section 5 extends the main model to consider the additional effects of false advertising when product quality is endogenous. This is important to consider because false advertising can reduce product quality investment by limiting the available returns from high quality products. However, while we confirm that such an 'investment' effect prompts the policymaker to select a weakly higher penalty, we further show how a positive level of false advertising can remain optimal for buyer and social welfare. In addition, once quality is endogenous, we also show how a penalty increase can raise the probability of false advertising.

Finally, Section 6 considers some robustness issues. First, we allow for an arbitrary number of quality types. Qualitatively the analysis remains the same as for the two-type case except now the policymaker must also decide which types can engage in false advertising. We show how false advertising can remain optimal for types with 'moderate' quality. Second, we let the firm types vary in marginal costs. Here, the resulting equilibrium and policy results remain qualitatively robust after making stronger requirements on buyer beliefs. Third, we examine a competitive context where an incumbent faces an entrant with private product quality. We demonstrate an equilibrium with false advertising that is qualitatively similar to monopoly, and provide related policy results for buyer surplus, profits, and total welfare.

Related Literature: The advertising literature typically focuses on truthful advertising (e.g. Anderson and Renault 2006, Johnson and Myatt 2006[4]). In earlier work, Nelson (1974) offered a seminal discussion of false advertising and how regulation may increase its credibility. However, since then, false advertising and its regulation has only been considered by some recent papers.

Some papers assume that buyers are naive and so believe all claims (e.g. Glaeser and Ujhelyi 2010, Hattori and Higashida 2012). Here, false advertising raises firms' profits, lowers consumer surplus, and can increase total welfare by offsetting the output distortion from imperfect competition. Other papers seek to endogenize advertising credibility with rational buyers. One set of papers, including some within marketing, introduce heterogeneous tastes so that claims can gain credibility by forfeiting revenues from some buyers (e.g. Chakraborty and Harbaugh 2014). However, these papers place little emphasis on policy. Another set of papers examine legal penalties in ways more related to our paper. Corts (2013, 2014a, 2014b) examines an information acquisition setting where a monopolist knows its type but finds it costly to learn its precise product quality. Total welfare is assumed to be increasing in the fineness of consumer information. For low learning costs, total welfare is maximized with high penalties that induce the firm to learn its quality and eradicate false advertising. However when learning costs are higher, such that the firm chooses not to know its quality, high penalties can damage total welfare by discouraging the use of advertising. Instead, lower penalties can raise welfare by inducing the firm to use speculative quality claims to signal its type even when such claims are false with positive probability. Under a different mechanism, Piccolo *et al* (2015) demonstrate that false advertising can maximize buyer surplus in a duopoly where the firms have different product qualities. They show that it is always optimal to use zero 'laissez-faire' penalties to induce full pooling rather than full separation. Intuitively, this creates a downward competitive pressure on prices by making the firms appear undifferentiated. However, after studying a monopoly version of the model, the authors suggest that this result can only arise in a competitive context.[5]

Our research differs from this previous literature in a number of important respects. Firstly, we analyze a richer class of semi-pooling equilibria. Unlike the full pooling and separating equilibria in past models, these equilibria allow for false claims to arise in equilibrium, and to inflate buyers' beliefs beyond their priors. Secondly, by using a more general form of demand, we highlight a novel role for demand curvature not present in other papers. This also enables us to provide precise conditions for when false advertising is beneficial, and to show how the

[4]See also the comprehensive literature reviews by Bagwell (2007) and Renault (2015).

[5]In more distant work, Daughety and Reinganum (1997) study how punitive damages for false safety claims can help support separating equilibria when firms face liability losses for product defects, Barigozzi *et al* (2009) examine false comparative advertising, and Drugov and Troya-Martinez (2015) analyze false advice where firms can also choose the vagueness of their claims.

optimal penalty varies intuitively with demand and market parameters. Thirdly, we show that competition is not necessary for false advertising to maximize buyer surplus. Moreover rather than focusing on just one objective, we provide a unified set of results for buyer surplus, total welfare and firm profits.

Our paper is related to a number of broader areas. First, it adds to the growing literature on the economics of consumer protection policy which considers other topics, such as high-pressure sales tactics (Armstrong and Zhou 2016) and refund rights (Inderst and Ottaviani 2013). Second, our model relates to a number of communication papers that study equilibrium lying and persuasion under full rationality (e.g. Kartik 2009, and Kamenica and Gentzkow 2011). In contrast, we study policy-related lying costs within a specific advertising context, where a third-party influences not only the amount of information that is communicated to buyers but also indirectly the price that they pay. Third, our paper is related to the literatures on price signaling and quality disclosure (see Bagwell 2007 and Dranove and Jin 2010 for respective reviews). These literatures also note that total welfare can be maximized by some form of pooling due to either i) the output distortions from price signaling under separation, or ii) positive costs of disclosure. However, our results provide clear conditions to characterize when (partial) pooling can be buyer- or welfare-optimal even when truthful disclosure is costless, and when full separation does not require the firm to distort its output away from the full information level.

2 Model

A monopolist sells one product to a unit mass of potential buyers. The monopolist is privately informed about its product quality q. Specifically, the product is of low quality L with probability $x \in (0, 1)$, and of high quality H with probability $1 - x$, where $-\infty < L < H < \infty$. Average ex ante quality is then defined as $\bar{q} = xL + (1 - x)H$. For our main analysis we assume that marginal costs are independent of quality and normalized to zero. Each buyer has a unit demand and values a given product of quality q at $q + \varepsilon$, where ε is a buyer's privately known match with the product. This match is drawn independently across buyers using a distribution function $G(\varepsilon)$ with support $[a, b]$ where $-\infty \leq a < b \leq \infty$. The associated density $g(\varepsilon)$ is strictly positive, continuously differentiable, and has an increasing hazard rate.

The monopolist sends a publicly observable advertisement or 'report' $r \in \{L, H\}$ at no cost, where a report $r = z$ is equivalent to a claim "Product quality is z". The binary report space is without loss because there are only two firm types and reports are costless. A policymaker is able to verify any advertised claim and impose a penalty ϕ if it is false, where false advertising is defined as the use of a high quality report $r = H$, by a firm with low quality $q = L$. The penalty ϕ can be interpreted as an expected penalty if claims are only verified probabilistically. The policymaker can costlessly choose any level of penalty, $\phi \geq 0$, in order to maximize one of three possible objectives: buyer surplus, total profit, or total welfare. Any penalties that involve a fine go to the policymaker.[6]

The timing of the game is as follows. At stage 1 the policymaker publicly commits to a penalty ϕ for false advertising. At stage 2 the monopolist privately learns its quality. It then announces a price p and issues a report $r \in \{L, H\}$. At stage 3 buyers decide whether to buy the product, taking into account ϕ as well as the firm's price and report. Finally at stage 4 the policymaker verifies the advertised claim and administers the penalty, ϕ, if it is false. The solution concept is Perfect Bayesian Equilibrium (PBE). All omitted proofs are included in the appendix unless stated otherwise.

3 Equilibrium Analysis

3.1 Benchmark with Known Quality

As a first step, consider a benchmark case in which the firm is known to have quality q. Quality claims are then redundant because it is weakly optimal for the firm to use truthful advertising. An individual buyer purchases the product if and only if $\varepsilon \geq p - q$ such that demand equals $D(p-q) = 1 - G(p-q)$. The firm then chooses its price to maximize $p\left[1 - G\left(p - q\right)\right]$, and so:

Lemma 1. *Suppose the firm is known to have quality q, and define $\underset{\sim}{q} = -b$ and $\tilde{q} = -a + 1/g(a)$. The firm's optimal price, $p^*(q)$, is increasing in q and satisfies:*

$$p^*(q) = \begin{cases} 0 & \text{if } q \leq \underset{\sim}{q} \\ \frac{1-G(p^*(q)-q)}{g(p^*(q)-q)} & \text{if } q \in \left(\underset{\sim}{q}, \tilde{q}\right) \\ a + q & \text{if } q \geq \tilde{q} \end{cases} \tag{1}$$

When $q \leq \underset{\sim}{q}$, quality is so low that the firm would make zero sales even if it priced at marginal cost. The market is inactive, and we normalize the firm's price to zero without loss. When instead $q \in (\underset{\sim}{q}, \tilde{q})$, the firm optimally sells to some but not all buyers such that $p^*(q)$ satisfies the usual monopoly first order condition. Finally if $q \geq \tilde{q}$, quality is so high that the firm optimally sells to all potential customers by pricing at the willingness-to-pay of the marginal buyer, $a+q$, such that the market is 'covered'. However, for some distributions, $\tilde{q} = \infty$ and so this final case is redundant. Henceforth to avoid some uninteresting cases, let $\bar{q} > \underset{\sim}{q}$ (or $\bar{q} + b > 0$) such that a product of average quality always has some positive value.

[6] More broadly, one could also interpret ϕ to include any direct costs of making a false claim. For instance, as in the case of Volkswagen, these could include the costs of falsifying emissions tests.

Our later analysis will consider how the optimal price varies with quality, and we will sometimes refer to $dp^*(q)/dq$ as 'quality pass-through'. Firstly for $q \in (\underset{\sim}{q}, \tilde{q})$, after differentiating the first order condition:

$$\frac{dp^*(q)}{dq} = \frac{1 - \sigma(p^*(q) - q)}{2 - \sigma(p^*(q) - q))}, \tag{2}$$

where $\sigma(\psi) = -[1 - G(\psi)]g'(\psi)/g(\psi)^2$ is the curvature of demand (see Aguirre *et al* 2010, and Weyl and Fabinger 2013). It then follows that $dp^*(q)/dq \in [0, 1)$ because our assumption of an increasing hazard rate implies that $D(p - q)$ is logconcave in price, such that $\sigma(\psi) \leq 1$. Intuitively, an increase in quality q produces a parallel outward shift in the inverse demand curve, and the firm optimally responds by both charging a higher price and by selling to strictly more buyers.[7] Secondly, where appropriate, when $q \geq \tilde{q}$ quality pass-through is one.

The equilibrium profit earned by a firm of known quality q can be written as

$$\pi^*(q) = p^*(q) \left[1 - G\left(p^*(q) - q\right)\right]. \tag{3}$$

It is straightforward to show that $\pi^*(q)$ is increasing and convex in q given that $g(\varepsilon)$ has an increasing hazard rate. Finally, buyer surplus can be expressed as

$$v^*(q) = \int_{p^*(q)}^{b+q} \left[1 - G\left(z - q\right)\right] dz. \tag{4}$$

Observe that $v^*(q) = 0$ when $q \leq \underset{\sim}{q}$ because no buyer purchases the product, but that buyer surplus is positive and weakly increasing in $q > \underset{\sim}{q}$. We further discuss the shape of $v^*(q)$ in Section 4.1 below.

3.2 Privately-Known Quality

Henceforth we assume that the firm is privately informed about its quality. As is typical in signaling games, there exists a large number of PBE because buyers can attribute any off-path claim or price to the low type. We therefore proceed as follows. Firstly we only consider PBE in which the high type makes a truthful claim with probability one. Although there do exist PBE in which the high type reports $r = L$, these equilibria require buyers to perversely believe that a firm with a higher claim has lower quality; such beliefs do not satisfy common refinements such as D1. Secondly then, if a firm reports $r = L$ it is reasonable for buyers to believe that quality is low with probability one. Therefore if the low type reports $r = L$ it should charge its full information price $p^*(L)$, and hence earn $\pi^*(L)$. If instead the low type reports $r = H$, it

[7] Weyl and Fabinger (2013) note that the optimal price change from any outward unit shift in inverse demand (such as an increase in quality within our model) equals one minus cost pass-through.

must charge the same price as the high type, otherwise buyers would correctly infer its type and it would earn (weakly) less than $\pi^*(L)$. Thirdly, we need to determine the price charged by the firm when it reports $r = H$. Many such prices can form a PBE even after applying standard refinements such as D1, due to the symmetry in marginal costs.[8] Therefore to make further progress in selecting amongst these prices, we restrict attention to buyer beliefs which depend only on the firm's claim and not its price. This implies that the firm should charge $p^*(q_H^e)$ when it makes a high report, where $q_H^e \equiv E(q|r = H)$ denotes buyer beliefs after observing a high report. One justification for this restriction is as follows. Notice that conditional on sending a high report, the payoff functions of the two types differ only by the penalty ϕ. Therefore buyers might reasonably expect the two types to price in the same way after reporting $r = H$, and therefore avoid making any price-based inferences.[9] We can then state:

Proposition 1. *Suppose a high type always reports truthfully, and buyer beliefs depend only on the firm's claim. There exists a unique PBE*[10]*, in which:*
i) A high type claims $r = H$ *and charges* $p^*(q_H^e)$.
ii) A low type randomizes. With probability y^* *it claims* $r = H$ *and charges* $p^*(q_H^e)$. *With probability* $1 - y^*$ *it claims* $r = L$ *and charges* $p^*(L)$.
- When $\phi \leq \phi_1 \equiv \pi^*(\bar{q}) - \pi^*(L)$, $y^* = 1$
- When $\phi \geq \phi_0 \equiv \pi^*(H) - \pi^*(L)$, $y^* = 0$
- When $\phi \in (\phi_1, \phi_0)$, $y^* \in (0, 1)$ *and uniquely solves*

$$\pi^*(q_H^e) - \phi = \pi^*(L), \tag{5}$$

$$\text{where} \quad q_H^e = \frac{xy^*L + (1-x)H}{1 - x + xy^*}. \tag{6}$$

iii) Buyer beliefs are $\Pr(q = H|r = L) = 0$ *and* $\Pr(q = H|r = H) = \frac{1-x}{1-x+xy^*}$.

[8] When $\phi \geq \pi^*(H) - \pi^*(L)$, D1 selects a fully separating equilibrium in which the low type plays $\{r = L, p = p^*(L)\}$, whilst the high type plays $\{r = H, p^*(H)\}$. However when $\phi < \pi^*(H) - \pi^*(L)$ many equilibrium outcomes are consistent with D1. Among others, this includes a continuum of semi-pooling equilibria, where for different levels of some price p^{sp} the high type plays $\{r = H, p = p^{sp}\}$, and the low type randomizes between $\{r = L, p = p^*(L)\}$ and $\{r = H, p = p^{sp}\}$.

[9] Another justification is as follows. First, consider a different game G in which low and high types only signal via their price, and have payoff functions $\Pi(p, q^e) - \phi$ and $\Pi(p, q^e)$ respectively where q^e is buyers' belief about quality. This game has many PBE and D1 has no bite. However Mailath *et al*'s (1993) Undefeated refinement selects a unique pooling equilibrium with price $\arg\max \Pi(p, q^e)$, and is consistent with price-independent buyer beliefs. Now return to our game. Suppose buyers expect the low type to report $r = H$ with probability $y \in (0, 1]$, and interpret any off-path price as a possible signal about the firm's type but not about y. Then conditional on reporting $r = H$, the firm can be thought of as playing game G when choosing its price, and thus both types should pool on $p^*(q_H^e)$, with buyers' beliefs being independent of price.

[10] In particular, the equilibrium is unique up to off-path beliefs: when $\phi < \phi_1$ the claim $r = L$ is off-path, and a range of beliefs $\Pr(q = L|r = L)$ lead to the same equilibrium play.

The exact nature of the equilibrium in Proposition 1 depends intuitively on whether the penalty for false advertising is 'low', 'medium', or 'high'. Firstly, if $\phi \geq \phi_0 \equiv \pi^*(H) - \pi^*(L)$ the equilibrium has full separation, with both types reporting truthfully because the penalty is sufficiently high that truth-telling is a dominant strategy for the low type. Advertising is perfectly informative and claims are fully believed by buyers. Consequently, each type $q \in \{L, H\}$ sets its full information price $p^*(q)$, and earns its full information profit $\pi^*(q)$. Secondly, if $\phi \leq \phi_1 \equiv \pi^*(\bar{q}) - \pi^*(L)$ the equilibrium has full pooling, with both types sending a high report because the penalty is sufficiently low that false advertising is a dominant strategy for the low type. Upon seeing a high claim buyers then maintain their prior belief that the firm's average quality is $\bar{q}$. As such, both types set $p^*(\bar{q})$ and earn profits of $\pi^*(\bar{q})$ and $\pi^*(\bar{q}) - \phi$ respectively. Finally and most interestingly, if $\phi \in (\phi_1, \phi_0)$ the equilibrium is semi-pooling. Here the low type randomizes, making a false report with probability $y^* \in (0, 1)$ and a truthful report with complementary probability $1 - y^*$. Upon seeing a high report buyers therefore update their belief about quality to

$$q_H^e \equiv E(q|r = H) = \frac{xy^*L + (1 - x)H}{1 - x + xy^*} \in (\bar{q}, H).$$

Consequently the high type reports $r = H$ and charges $p^*(q_H^e)$, and the low type randomizes between $r = L$ and price $p^*(L)$, and report $r = H$ and price $p^*(q_H^e)$. The probability with which the low type lies, y^*, satisfies equation (5) and ensures that the payoffs from false advertising and truth-telling are equal. If instead the low type lied with a probability above (below) y^*, buyers' expected quality q_H^e would be too low (high) and hence the low type would have a strict preference to report truthfully (falsely), yielding a contradiction. Hence randomization with probability y^* is an essential feature of this equilibrium.

Notice that false advertising exists in both the pooling and the semi-pooling equilibria. However in a pooling equilibrium, buyers effectively ignore high claims and simply maintain their prior belief. On the other hand, semi-pooling equilibria have the more appealing feature that high claims do provide buyers with some useful information, and thus induce them to positively update their belief, with $q_H^e > \bar{q}$.

4 The Effects of Policy

First consider the effects of policy on the level of false advertising, y^*. By using equations (5) and (6) it follows that:

Lemma 2. *The level of false advertising y^*, is continuous and weakly decreasing in the level of penalty ϕ.*

This feature of the model is useful analytically in the subsquent sections, and ensures that stronger policy smoothly increases the informativeness of advertising. When $\phi > \phi_0$ or $\phi < \phi_1$, a low quality firm has a strict preference for truth-telling or lying respectively, and so small changes in ϕ have no effect. However when $\phi \in [\phi_1, \phi_0]$, the probability of false advertising y^* satisfies the indifference condition (5) and is strictly decreasing in ϕ from 1 to 0. Intuitively, to maintain indifference of the low type as ϕ increases, high reports must become more credible. Since buyers are Bayesian, this is only possible if y^* is strictly lower.

4.1 Buyer Surplus

We now consider the effects of policy on a variety of welfare measures, starting with buyer surplus. Using Proposition 1 we can write expected buyer surplus as

$$E(v) = x(1-y^*)\int_{p^*(L)-L}^{b} (L+\varepsilon - p^*(L))\,\mathrm{d}G(\varepsilon) + xy^*\int_{p^*(q_H^e)-q_H^e}^{b} (L+\varepsilon - p^*(q_H^e))\,\mathrm{d}G(\varepsilon)$$
$$+(1-x)\int_{p^*(q_H^e)-q_H^e}^{b} (H+\varepsilon - p^*(q_H^e))\,\mathrm{d}G(\varepsilon). \tag{7}$$

In words, with probability $x(1-y^*)$ the firm sends a low report and charges $p^*(L)$. Buyers correctly infer low quality, buy if $\varepsilon \geq p^*(L)-L$, and receive $L+\varepsilon-p^*(L)$. Then with probability $1-x+xy^*$ the firm sends a high report and charges $p^*(q_H^e)$. Buyers update their beliefs according to equation (6), and buy if $\varepsilon \geq p^*(q_H^e) - q_H^e$. With conditional probability $xy^*/\left(1-x+xy^*\right)$, the product is low quality, and buyers receive $L+\varepsilon-p^*(q_H^e)$. With conditional probability $(1-x)/\left(1-x+xy^*\right)$, the product is high quality, and buyers receive $H+\varepsilon-p^*(q_H^e)$. After collecting terms and using the definition of $v^*(q)$ in equation (4), the above expression simplifies as follows, where $E(v)$ is just a convex combination of $v^*(L)$ and $v^*(q_H^e)$.

$$E(v) = x(1-y^*)v^*(L) + (xy^* + 1 - x)v^*(q_H^e). \tag{8}$$

We now exploit the smooth feature of our equilibrium to investigate the effect of a marginal increase in the penalty ϕ. In particular Lemma 3 will provide conditions under which an increase in the penalty ϕ leads to a *reduction* in expected buyer surplus. To understand why this can happen, recall that the level of false advertising y^* is a decreasing function of ϕ, and note that an increase in y^* produces two effects. On the one hand, buyers are more likely to receive a false advert and so be persuaded to buy a low quality product at an inflated price $p^*(q_H^e) > p^*(L)$. However on the other hand, the increase in lying damages the credibility of high quality claims, and forces any type making such a claim to reduce its price. In more detail, using equation (8) one can write

$$\frac{\partial E(v)}{\partial y^*} = \underbrace{x\left[v^*\left(q_H^e\right) - \left(q_H^e - L\right) D\left(p^*(q_H^e) - q_H^e\right) - v^*(L)\right]}_{\text{'Persuasion' effect}} \\ \underbrace{- (1 - x + xy^*) D(p^*(q_H^e) - q_H^e) \frac{\partial p^*(q_H^e)}{\partial q_H^e} \frac{\partial q_H^e}{\partial y^*}}_{\text{'Price' effect}}. \quad (9)$$

The first term is a 'persuasion' effect. Conditional on the firm having low quality (which occurs with probability x), a marginal increase in lying replaces the surplus that the buyer would have received if the firm had told the truth, $v^*(L)$, with the surplus associated with false advertising, $v^*\left(q_H^e\right) - \left(q_H^e - L\right) D\left(p^*(q_H^e) - q_H^e\right)$. To explain this latter surplus, note that after observing a high report, buyers update their beliefs to q_H^e, and expect to receive a surplus $v^*\left(q_H^e\right)$. However since quality is low, each of the $D\left(p^*(q_H^e) - q_H^e\right)$ units bought is worth $q_H^e - L$ less than anticipated. This harms buyers by prompting them to pay too much and to potentially buy too many units of a low quality product. The second term in (9) is a 'price' effect. A marginal increase in lying decreases the probability that a high claim is true, and so causes rational buyers to revise their belief q_H^e downwards. This lowers the level of monopoly power and induces any type with a high report to reduce its price by $-\left(\partial p^*(q_H^e)/\partial q_H^e\right) \times \left(\partial q_H^e/\partial y^*\right)$. Hence conditional on the firm sending a high report (which occurs with probability $1 - x + xy^*$), buyer surplus is strictly higher on each of the $D(p^*(q_H^e) - q_H^e)$ inframarginal units bought. Importantly for our later results, this effect is more powerful when quality pass-through $\partial p^*(q_H^e)/\partial q_H^e$ is larger.

In order to determine which of the persuasion and price effects dominates, we impose the following regularity condition on demand curvature:

Condition 1. Let $z_v(\psi) = -\sigma'(\psi) + [2 - \sigma(\psi)]g(\psi)/[1 - G(\psi)]$. The demand function satisfies either i) $\tilde{q} < \infty$ and $z_v(\psi) > 0$ for all $\psi \in (a, b)$, or ii) $\underset{\sim}{q} = -\infty$, $\tilde{q} = \infty$, $z_v(\psi)$ changes from negative to positive at exactly one value of $\psi \in (a, b)$, and $\lim_{\psi \to a} \sigma(\psi) = -\infty$.

Condition 1 is satisfied by a wide class of demand functions, and ensures that $v^*(q)$ is s-shaped in quality. For convenience, we denote

$$\hat{q}_v = \sup\left\{q \in (\underset{\sim}{q}, \tilde{q}) \,:\, z_v(p^*(q) - q) > 0\right\} \quad (10)$$

as the finite quality level at which $v^*(q)$ changes from being strictly convex to concave. Condition 1i is satisfied by a rich class of demands that exhibit constant curvature, which includes linear and exponential demand (see Bulow and Pfleiderer 1983). Here $\hat{q}_v = \tilde{q}$ such that $v^*(q)$ is strictly convex for all $q \in (\underset{\sim}{q}, \tilde{q})$, but independent of quality and equal to $\int_a^b [1 - G(z)]\, dz$ when

$q > \tilde{q}$. Alternatively, Condition 1ii is satisfied by many demands with increasing curvature - including those derived from the Normal, Logistic, Type I Extreme Value and Weibull distributions. For these demands $\hat{q}_v$ solves $z_v\left(p^*(q) - q\right) = 0$, and $v^*(q)$ is strictly convex for $q < \hat{q}_v$ but strictly concave for $q > \hat{q}_v$. Further details are provided in Section A of the Supplementary Appendix.[11] We can then state:

Lemma 3. *Consider $\phi \in [\phi_1, \phi_0]$ and suppose that Condition 1 holds.*
i) If $L < \hat{q}_v$ expected buyer surplus is quasiconcave in ϕ. In particular there exists a threshold $q^*(L) \geq \hat{q}_v$ (which is weakly decreasing in L and satisfies $\lim_{L \to \hat{q}_v} q^*(L) = \hat{q}_v$) such that expected buyer surplus is strictly increasing in ϕ if $q_H^e < q^*(L)$, but strictly decreasing in ϕ if $q_H^e > q^*(L)$.
ii) If $L > \hat{q}_v$ expected buyer surplus is weakly decreasing in ϕ.

Lemma 3 is interpreted as follows. First consider $L < \hat{q}_v$. When $q_H^e < q^*(L)$ buyers are relatively pessimistic upon seeing a high report, quality pass-through is relatively weak, and so the persuasion effect dominates. When instead $q_H^e > q^*(L)$ buyers are relatively optimistic upon seeing a high report, quality pass-through is relatively strong, and so the price effect dominates. Hence a stronger policy benefits buyers in the former situation, but harms them in the latter. Second consider $L > \hat{q}_v$. If demand satisfies Condition 1i, the price and persuasion effects cancel such that $\partial E(v)/\partial \phi = 0$. Intuitively the market is fully covered irrespective of the firm's claim, and buyers pay either $a + L$ following a low claim, or $a + q_H^e$ following a high claim. The average price paid is therefore $a + \bar{q}$ which is independent of ϕ. However if instead demand satisfies Condition 1ii, the price effect strictly dominates such that $\partial E(v)/\partial \phi < 0$.

We now consider the optimal level of penalty, ϕ^*. To ease exposition, we henceforth focus on the (more interesting) case where $L < \hat{q}_v$. Recalling Lemma 3, we find that:

Proposition 2. *Fix $L < \hat{q}_v$ and suppose that Condition 1 holds. The buyer-optimal penalty, ϕ^*, is characterized as follows:*
i) When $H \leq q^*(L)$, $\phi^* \geq \phi_0$ such that $y^* = 0$.
ii) When $\bar{q} < q^*(L) < H$, $\phi^* = \pi^*(q^*(L)) - \pi^*(L)$ such that $y^* = \frac{(H - q^*(L))(1-x)}{(H-q^*(L))(1-x) + q^*(L) - \bar{q}} \in (0,1)$.
iii) When $q^*(L) \leq \bar{q}$, $\phi^* \leq \phi_1$ such that $y^* = 1$.

Proposition 2 provides a range of demand and parameter conditions where a buyer-oriented policymaker refrains from eradicating false advertising. Recall from Lemma 3 that for $L < \hat{q}_v$, a marginal decrease in false advertising increases buyer surplus if and only if buyers are relatively pessimistic about high claims, with $q_H^e < q^*(L)$. Therefore when $H \leq q^*(L)$ buyer surplus is globally decreasing in y^* and the policymaker optimally eliminates false advertising. However

[11] This appendix may also prove useful for other wider literatures. For instance, a recent literature on the welfare effects of third-degree price discrimination uses a restriction related to Condition 1 which ensures that buyer surplus is convex with respect to marginal cost (e.g. Chen and Schwartz 2015 and Cowan 2012).

when $\bar{q} < q^*(L) < H$ buyer surplus is quasiconcave and maximized at some $y^* \in (0,1)$, such that the optimal penalty tolerates some false advertising. Finally when $q^*(L) \leq \bar{q}$ buyer surplus is globally increasing in y^* and so the policymaker fully permits false advertising.

The fact that a positive level of false advertising can generate a higher buyer surplus than under full information (where $y^* = 0$) gives several policy implications. First, any instinctive per se implementation of strong penalties or blanket prohibitions on false advertising may actually limit buyer surplus. Second, the optimal use of advertising penalties is superior to an outright ban on low quality products. Such a ban only generates a surplus $E(v) = (1-x)v^*(H)$, which is weakly less than the surplus under full information.

Finally, we further detail the conditions under which positive false advertising is optimal.

Corollary 1. *Given Condition 1 and $L < \hat{q}_v$, the buyer-optimal level of false advertising is increasing in L, H, and $(1-x)$.*

When product quality levels are higher, or when the probability of a high type is larger, policy should allow a higher level of false advertising, y^*. Intuitively, when the monopolist's product quality technology is relatively 'healthy', the expected quality from a high claim, q_H^e, is relatively high such that the price effect becomes relatively more powerful. On the contrary, when the product quality technology is less 'healthy', the persuasion effect becomes especially harmful.

4.2 Profits

We now examine the effect of policy on profits. To begin, consider each individual firm type:

$$E(\pi_L) = \begin{cases} \pi^*(\bar{q}) - \phi & \text{if } \phi < \phi_1 \\ \pi^*(L) & \text{if } \phi \geq \phi_1 \end{cases} \quad \text{and} \quad E(\pi_H) = \begin{cases} \pi^*(\bar{q}) & \text{if } \phi < \phi_1 \\ \pi^*(L) + \phi & \text{if } \phi \in [\phi_1, \phi_0] \\ \pi^*(H) & \text{if } \phi > \phi_0 \end{cases} \tag{11}$$

This is explained as follows. When $\phi < \phi_1$ the equilibrium has full pooling such that each type earns $\pi^*(\bar{q})$, but the low type also incurs a penalty ϕ. When $\phi \in [\phi_1, \phi_0]$ the low type is indifferent between lying and truth-telling, and so earns $\pi^*(L)$. The high type, meanwhile, earns $\pi^*(q_H^e)$ which is equal to $\pi^*(L) + \phi$ from (5). Finally when $\phi > \phi_0$ the equilibrium has full separation, and so each type earns its full information payoff.

Remark 1. An increase in ϕ reduces $E(\pi_L)$, but increases $E(\pi_H)$.

Intuitively, stronger regulation increases the high type's payoff because it leads buyers to update more optimistically upon seeing a high claim. However tougher regulation hurts a low type because it becomes costlier to mimic a high type. Now consider expected equilibrium profit, $E(\Pi) = xE(\pi_L) + (1-x)E(\pi_H)$:

Proposition 3. *Expected profit is quasiconvex in ϕ and minimized at $\phi = \phi_1$. In addition:*
i) If $L < \tilde{q}$, expected profit is maximized by $\phi^* \geq \phi_0$.
ii) If $L \geq \tilde{q}$, expected profit is maximized by either $\phi^* = 0$ or $\phi^* \geq \phi_0$.

A small increase in regulation can either benefit or harm the monopolist, depending upon how existing regulation ϕ compares with ϕ_1. In addition, it is straightforward to see from (11) that $\phi \in (0, \phi_0)$ is strictly dominated under an expected profit objective. This implies that the penalty should never be paid in equilibrium. Then, given the convexity of $\pi^*(q)$, full separation with $\phi^* \geq \phi_0$ is always weakly optimal. Intuitively, strong regulation allows the firm to extract buyer surplus more effectively when it has high quality. Hence, if the monopolist could credibly commit to effective self-regulation (perhaps through a third party), Proposition 3 implies that it would weakly prefer to avoid using false advertising. In some circumstances, such as when $L < \hat{q}_v$ and $H < q^*(L)$, such self-regulation might be acceptable to buyers because the monopolist's preferred level of penalty coincides with that of the buyers. This may offer some support for Europe's industry-led regulation. However, in other circumstances self-regulation would go against buyers' preferences e.g. when $L < \hat{q}_v$ and $H > q^*(L)$. Here, contrary to any concerns that self-regulation may be too lax, the monopolist's preferred level of penalty is strictly *higher* than buyers'.

4.3 Total Welfare

We now consider total welfare. Suppose that the penalty, ϕ, is in the form of a fine which is as valuable to the policymaker as it is to the firm. Using Proposition 1, we can write expected total welfare as follows.

$$E(w) = x(1-y^*)\left[v^*(L) + \pi^*(L)\right] + (1 - x + xy^*)\left[v^*(q_H^e) + \pi^*(q_H^e)\right]. \qquad (12)$$

Notice that this expression is not just the summation of expected buyer surplus in (8), and weighted firm-type profits in (11), because by assumption the penalty has social value.

We will shortly provide conditions under which a stronger penalty ϕ reduces expected total welfare. We start with some intuition and then formally state the result. The firm uses its monopoly power to restrict output below the socially efficient level. A marginal increase in the level of false advertising y^* then changes this output distortion in two ways. First, it lowers the credibility of any high claim, and so forces any type with such a claim to further reduce its output below the socially optimal level. Second however, it also induces buyers to over-estimate a low type's quality, thereby causing the low type to increase its output. Under certain circumstances this latter output expansion can raise welfare and dominate the former effect[12]. In more detail:

$$\frac{\partial E(w)}{\partial y^*} = \underbrace{x\Big[v^*\left(q_H^e\right) - v^*\left(L\right) - (q_H^e - L)D(p^*(q_H^e) - q_H^e) + \pi^*(q_H^e) - \pi^*\left(L\right)\Big]}_{\text{Output expansion by a firm with } q=L}$$
$$+ \underbrace{(1 - x + xy)D(p^*(q_H^e) - q_H^e)\left(1 - \frac{\partial p^*(q_H^e)}{\partial q_H^e}\right)\frac{\partial q_H^e}{\partial y^*}}_{\text{Output contraction by a firm with } r=H} \quad (13)$$

The first term in (13) represents the change in welfare when a low type moves from reporting $r = L$ and generating a total surplus of $v^*\left(L\right) + \pi^*\left(L\right)$, to claiming $r = H$ and generating a surplus of $v^*\left(q_H^e\right) - \left(q_H^e - L\right) D(p^*(q_H^e) - q_H^e) + \pi^*\left(q_H^e\right)$. This term is positive if and only if L is above a certain threshold. Intuitively, the low type's socially optimal output level $D(-L)$ is increasing in L. Moreover when a low type engages in false advertising, its output increases from $D\left(p^*(L) - L\right) \leq D(-L)$ to $D(p^*(q_H^e) - q_H^e)$. Therefore if L is relatively small, this 'output expansion effect' goes far beyond the efficient level and so is bad for welfare. However if L is relatively large, the output expansion effect brings the low type closer to the efficient level, and so is good for welfare. The second term in (13) represents the change in surplus generated by a firm that claims to have high quality, following a small increase in y^*. As explained above, this is unambiguously negative because an increase in y^* reduces the credibility (and hence output) of a firm that reports $r = H$. Ceteris paribus, this 'output contraction effect' is smaller when quality pass-through $\partial p^*(q_H^e)/\partial q_H^e$ is larger since in that case the firm's output is less sensitive to buyers' belief about its quality.

To determine which of these two output effects dominates, we impose a regularity condition that differs slightly to the one used earlier.

Condition 2. Let $z_w(\psi) = -\sigma'(\psi) + [2-\sigma(\psi)][3-\sigma(\psi)]g(\psi)/[1-G(\psi)]$. The demand function satisfies either i) $\tilde{q} < \infty$ and $z_w(\psi) > 0$ for all $\psi \in (a,b)$, or ii) $\underset{\sim}{q} = -\infty$, $\tilde{q} = \infty$, $z_w(\psi)$ changes from negative to positive at exactly one value of $\psi \in (a,b)$, and $\lim_{\psi \to a} \sigma(\psi) = -\infty$.

Condition 2 ensures that $w^*(q) \equiv v^*(q) + \pi^*(q)$ is s-shaped in quality, where

$$\hat{q}_w = \sup\Big\{q \in (\underset{\sim}{q}, \tilde{q}) \;:\; z_w(p^*(q) - q) > 0\Big\} \quad (14)$$

denotes the critical quality level at which $w^*(q)$ changes from being strictly convex to concave.

[12]The output distortions on the two types' output are reminiscent of the output effects in third-degree price discrimination (e.g. Aguirre *et al* 2010).

The condition is again satisfied by a wide range of commonly-used distribution functions.[13] We can then state:

Lemma 4. *Consider $\phi \in [\phi_1, \phi_0]$ and suppose that Condition 2 holds.*
i) When $q_H^e < \hat{q}_w$ expected total welfare is strictly increasing in ϕ.
ii) When $L < \hat{q}_w < q_H^e$ expected total welfare is quasiconcave in ϕ. In particular there exists a threshold $L^*(q_H^e) < \hat{q}_w$ (which is weakly decreasing in q_H^e) such that expected total welfare is strictly increasing in ϕ if $L < L^*(q_H^e)$, but strictly decreasing in ϕ if $L > L^*(q_H^e)$.
iii) When $L > \hat{q}_w$ expected total welfare is weakly decreasing in ϕ.

Lemma 4 can be understood as follows. When $q_H^e < \hat{q}_w$ quality pass-through is relatively small, such that the output contraction effect dominates, and so $E(w)$ decreases in the level of false advertising y^*. When $L < \hat{q}_w < q_H^e$ quality pass-through is relatively stronger, and so the output contraction effect is weaker. A small increase in y^* therefore raises welfare provided L is sufficiently large, such that the expansion in the low type's output is not (too) excessive. Finally when L is large with $L > \hat{q}_w$, an increase in the penalty can never raise welfare as the output expansion always weakly dominates.

In a model with one firm type and naive buyers, Glaeser and Ujhelyi (2010) show that some false advertising always improves total welfare by increasing output towards the social optimum. Our welfare result with two types and rational buyers is more complex for the following reasons. First, since in our model the low type must pool with the high type, the increase in its output, $D(p^*(q_H^e) - q_H^e) - D(p^*(L) - L)$, may actually go beyond the social optimum and reduce welfare. Second, since in our model false advertising reduces the credibility of high claims, there is another welfare-reducing 'output contraction effect' which is not present in their paper.

Now consider the implications for the optimal penalty. To ease exposition, we focus on the (more interesting) case where $L < \hat{q}_w$. First, note that the policymaker will always eliminate false advertising with $\phi^* \geq \phi_0$ when $H < \hat{q}_w$. This follows from Lemma 4 because $q_H^e < \hat{q}_w$ for all $y^* \in [0, 1]$. For the remaining cases:

Proposition 4. *Fix $L < \hat{q}_w < H$ and suppose that Condition 2 holds. The welfare-optimal penalty, ϕ^*, is characterized as follows:*
i) When $L \leq L^*(H)$, $\phi^* \geq \phi_0$ such that $y^* = 0$.
ii) When $L \in (L^*(H), L^*(\bar{q}))$, ϕ^* induces $y^* \in (0, 1)$ such that $q_H^e = q^{**}$ where $L = L^*(q^{**})$.
iii) When $L \in [L^*(\bar{q}), \hat{q}_w)$, $\phi^* \leq \phi_1$ such that $y^* = 1$.

[13] Condition 2i holds for all demands with constant curvature, where $\hat{q}_w = \tilde{q} < \infty$ such that $w^*(q)$ is strictly convex for $q \in (\underset{\sim}{q}, \hat{q}_w)$ and linear for $q \geq \tilde{q}_w$. Condition 2ii is satisfied by some demands with increasing curvature, including the Normal, Weibull, and Type I Extreme Value, where $\hat{q}_w$ solves $z_w\,(p^*(q) - q) = 0$ such that $w^*(q)$ is strictly convex for $q < \hat{q}_w$ but strictly concave for $q > \hat{q}_w$. See Section A of the Supplementary Appendix.

Proposition 4 provides parameter conditions to show when a positive level of false advertising is welfare-optimal. Even though full separation with $\phi^* \geq \phi_0$ involves full information prices and no output distortions, it may still be optimal to use false advertising to induce some form of pooling. In line with intuition, one can show that the optimal level of false advertising is weakly lower than that under a buyer surplus objective. However, the optimal level of false advertising remains increasing in the 'healthiness' of the market (e.g. L, H, and $(1-x)$).[14]

5 Endogenous Quality Investment

We now extend the main model to examine some additional effects of false advertising in a market with endogenous product quality. Such effects have been largely ignored within the literature. However they are important to consider because the existence of false advertising may reduce the incentives to invest in product quality by limiting the credibility of advertising. Suppose that the firm is initially endowed with low quality L, but can upgrade to high quality H by paying an investment cost C. This cost is drawn privately from a distribution $F(C)$ on $(0,\infty)$, with corresponding density $f(C) > 0$. The move order is then as follows. At stage 1 the policymaker commits to a penalty ϕ. At stage 2 the firm learns its investment cost C, and privately chooses whether to upgrade. It also announces its report and price. The game then proceeds as in the main model, with buyers making their purchase decisions, and the policymaker instigating any potential penalties. Let $x^*(\phi)$ denote the endogenous probability that the firm has low quality.

There always exists a trivial equilibrium in which $x^*(\phi) = 1$. If buyers believe that product quality is low for all reports and prices, the firm has no incentive to invest. However, in general, there also exist other alternative PBE. Henceforth, we restrict attention to PBE where, as before, buyer beliefs only depend upon the firm's claim. Moreover whenever possible, we select an equilibrium where the firm invests with positive probability.

Lemma 5. *i) When $\phi = 0$ all equilibria have $x^*(\phi) = 1$. ii) When $\phi \in (0, \phi_0]$ there is a unique equilibrium (up to off-path beliefs) satisfying our restrictions, with $x^* = 1 - F(\phi) \in (0,1)$ and $r(H) = H$.*

Intuitively, an increase in ϕ induces investment by widening the gap in profits earned by high and low quality firms. In more detail, when $\phi = 0$ buyers cannot distinguish between high and low quality. The firm earns the same profit regardless and therefore chooses not to

[14]Finally, we note that false advertising can remain optimal even when a fraction τ of the penalty is 'lost' and does not contribute to total welfare. For instance when $L^*(\bar{q}) < L < \hat{q}_w \leq \bar{q}$ the optimal penalty induces $y^* = 1$ for all $\tau \in [0,1]$. The only difference is that when $\tau = 0$ any $\phi \in [0, \phi_1]$ maximizes total welfare, whereas when $\tau > 0$ it is more attractive to reduce the penalties incurred, so $\phi^* = 0$ is the unique optimum.

invest. Alternatively when $\phi \geq \phi_0$, claims are fully credible. A low quality firm reports $r = L$ and earns $\pi^*(L)$, whilst a high quality firm reports $r = H$ and earns $\pi^*(H)$. Since the gains from investing are $\pi^*(H) - \pi^*(L) \equiv \phi_0$, the firm upgrades if and only if $C \leq \phi_0$. Finally when $\phi \in (0, \phi_0)$, the level of false advertising is necessarily positive for the same reason as in the main model. This further implies that a high quality product earns ϕ more than a low quality product such that the firm invests with probability $F(\phi)$. However unlike the main model, the probability of false advertising y^* is not necessarily decreasing everywhere in ϕ. Recall the definition $q_H^e \equiv E(q|r = H)$. Intuitively, an increase in ϕ can enhance advertising credibility and cause investment to increase by so much that, ceteris paribus, the net gains from false advertising, $\pi^*(q_H^e) - \phi$, actually rise, and prompt a higher y^*. In his seminal discussion, Nelson (1974) suggested that advertising policy may increase the credibility of false advertising. Here, we formalize an even stronger relationship - policy can provide so much credibility that parameters exist where the probability of false advertising is increasing in the level of penalty. Nevertheless, despite any potential increase in y^*, stronger penalties still always induce a larger expected quality, q_H^e. Now consider the optimal penalty:

Proposition 5. *Suppose Condition 1 holds and that $L < \hat{q}_v$. A buyer-orientated policymaker i) always sets $\phi > 0$, and ii) sets $\phi < \phi_0$ such that $y^* > 0$ provided $H > q^*(L)$ and $f(\phi_0)/F(\phi_0)$ is sufficiently small.*

To understand this result, rewrite (8) from earlier using $x \equiv x^*(\phi)$ as

$$E(v) = v^*(L) + \underbrace{(H - L)\frac{v^*(q_H^e) - v^*(L)}{q_H^e - L}}_{\text{Price/Persuasion terms}} \times \underbrace{(1 - x^*(\phi))}_{\text{Investment term}} . \tag{15}$$

The second term captures the tradeoff between the price and persuasion effects. As in the main model, Condition 1 ensures that this term is increasing in ϕ if and only if $q_H^e < q^*(L)$. The third term relates to a new 'investment effect'. A high quality product generates more buyer surplus than a low quality product. Therefore ceteris paribus, an increase in ϕ is beneficial since it prompts a higher level of investment. Proposition 5 is then explained as follows. Firstly, unlike in the main model, $\phi = 0$ is never optimal because the firm then never invests and so buyers get only $v^*(L)$. Alternatively, for any $\phi > 0$, the firm invests with positive probability and so from (15) buyer surplus strictly exceeds $v^*(L)$. Secondly though, despite this new investment effect, policy may still refrain from completely eliminating false advertising. In particular, this is the case when $H > q^*(L)$ and $f(\phi_0)/F(\phi_0)$ is relatively small. Intuitively the latter restriction on $f(C)$ implies that starting from strong regulation, $\phi = \phi_0$, a small decrease in ϕ only has a small effect on the investment probability, such that the combined price and persuasion effects dominate. Consequently, as in the main model, false advertising can sometimes benefit buyers.[15]

Finally, one can compare the optimal penalty with that under exogenous quality. In particular, let ϕ^*_{en} denote the optimal penalty with endogenous quality, and impose a technical condition $f(\phi)(H-L) < 1$ to ensure it is unique. Then, to make a comparison, let $x^*(\phi^*_{en})$ be the proportion of low types under exogenous quality, and denote the associated optimal penalty by ϕ^*_{ex}. One can then prove that $\phi^*_{ex} \leq \phi^*_{en}$, such that the optimal penalty is stronger when quality is endogenous due to the existence of the investment effect.

6 Robustness

This final section shows how the results of the main model are robust to i) an arbitrary number of quality types, ii) asymmetric costs, and iii) competition.

6.1 An Arbitrary Number of Types

Suppose there are now $n > 2$ quality levels, denoted by $q_1 < ... < q_n$, and that the firm has quality q_i with probability $x_i \in (0,1)$. To simplify the exposition, let $q_2 > \underset{\sim}{q}$ and $q_{n-1} < \hat{q}_v$ (relaxing these assumptions is straightforward, but adds no new insights). Marginal cost is the same for all types and normalized to zero, while ex ante expected quality is again denoted by $\bar{q} = \sum x_i q_i$. The firm may send any report from the set $Q = \{q_1, ..., q_n\}$, and the policymaker can commit to a richer penalty $\phi(q,r) \geq 0$, which depends on both the firm's actual and reported qualities. We assume that the firm can only be fined if it over-reports its quality i.e. $\phi(q,r) = 0$ for all $r \leq q$. The game and move order are otherwise unchanged.

As usual, for any particular penalty $\phi(q,r)$ there may exist a large number of PBE. Therefore, for reasons analogous to the main model, we continue to restrict attention to PBE in which i) $r(q) \geq q \; \forall q$ such that no type under-reports its quality, and ii) buyer beliefs depend on the firm's claim but not its price. Notice that in any PBE satisfying these restrictions, the penalty function $\phi(q,r)$ induces a mapping from quality types into reports. It is then convenient to let $y^*_{i,j}$ be the probability that a firm of type i claims to have quality j; hence $y^*_{i,i}$ denotes the probability that firm type i sends a truthful report. Letting $\mathbf{y}^*$ be the (triangular) matrix of such probabilities, we may then state:

[15] Similarly, one can show that a welfare-maximizing policymaker may also refrain from completely eradicating false advertising.

Lemma 6. *The optimal penalty can be derived in two steps:*
i) First, choose the matrix of probabilities $\mathbf{y}^*$ *which maximizes the policymaker's objective.*
ii) Second, there exists a penalty function $\phi(q,r)$ *which induces the policymaker's optimal* $\mathbf{y}^*$ *as the unique equilibrium outcome of the game.*

Thus conceptually the problem is similar to the two-type case. In particular, analogous to Lemma 2, we can work with the matrix of report probabilities $\mathbf{y}^*$, and be sure that at least one penalty function can implement the desired $\mathbf{y}^*$. Now consider optimal penalties. Given our main model, it is not surprising that under certain conditions the policymaker will permit some false advertising. However once there is an arbitrary number of types, the policymaker also has to decide *which* quality types will be allowed to engage in false advertising, and which quality level(s) they will mimic. To simplify the exposition we now focus on distributions satisfying Conditions 1i and 2i for which $\hat{q}_v = \hat{q}_w = \tilde{q}$[16]:

Proposition 6. *Suppose the match distribution satisfies Conditions 1i and 2i. The optimal report probabilities are as follows:*
i) Buyer surplus. (a) When $q_n \leq \tilde{q}$ *it is maximized by* $y^*_{i,i} = 1$ *for all* i*. (b) When* $q_n > \tilde{q} > \bar{q}$ *there exists a critical type* i^* *satisfying* $E(q|q \geq q_{i^*}) \leq \tilde{q} < E(q|q \geq q_{i^*+1})$*, such that the optimal solution has* $y^*_{i,i} = 1$ *for all* $i < i^*$*,* $y^*_{i,n} = 1$ *for all* $i > i^*$*, and* $y^*_{i^*,i^*} = 1 - y^*_{i^*,n}$ *where* $y^*_{i^*,n}$ *satisfies:*

$$\frac{x_{i^*} y^*_{i^*,n} q_{i^*} + \sum_{i=i^*+1}^{n} x_i q_i}{x_{i^*} y^*_{i^*,n} + \sum_{i=i^*+1}^{n} x_i} = \tilde{q}.$$

(c) When $\bar{q} \geq \tilde{q}$ *it is maximized by* $y^*_{i,n} = 1$ *for all* i*.*
ii) Profit is maximized by $y^*_{i,i} = 1$ *for all* i*.*
iii) Total welfare. There exists a threshold L^* *such that: (a) When* $q_n \leq \tilde{q}$ *it is maximized by* $y^*_{i,i} = 1$ *for all* i*. (b) When* $q_n > \tilde{q}$ *and* $q_{i^*} \geq L^*$ *it is maximized by the buyer-optimal matrix. (c) When* $q_n > \tilde{q}$ *and* $q_{i^*} < L^*$*, it is maximized by* $y^*_{i,i} = 1$ *for all* i *with* $q_i < L^*$*, and* $y^*_{i,n} = 1$ *for all* i *with* $q_i \geq L^*$*.*

The policymaker induces each firm type to either report truthfully or to claim to have the highest possible quality q_n. Similar to the two-type model, in many cases one firm type is required to randomize over its report. Whether buyers gain from false advertising depends upon how the highest quality type q_n compares with $\tilde{q}$. If $q_n \leq \tilde{q}$ the persuasion effect dominates, such that buyers are better off if the firm truthfully reveals its quality. However if $q_n > \tilde{q}$, the highest type has a lot of market power, and so lower types are pooled with it to generate a beneficial price effect. In order to minimize the negative persuasion effect, this pooling is done

[16] The optimal pattern of false advertising is qualitatively the same for distributions satisfying the alternative Conditions 1ii and 2ii. Further details are available on request.

from the top i.e. first the q_{n-1} type is pooled, then the q_{n-2} type, and so forth, until either $E(q|r = q_n) = \tilde{q}$ or no more types are left to pool. Hence the optimum has full pooling when $\bar{q} > \tilde{q}$, and semi-pooling when $\bar{q} < \tilde{q} < H$. In the latter case, the policymaker permits 'small' lies by types close to q_n, whilst forbidding 'large' lies by types at or close to q_1.

Policy under a total welfare objective also depends upon whether $q_n \gtrless \tilde{q}$. When $q_n \leq \tilde{q}$ a welfare-maximizing policy involves truthful advertising and so coincides with what is optimal for both buyers and the firm. However when $q_n > \tilde{q}$ a welfare-oriented policymaker may allow some lower types to use false advertising in order to raise their output. As with buyer surplus, types with quality closer to q_n are more likely to be allowed to use false advertising since their socially-optimal output levels are highest. Overall, the main insights from the two-type model carry over into this richer multi-type environment.

6.2 Asymmetric Costs

Returning to the two-type case, we now permit the types to differ in marginal costs. In particular, suppose that a product of quality q now has constant marginal cost $c(q)$ with $c'(q) \in (0, 1)$ and $c''(q) = 0$. Let $\pi(p, q^e; i) = (p - c(i))[1 - G(p - q^e)]$ be the profit earned with price p, expected quality q^e, and actual quality $i \in \{L, H\}$, and denote $p^*(q^e; i) = \arg\max_p \pi(p, q^e; i)$ and $\pi^*(i) = \pi(p^*(i; i), i; i)$.

Common refinements now make sharper predictions, but as we discuss, these are unappealing in the context of our framework. First, refinements like D1 uniquely select a least-cost separating equilibrium. Here, the low type always reports truthfully and sets its full-information price $p^*(L; L)$. When $\pi^*(L) \geq \pi(p^*(H; H), H; L) - \phi$, the high type sends $r = H$ and charges its full-information price $p^*(H; H)$. When $\pi^*(L) < \pi(p^*(H; H), H; L) - \phi$, the high type sends $r = H$ and charges the highest price consistent with $\pi(p, H; L) - \phi = \pi^*(L)$ such that the low type just prefers to advertise truthfully. However, for our purposes, a major disadvantage of this equilibrium is that the low type always reports truthfully, and hence no false advertising occurs. Mailath *et al* (1993) also provide a number of wider arguments against the de facto use of refinements like D1 to select least-cost separating equilibria. Second, Mailath *et al*'s (1993) Undefeated Equilibrium refinement has relatively little bite. Nevertheless, a stronger version proposed by Mezzetti and Tsoulouhas (2000) selects either a full separating equilibrium or a full pooling equilibrium, depending on the penalty ϕ and the ex ante probability that quality is low. However, as discussed earlier, while a full pooling equilibrium does allow for false advertising, it does not allow for the potentially more realistic case where false claims inflate buyers' beliefs beyond their priors.

Therefore we now characterize an alternative semi-pooling equilibrium in a similar style to the main model. First, we only consider equilibria where the high type reports truthfully. Since it is costless for a high type to claim $r = H$, only 'perverse' buyer beliefs could force a high type to make a low claim. Secondly, if the low type reports $r = H$ in equilibrium, then it must charge the same price as the high type otherwise buyers would correctly infer its type. However, conditional on reporting $r = H$, the two types' payoffs now differ by more than just a constant, and so the argument we used in the main model must be modified. Thirdly then, we proceed as follows. Suppose that buyers expect the low type to report $r = H$ with probability $y^* \in (0, 1]$, and interpret any off-path price as a possible signal about the firm's type but not about y^*. We then restrict attention to PBE where, conditional on reporting $r = H$, the firm charges the high type's preferred price, $p^*(q_H^e; H)$. One justification for this is that since the high type is the one being mimicked, it should have some 'leadership' in choosing its preferred pooling price.[17] We may then state:

Lemma 7. *Suppose $c(H) - c(L)$ is not too large. There is a unique semi-pooling PBE (up to off-path beliefs) satisfying our restrictions, in which:*
i) A high type firm claims $r = H$ and charges $p^*(q_H^e; H)$.
ii) A low type firm randomizes. With probability y^* it claims $r = H$ and charges $p^*(q_H^e; H)$. With probability $1 - y^*$ it claims $r = L$ and charges $p^*(L; L)$.
- When $\phi \leq \phi_1' \equiv \pi^*(p^*(\bar{q}; H), \bar{q}; L) - \pi^*(L)$, $y^* = 1$.
- When $\phi \geq \phi_0' = \pi^*(p^*(H; H), H; L) - \pi^*(L)$, $y^* = 0$.
- When $\phi \in (\phi_1', \phi_0')$, $y^* \in (0, 1)$ and uniquely solves

$$\pi^*(p^*(q_H^e; H), q_H^e; L) - \phi = \pi^*(L). \tag{16}$$

iii) q_H^e is given by (6). Buyer beliefs are such that $\Pr(q = H | \{r, p\} = \{H, p^*(q_H^e; H)\}) = \frac{1-x}{1-x+xy^*}$ and $\Pr(q = H | \{r, p\} \neq \{H, p^*(q_H^e; H)\}) = 0$.

This equilibrium is qualitatively similar to the one derived earlier in the main model. In particular for intermediate values of the penalty ϕ false advertising arises in equilibrium, and yet high advertised claims are still partially informative and thus induce buyers to positively update their prior belief about quality. We also note that equilibrium play varies smoothly with both ϕ and $c(H)$, and converges to that of the main model as $c(H) \to c(L)$.

Finally, we now briefly comment on the implications for policy. Suppose that $c(H) - c(L)$ is not too large. Relative to symmetric costs, there is now an additional reason to eradicate

[17] Another justification is as follows. Suppose again that buyers interpret any off-path price as a possible signal about the firm's type but not about y^*. Then, provided $c(H) - c(L)$ is not too large, Mezzetti and Tsoulouhas's (2000) Strongly Undefeated Equilibrium refinement uniquely predicts that the firm should charge $p^*(q_H^e; H)$ when reporting $r = H$.

false advertising. Under asymmetric costs, a lying low type must distort its price upwards at $p^*\left(q_H^e; H\right)$ instead of its preferred (lower) price $p^*\left(q_H^e; L\right)$. This distortion provides a further loss to buyer surplus, total welfare, and ex ante profits. Nevertheless under Conditions 1 and 2, it remains true that for certain values of L, H, and x both a buyer- and a welfare-oriented policymaker would permit a strictly positive level of false advertising.

6.3 Competition

This final subsection introduces competition into the main model. Suppose an established incumbent, I, with quality q_I, competes against an entrant, E, with quality, q_E. Product differentiation is modeled using a Hotelling line such that a buyer with location $z \in [0, 1]$ can gain $U_I(z) = q_I - p_I - tz$ or $U_E(z) = q_E - p_E - t(1 - z)$ from trading with the respective firms. While the incumbent's product quality is known, the entrant's quality is private information. Specifically, the entrant's product quality equals L with probability $x \in (0, 1)$ and $H > L$ with probability $1-x$, such that the entrant's ex ante average quality level equals $\bar{q} = xL+(1-x)H$. Let all marginal costs be zero, and the buyers' outside option be sufficiently poor such that buyers always buy. The game then proceeds with i) the policymaker publicly selecting ϕ, ii) the entrant learning its quality and issuing a report $r \in \{L, H\}$, iii) the entrant and incumbent simultaneously selecting their prices, p_E and p_I, iv) buyers making their purchase decisions, and v) the policymaker administering any potential penalties.

To begin, consider a benchmark case where q_E is public information. The Nash equilibrium price charged by firm $i \in \{I, E\}$ is then

$$p_i^*(q_i, q_{-i}) = \begin{cases} 0 & \text{if } q_i \leq \underset{\sim}{q_i} \\ t + \left(\frac{q_i - q_{-i}}{3}\right) & \text{if } q_i \in \left(\underset{\sim}{q_i}, \tilde{q}_i\right) \\ q_i - q_{-i} - t & \text{if } q_i \geq \tilde{q}_i \end{cases} \tag{17}$$

where $\underset{\sim}{q_i} = q_{-i} - 3t$ and $\tilde{q}_i = q_{-i} + 3t$. Intuitively, when $q_i \leq \underset{\sim}{q_i}$ firm i is uncompetitive so its price is driven down to marginal cost. When instead $q_i \in (\underset{\sim}{q_i}, \tilde{q}_i)$, both firms are active. Here, an increase in q_i shifts out firm i's demand curve at the expense of its rival, prompting firm i to charge more and its rival to charge less. Finally, when $q_i \geq \tilde{q}_i$, firm i's product is so strong that it monopolizes the whole market; firm i then sets its price such that the marginal buyer is indifferent about buying from it. In addition, one can show that firm i's equilibrium profit, $\pi_i^*(q_i, q_{-i})$, is increasing in its own quality q_i, and decreasing in that of its rival q_{-i}.

Now let the entrant's product quality be private information. As consistent with the main model, we restrict attention to PBE where i) the high entrant type always issues a high report, and ii) buyer beliefs do not depend on price.

Lemma 8. *There exists a unique semi-pooling equilibrium (up to off-path beliefs) satisfying our restrictions, in which:*
i) The probability with which a low quality entrant reports $r = H$, y^*, is the same as in Proposition 1ii) after replacing $\pi^*(z)$ with $\pi^*_E(z, q_I)$.
ii) The firms charge $p^*_E(L, q_I)$ and $p^*_I(q_I, L)$ respectively if the entrant reports $r = L$; and $p^*_E(q^e_H, q_I)$ and $p^*_I(q_I, q^e_H)$ if the entrant reports $r = H$.
iii) Buyer beliefs are similar to Proposition 1iii) i.e. following a report $r = H$, buyers expect the entrant to have quality $q^e_H = \frac{xy^*L+(1-x)H}{1-x+xy^*}$.

There exists a semi-pooling equilibrium which is qualitatively the same as under monopoly where a low quality entrant randomizes between lying and reporting truthfully. Buyers update their beliefs about entrant quality accordingly, and conditional on those beliefs, the two firms charge Nash equilibrium prices.

Now consider the optimal penalty, starting with an industry profits objective. One can verify that both the entrant and the incumbent weakly prefer $y^* = 0$, with a strict preference whenever $L < \tilde{q}_E$. Hence, like the monopoly case, an industry self-regulator would choose to completely eliminate false advertising with the use of a tough policy $\phi^* \geq \phi_0 \equiv \pi^*_E(H, q_I) - \pi^*_E(L, q_I)$.

We now consider buyer surplus and total welfare. Here, in order to demonstrate that a policymaker may still permit a positive level of false advertising, it is sufficient to focus on the case where the entrant always has positive market share, with $L \geq \underset{\sim}{q_E}$. First, consider buyer surplus.

Proposition 7. *When $L \geq \underset{\sim}{q_E}$ the buyer-optimal level of false advertising y^* is the same as in Proposition 2 after replacing $\hat{q}_v$ and $q^*(L)$ with $\tilde{q}_E$.*

To understand this result, write

$$\frac{\partial E(v)}{\partial y^*} = x\left[v^*(q_I, q^e_H) - (q^e_H - L)D^*_E(q^e_H, q_I) - v^*(q_I, L)\right] \\ - (1 - x + xy^*)\frac{\partial q^e_H}{\partial y^*}\left[D^*_I(q_I, q^e_H)\frac{\partial p^*_I(q_I, q^e_H)}{\partial q^e_H} + D^*_E(q^e_H, q_I)\frac{\partial p^*_E(q^e_H, q_I)}{\partial q^e_H}\right], \quad (18)$$

where $D^*_I(q_I, q^e_H)$ and $D^*_E(q^e_H, q_I)$ are the respective equilibrium demands. The first term is a revised 'persuasion' effect which measures the change in buyer surplus generated by a low quality entrant when it changes its report from $r = L$ to $r = H$. As usual, a low type entrant uses false advertising to induce buyers to buy too many units, at an inflated price. However, false advertising now also allows the low type entrant to compete more effectively, which reduces the incumbent's price from $p^*_I(q_I, L)$ to $p^*_I(q_I, q^e_H)$. The second term is a revised 'price' effect. Conditional on the entrant using a high claim, an increase in lying reduces q^e_H and prompts the entrant to charge a (weakly) lower price, but allows the incumbent to select a (weakly) higher price. This net price effect need no longer benefit buyers - it is beneficial if and only if the entrant's market share exceeds 0.5, which is equivalent to $q_H^e \geq q_I$. However, in aggregate, the optimal penalty remains qualitatively similar to that under monopoly. In particular, false advertising remains optimal when $H > \tilde{q}_E$ in order to weaken the high type entrant's market power.

Now consider total welfare. From above, it is optimal to set $\phi^* \geq \phi_0$ to induce $y^* = 0$ when $H \leq \tilde{q}_E$ as this is preferred by all parties. For the remaining cases, we can state:

Proposition 8. *When $L \geq \underset{\sim}{q}_E$ the welfare-optimal y^* is the same as in Proposition 4 after substituting $\tilde{q}_E$ for $\hat{q}_w$, and $L_E^* = q_I + 3t/5$ for $L^*(H)$ and $L^*(\bar{q})$.*

This can be understood as follows. Firstly, an increase in y^* expands the output of a low type entrant. As with monopoly, this can either increase or decrease welfare depending on the level of L. Secondly, an increase in y^* reduces q_H^e, and therefore decreases the output of an entrant who reports $r = H$. However unlike monopoly, this second effect can actually increase welfare because, under competition, the firm with the highest (expected) quality uses its market power to restrict its output below the socially efficient level. Hence when $q_H^e \in (\underset{\sim}{q}_E, q_I)$ an entrant who reports $r = H$ actually overproduces, and so a small reduction in its output is socially beneficial. The proposition then shows that the aggregate of these two effects is qualitatively similar to monopoly. In particular, false advertising is used if and only if $H > \tilde{q}_E$ and L is relatively large with $L \geq L_E^*$ in order to raise the output of a low type entrant.

7 Conclusions

Despite its prevalence and policy importance, false advertising remains under-studied. This paper presents a model where false claims arise in equilibrium and actively influence the purchase decisions of rational buyers. By utilizing some results on demand curvature, the paper provides precise conditions under which buyers and society benefit from a positive level of false advertising due its effects in counteracting monopoly power. These results are extended in various directions, including cases where the firm faces competition or where product quality is endogenous.

Future work to further understand false advertising and advertising policy would be useful in a number of avenues. First, our paper focuses on markets where buyers are only able to assess a product's value with sufficient delay, such that policy plays a key role in regulating advertising. It would be interesting to extend our model to consider other contexts where regulation interacts with alternative sources of credibility, such as seller reputation. Second, our model considers a single market. More broadly, it would be useful to consider how a resource-constrained policymaker should regulate across multiple heterogeneous markets. Finally, much work remains in building on our analysis to study other types of false advertising, such as firms' misleading price claims.

References

[1] Aguirre I., S. Cowan and J. Vickers (2010) "Monopoly Price Discrimination and Demand Curvature" American Economic Review, 100, 1601-1615

[2] Anderson S. and R. Renault (2006) "Advertising Content" American Economic Review, 96, 93-113

[3] Armstrong M. and J. Zhou (2016) "Search Deterrence" Review of Economic Studies, 83, 26-57

[4] Bagwell K. (2007) "The Economic Analysis of Advertising" in the Handbook of Industrial Organization, Vol. 3, M. Armstrong and R. Porter (eds.), North-Holland: Amsterdam

[5] Barigozzi F., P.G. Garella, and M. Peitz (2009) "With a Little Help from My Enemy: Comparative Advertising as a Signal of Quality" Journal of Economics and Management Strategy, 18, 1071-1094

[6] Bulow J.I. and P. Pfleiderer (1983) "A Note on the Effect of Cost Changes on Prices" Journal of Political Economy, 91, 182-185

[7] Cawley J., M. Eisenberg and R. Avery (2013) "The Effect of Deceptive Advertising on Consumption of the Advertised Good and its Substitutes: The Case of Over-the-Counter Weight Loss Products" Working paper

[8] Chakraborty A. and R. Harbaugh (2014) "Persuasive Puffery" Marketing Science, 33, 382-400

[9] Chen Y. and M. Schwartz (2015) "Differential Pricing When Costs Differ: A Welfare Analysis" Rand Journal of Economics, 46, 442-460

[10] Corts K. (2013) "Prohibitions on False and Unsubstantiated Claims: Inducing the Acquisition and Revelation of Information Through Competition Policy" Journal of Law and Economics, 56, 453-486

[11] Corts K. (2014a) "The Social Value of Information on Product Quality" Economics Letters, 122, 140-143

[12] Corts K. (2014b) "Finite Optimal Penalties for False Advertising" Journal of Industrial Economics, 62, 661-681

[13] Cowan S. (2012) "Third-Degree Price Discrimination and Consumer Surplus" Journal of Industrial Economics, 60, 333-345

[14] Darke P.R. and R.J.B. Ritchie (2007) "The Defensive Consumer: Advertising, Deception, Defensive Processing, and Distrust" Journal of Marketing Research, 44, 114-127

[15] Daughety A.F. and J.F. Reinganum (1997) "Everybody Out of the Pool: Products Liability, Punitive Damages, and Competition" Journal of Law, Economics, and Organization, 13, 410-432

[16] Dranove D. and G. Jin (2010) "Quality Disclosure and Certification: Theory and Practice" Journal of Economic Literature, 48, 935-963

[17] Drugov M. and M. Troya-Martinez (2015) "Vague Lies and Lax Standards of Proof: On the Law and Economics of Advice" Working paper

[18] Fabinger M. and G. Weyl (2015) "A Tractable Approach to Pass-Through Patterns" Working paper

[19] Glaeser E. and G. Ujhelyi (2010) "Regulating Misinformation" Journal of Public Economics, 94, 247-257

[20] Hattori K. and K. Higashida (2012) "Misleading Advertising in Duopoly" Canadian Journal of Economics, 45, 1154-1187

[21] Inderst R. and M. Ottaviani (2013) "Sales Talk, Cancellation Terms, and the Role of Consumer Protection" Review of Economic Studies, 80, 1002-1026

[22] Johnson J.P and D.P. Myatt (2006) "On the Simple Economics of Advertising, Marketing, and Product Design" American Economic Review, 96, 756-784

[23] Kamenica E. and M. Gentzkow (2011) "Bayesian Persuasion" American Economic Review, 101, 2590-2615

[24] Kartik N. (2009) "Strategic Communication with Lying Costs" Review of Economic Studies, 76, 1359-1395

[25] Mailath G.J., M. Okuno-Fujiwara and A. Postlewaite (1993) "Belief-Based Refinements in Signaling Games" Journal of Economic Theory, 60, 241-276

[26] Mayzlin D., Y. Dover and J. Chevalier (2014) "Promotional Reviews: An Empirical Investigation of Online Review Manipulation" American Economic Review, 104, 2421-55

[27] Mezzetti C. and T. Tsoulouhas (2000) "Gathering Information Before Signing a Contract with a Privately Informed Principal" International Journal of Industrial Organization, 18, 667-689

[28] Nelson P. (1974) "Advertising as Information" Journal of Political Economy, 82, 729-754

[29] Piccolo S., P. Tedeschi and G. Ursino (2015) "How Limiting Deceptive Practices Harms Consumers" Rand Journal of Economics, 46, 611-624

[30] Rao A. and E. Wang (2015) "Demand for "Healthy" Products: False Claims in Advertising" Working paper

[31] Renault R. (2015) "Advertising in Markets" in the 'Handbook of Media Economics', S.P. Anderson, D. Stromberg, and J. Waldfogel (eds.), Elsevier

[32] Weyl G. and M. Fabinger (2013) "Pass-Through as an Economic Tool: Principles of Incidence under Imperfect Competition" Journal of Political Economy, 121, 528-583

[33] Zinman J. and E. Zitzewitz (2016) "Wintertime for Deceptive Advertising?" American Economic Journal: Applied Economics, 8, 177-192

Appendix

Proof of Lemma 1. i) If $q \le \underset{\sim}{q}$ demand is zero for all $p \ge 0$, so profit is weakly maximized at $p^* = 0$. ii) If $q > \underset{\sim}{q}$ profit is strictly increasing in $p < a + q$, therefore the optimal price must satisfy $p^* \ge a + q$. At an interior solution, the first order condition is

$$1 - pg\,(p-q)/\left[1 - G\,(p-q)\right] = 0. \tag{19}$$

a) When $q \in \left(\underset{\sim}{q}, \tilde{q}\right)$ the left-hand side of (19) is strictly positive at $p \to a + q$, strictly negative as $p \to b + q$, and strictly decreasing in p because $1 - G(\varepsilon)$ is logconcave. Hence a unique p^* solves equation (19). Define $\sigma(\psi) = -[1 - G(\psi)]g'(\psi)/g(\psi)^2$. Differentiating (19) gives $\partial p^*(q)/\partial q = (1 - \sigma(p^*(q) - q))/(2 - \sigma(p^*(q) - q))$, which lies in $[0, 1)$ because logconcavity of $1 - G(\varepsilon)$ implies $\sigma(\psi) \le 1$. b). When $q \ge \tilde{q}$ the lefthand side of (19) is strictly negative at all $p > a + q$ and hence $p^* = a + q$. □

Proof of Proposition 1. The proof proceeds in several steps. a) Beliefs depend only on the firm's report, so define $\beta_i^e = \Pr(q = H | r = i)$ and $q_i^e = (1 - \beta_i^e) L + \beta_i^e H$ for $i \in \{L, H\}$. b) Conditional on its report and buyer beliefs, the firm's price must maximize its profit. So given a report $r = i$ for $i \in \{L, H\}$, the firm charges $p^*(q_i^e)$. c) As $y^* = \Pr(r(L) = H)$, Bayes' rule implies $\beta_H^e = (1 - x)/(1 - x + xy^*)$, and $\beta_L^e = 0$ if $y^* < 1$. However, Bayes' rule places no restriction on β_L^e if $y^* = 1$. d) y^* must be consistent with the low type behaving optimally. Firstly given $y^* = 0$, $r = L$ is weakly dominant iff $\phi \ge \phi_0$. Secondly given $y^* = 1$, reporting $r = H$ is weakly dominant iff $\phi \le \pi^*(\bar{q}) - \pi^*(q_L^e)$ i.e. for any $\phi \le \phi_1$ given an appropriate off-path belief β_L^e. Thirdly given $y^* \in (0, 1)$, the low type must be indifferent between $r = L$ and $r = H$ i.e. (5) must hold. Moreover $y^* \in (0, 1)$ implies $q_H^e \in (\bar{q}, H)$, such that equation (5) cannot hold for $\phi \notin (\phi_1, \phi_0)$, but has a unique solution for any $\phi \in (\phi_1, \phi_0)$. e) Finally, given buyer beliefs, it is indeed optimal for the high type to report $r = H$. □

Proof of Lemma 3. Recall that y^* strictly decreases in $\phi \in [\phi_1, \phi_0]$. Using (6) and (8):

$$\frac{\partial E(v)}{\partial y^*} = x\left[v^*\left(q_H^e\right) - v^*(L) - \frac{dv^*\left(q_H^e\right)}{dq} \times \left(q_H^e - L\right)\right], \tag{20}$$

i) Consider $L < \hat{q}_v$. a) Under Condition 1i $dv^*(q)/dq$, $d^2v^*(q)/dq^2 > 0$ for $q \in (\underset{\sim}{q}, \tilde{q})$, and $dv^*(q)/dq = 0$ for $q > \tilde{q}$. Hence (20) is strictly negative (positive) for q_H^e below (above) $q^*(L) = \tilde{q}$. b) Under Condition 1ii we have the following results. First, (20) is strictly negative when $q_H^e \le \hat{q}_v$ because $d^2v^*(q)/dq^2 > 0$ for all $q \in (\underset{\sim}{q}, \hat{q}_v)$. Second, (20) is strictly increasing in

q^e_H when $q^e_H > \hat{q}_v$ because in that region $d^2v^*(q^e_H)/dq^2 < 0$. Third, (20) is strictly positive for sufficiently high q^e_H. To see this, note that $dv^*(q)/dq = [1 - G(p^*(q) - q)]/[2 - \sigma(p^*(q) - q)]$, hence $dv^*(L)/dq > 0$, and also $\lim_{q\to\infty} dv^*(q)/dq = 0$ because $\lim_{q\to\infty} p^*(q) - q = a$ and $\lim_{\psi\to a} \sigma(\psi) = -\infty$. Therefore since $v^*(q)$ is strictly convex for $q < \hat{q}_v$ and strictly concave for $q > \hat{q}_v$, we infer that for sufficiently high q^e_H we have $dv^*(q^e_H)/dq < dv^*(z)/dq$ for all $z \in (L, q^e_H)$. Rewriting (20) as $x \int_L^{q^e_H} ((dv^*(z)/dq) - (dv^*(q^e_H)/dq))\, dz$ shows that (20) is strictly positive for sufficiently high q^e_H. Fourth then, (20) has a unique root which we denote by $q^*(L) > \hat{q}_v$, and is strictly negative (positive) for q^e_H below (above) $q^*(L)$. Fifth, note that $q^*(L)$ is strictly decreasing in L because $dv^*(q^*(L))/dq > dv^*(L)/dq$ and so (20) is strictly increasing in L. Also note that $\lim_{L\to\hat{q}_v} q^*(L) = \hat{q}_v$. Finally since q^e_H increases in ϕ, it is immediate that under Condition 1 $E(v)$ is quasiconcave in ϕ. ii) Consider $L > \hat{q}_v$. Given Condition 1, $v^*(q)$ is weakly increasing and concave in $q \geq \hat{q}_v$, so (20) is weakly positive. □

Proof of Proposition 2. i) Note that $q^e_H \leq q^*(L)$ for all ϕ, so by Lemma 3 $E(v)$ is maximized at $\phi^* \geq \phi_0$. ii) Note that $q^e_H < q^*(L)$ when $\phi < \pi^*(q^*(L)) - \pi^*(L)$, and $q^e_H > q^*(L)$ when $\phi > \pi^*(q^*(L)) - \pi^*(L)$. Hence from Lemma, 3 $E(v)$ is maximized at $\phi^* = \pi^*(q^*(L)) - \pi^*(L)$ such that $q^e_H = q^*(L)$. iii) Note that $q^e_H \geq q^*(L)$ for all ϕ, hence by Lemma 3 $E(v)$ is maximized at $\phi^* \leq \phi_1$. Finally, Proposition 1 gives the associated optimal y^* for each case. □

Proof of Corollary 1. Using Proposition 2 optimal false advertising is

$$y^* = \min\left\{\max\left\{\frac{(H - q^*(L))(1-x)}{(H - q^*(L))(1-x) + q^*(L) - \bar{q}}, 0\right\}, 1\right\}. \tag{21}$$

Recall from the proof of Lemma 3 that $q^*(L)$ is weakly decreasing in L. Hence (21) is weakly increasing in L, H, and $(1-x)$. □

Proof of Proposition 3. Given $E(\Pi) = xE(\pi_L) + (1-x)E(\pi_H)$, it is immediate from (11) that a) $E(\Pi) = \pi^*(\bar{q}) - x\phi$ when $\phi < \phi_1$, b) $E(\Pi) = \pi^*(L) + (1-x)\phi$ when $\phi \in [\phi_1, \phi_0]$, and c) $E(\Pi) = x\pi^*(L) + (1-x)\pi^*(H)$ when $\phi > \phi_0$. Hence $E(\Pi)$ is quasiconvex, minimized at ϕ_1, and cannot be maximized at any $\phi \in (0, \phi_0)$. Then for part i), $\phi = \phi_0$ strictly dominates $\phi = 0$ because $\pi^*(q)$ is convex everywhere and strictly convex for $q \in \left(\underset{\sim}{q}, \tilde{q}\right)$. For part ii) note that $\pi^*(q) = a + q$ for all $q \geq \tilde{q}$, and hence $E(\Pi) = a + \bar{q}$ for any $\phi \in \{0\} \cup [\phi_0, \infty)$. □

Proof of Lemma 4. Recall that y^* strictly decreases in $\phi \in [\phi_1, \phi_0]$. Using equation (12):

$$\frac{\partial E(w)}{\partial y^*} = x\left[w^*\left(q_H^e\right) - w^*(L) - \frac{dw^*\left(q_H^e\right)}{dq} \times \left(q_H^e - L\right)\right]. \tag{22}$$

where $w^*(q) = v^*(q) + \pi^*(q)$. i) When $q_H^e < \hat{q}_w$ (22) is strictly negative because $w^*(q)$ is strictly convex for all $q \in \left(\underset{\sim}{q}, \hat{q}_w\right)$. ii) Consider $L < \hat{q}_w < q_H^e$, and define $\check{L} < \hat{q}_w$ as the unique solution to $dw^*(\check{L})/dq = dw^*(q_H^e)/dq$. First, (22) is strictly increasing in $L < \check{L}$ and strictly decreasing in $L > \check{L}$. Second, (22) is continuous in L around $\hat{q}_w$, and (weakly) positive at $L = \hat{q}_w$ because by Condition 2 $w^*(q)$ is weakly concave for $q > \hat{q}_w$. Third, (22) is strictly negative for sufficiently low L. To prove this, note that (22) is proportional to $\frac{w^*\left(q_H^e\right) - w^*(L)}{q_H^e - L} - \frac{dw^*\left(q_H^e\right)}{dq}$. Fixing q_H^e, there exists a $\delta > 0$ such that $\frac{dw^*\left(q_H^e\right)}{dq} > \delta$. Moreover $\frac{w^*\left(q_H^e\right) - w^*(L)}{q_H^e - L}$ is weakly less than $\frac{w^*\left(q_H^e\right)}{q_H^e - L}$, which in turn is strictly less than δ for sufficiently low L. Fourth then, (22) has a unique root $L^*(q_H^e) < \check{L}$, and is strictly negative (positive) for L below (above) $L^*(q_H^e)$. Note also that $L^*(q_H^e)$ is weakly decreasing in q_H^e.[18] Therefore, since q_H^e increases in ϕ, $E(w)$ is quasiconcave in ϕ. iii) When $L > \hat{q}_w$ (22) is weakly positive because $w^*(q)$ is weakly concave for all $q > \hat{q}_w$. □

Proof of Proposition 4. This follows directly from Lemma 4 and its proof. Recall that $L^*(q_H^e)$ is weakly decreasing in q_H^e. i) Since $L \leq L^*(q_H^e)$ for all $y^* \in [0, 1]$, $E(w)$ is maximized at $y^* = 0$. ii) Since $L > L^*(q_H^e)$ for $q_H^e > q^{**}$, and $L < L^*(q_H^e)$ for $q_H^e < q^{**}$, $E(w)$ is maximized by the unique y^* such that $q_H^e = q^{**}$. iii) Since $L \geq L^*(q_H^e)$ for all $y^* \in [0, 1]$, $E(w)$ is maximized at $y^* = 1$. □

Proof of Lemma 5. Part i) follows from arguments in the text. For part ii) look for an equilibrium in which a positive measure of types invest. Since $\pi^*(H) - \pi^*(L) < \infty$ not all types invest, hence $x^*(\phi) \in (0, 1)$. Since the firm's payoff following $r = L$ is independent of q, any firm reporting $r = L$ must have low quality; equivalently, $r(H) = H$. Firstly, in any equilibrium with $y^* = 0$ a firm with $q = L$ earns ϕ_0 less than a firm with $q = H$. Secondly, in any equilibrium with $y^* > 0$ a firm with $q = L$ earns ϕ less than a firm with $q = H$, such that $x^*(\phi) = 1 - F(\phi)$. a) Consider $\phi = \phi_0$. There is clearly an equilibrium with $y^* = 0$. There is no equilibrium with $y^* > 0$, since $\pi^*(q_H^e) - \phi_0 < \pi^*(H) - \phi_0 = \pi^*(L)$, such that no firm with $q = L$ would want to report $r = H$. b) Consider $\phi \in (0, \phi_0)$. There is no equilibrium with $y^* = 0$, since $\pi^*(q_H^e) - \phi = \pi^*(H) - \phi > \pi^*(L)$, such that a firm with $q = L$ would deviate and

[18] Under Condition 2i $w^*(q)$ is linear in $q > \hat{q}_w = \tilde{q}$ such that $L^*(q_H^e)$ in invariant to q_H^e and is the lowest solution to $v^*(\tilde{q}) - v^*(L) + a - \pi^*(L) + L = 0$. Under Condition 2ii $L^*(q_H^e)$ is strictly decreasing in q_H^e because the righthand side of (22) is strictly increasing in both $L = L^*(q_H^e)$ and q_H^e.

report $r = H$. Therefore look for an equilibrium with $y^* > 0$: the gain to a firm with $q = L$ from reporting $r = H$ instead of $r = L$ is $\pi^*(q_H^e) - \phi - \pi^*(L)$. Using $x^*(\phi) = 1 - F(\phi)$, this equals:

$$\pi^* \left(L + (H - L) \frac{F(\phi)}{F(\phi) + y^* (1 - F(\phi))} \right) - \phi - \pi^*(L). \tag{23}$$

This is continuous and strictly decreasing in y^*, and is strictly positive at $y^* = 0$. If (23) is weakly positive at $y^* = 1$ it is strictly positive at all $y^* \in [0, 1)$, hence there is a unique equilibrium with $y^* = 1$. If (23) is strictly negative at $y^* = 1$, there exists a unique equilibrium with $y^* \in (0, 1)$ which makes (23) equal to zero. □

Proof of Proposition 5. i) The proof that $\phi = 0$ is never optimal is given in the text after the proposition. ii) It is enough to show that $\partial E(v)/\partial\phi|_{\phi=\phi_0} < 0$. Note that for $\phi > 0$,

$$\pi^* \left(q_H^e(\phi)\right) = \max \left\{\pi^*(L) + \phi, \pi^* \left(L + (H - L)F(\phi)\right)\right\}, \tag{24}$$

where the first part applies when $y^* \in (0, 1)$, and the second part applies when $y^* = 1$. Equation (24) implies that for some small $\delta > 0$, $\pi^* \left(q_H^e(\phi)\right) = \pi^*(L) + \phi$ for all $\phi \in [\phi_0 - \delta, \phi_0]$. Using $dq_H^e/d\phi = 1/\left(d\pi^*(q_H^e)/dq\right)$ and equation (15), $\partial E(v)/\partial\phi|_{\phi=\phi_0}$ is proportional to

$$\frac{(H - L)\left(dv^*(H)/dq\right) - \left(v^*(H) - v^*(L)\right)}{\left(d\pi^*(H)/dq\right)\left(v^*(H) - v^*(L)\right)(H - L)} + \frac{f(\phi_0)}{F(\phi_0)}.$$

The first term is strictly negative since $H > q^*(L)$, and dominates the second term provided $f(\phi_0)/F(\phi_0)$ is sufficiently small. □

All remaining proofs for the paper are in Section B of the Supplementary Appendix.

Supplementary Appendix

Section A: Further Information on Conditions 1 and 2

This section provides further details on Conditions 1 and 2.

Claim 1. Condition 1 (resp. Condition 2) ensures that buyer surplus (resp. total welfare) is strictly convex for $q \in (\underset{\sim}{q}, \hat{q}_v)$ (resp. $q \in (\underset{\sim}{q}, \hat{q}_w)$), and weakly concave for q above $\hat{q}_v$ (resp. $\hat{q}_w$).

Proof. Using the definitions of $p^*(q)$ and $v^*(q)$ in equations (1) and (4), and also the definition $w^*(q) = v^*(q) + \pi^*(q)$, we have that $d^2v^*(q)/dq^2 \propto z_v(p^*(q) - q)$ and $d^2w^*(q)/dq^2 \propto z_w(p^*(q) - q)$ for all $q \in (\underset{\sim}{q}, \tilde{q})$. Then note that since $1 - G(\varepsilon)$ is logconcave, $p^*(q) - q$ is strictly decreasing in q, with $\lim_{q \to \underset{\sim}{q}} p^*(q) - q = b$ and $\lim_{q \to \tilde{q}} p^*(q) - q = a$. Finally note that for $\tilde{q} < \infty$, $v^*(q)$ and $w^*(q)$ are both linear (and so weakly concave) for all $q > \tilde{q}$. □

Now consider the following generalized setting in which demand equals $s\left[1 - G\left(\frac{p-q-\mu}{m}\right)\right]$, where μ is a location parameter and $m, s \in (0, \infty)$ are stretch parameters (Weyl and Tirole 2012). This corresponds to a setting in which a mass $s > 0$ of buyers have unit demand, and each buyer's valuation is given by $q + \mu + m\varepsilon$ with ε distributed according to $G(\varepsilon)$. In the main text we focus on the case $\mu = 0$ and $m = s = 1$. However in fact:

Claim 2. If Conditions 1 and 2 hold for a demand $1 - G(p - q)$, they also hold for any generalized demand of the form $s\left[1 - G\left(\frac{p-q-\mu}{m}\right)\right]$.

Proof. Consider Condition 1. The market coverage point for this generalized demand is $\tilde{q}(s, m, \mu) = \mu + m(-a + 1/g(a))$, hence $\tilde{q}(s, m, \mu) < \infty$ if and only if $\tilde{q} < \infty$. Also the other threshold $\underset{\sim}{q}(s, m, \mu)$ satisfies $\underset{\sim}{q}(s, m, \mu) = -\infty$ if and only if $\underset{\sim}{q} = -\infty$. Let $\sigma(\psi; s, m, \mu)$ be the curvature of the generalized demand form. We may then write the analogue of $z_v(\psi)$ for this new demand as

$$z_v(\psi; s, m, \mu) = -\frac{d\sigma(\psi; s, m, \mu)}{d\psi} + [2 - \sigma(\psi; s, m, \mu)] \left[\frac{dsG\left(\frac{\psi-\mu}{m}\right)}{d\psi} \Big/ s\left[1 - G\left(\frac{\psi - \mu}{m}\right)\right]\right].$$

After solving for $\sigma(\psi; s, m, \mu)$ and substituting it in, then canceling terms:

$$z_v(\psi; s, m, \mu) = \frac{1}{m}\left[-\sigma'\left(\frac{\psi - \mu}{m}\right) + \left[2 - \sigma\left(\frac{\psi - \mu}{m}\right)\right] \frac{g\left(\frac{\psi-\mu}{m}\right)}{1 - G\left(\frac{\psi-\mu}{m}\right)}\right] \propto z_v\left(\frac{\psi - \mu}{m}\right).$$

Hence $z_v(\psi; s, m, \mu)$ satisfies Condition 1 if and only if $z_v(\psi)$ satisfies it. The proof for Condition 2 is very similar and so is omitted. □

Specific Examples

We now show that Conditions 1 and 2 are satisfied by a wide range of common demand curves. In light of Claim 2 it is sufficient to focus on the case $s = m = 1$ and $\mu = 0$. For further related background material, including a proof that demands with distributions 2-6 below have increasing curvature, see Fabinger and Weyl (2015) and their associated online appendix.

1. Generalized Pareto Distribution: $G(\psi) = 1 - \left(1 - \frac{(1-\sigma)\psi}{(2-\sigma)}\right)^{\frac{1}{1-\sigma}}$ on $\left[0, \frac{2-\sigma}{1-\sigma}\right)$ for $\sigma < 1$, and $G(\psi) = 1 - e^{-\psi}$ on $[0, \infty)$ for $\sigma = 1$. Special cases include the Uniform ($\sigma = 0$) and Exponential ($\sigma = 1$) distributions. Note that $\tilde{q} = (2 - \sigma) < \infty$ and $\sigma(\psi) = \sigma$. Hence Conditions 1i and 2i are satisfied, because $z_v(\psi) = (2 - \sigma) > 0$ and $z_w(\psi) = (3 - \sigma)(2 - \sigma) > 0$.

2. Normal: $G(\psi) = \int_{-\infty}^{\psi} \frac{e^{-\psi^2/2}}{\sqrt{2\pi}} dx$ on $(-\infty, \infty)$. Note that $\underset{\sim}{q} = -\infty$, $\tilde{q} = \infty$, and $\sigma(\psi) = \frac{\psi[1-G(\psi)]}{g(\psi)}$ because $g'(\psi) = -\psi g(\psi)$. Hence $\lim_{\psi \to a} \sigma(\psi) = -\infty$. Moreover

$$z_v(\psi) \propto 2\left(\frac{g(\psi)}{1 - G(\psi)}\right)^2 - 1 - \psi^2. \tag{25}$$

Condition 1ii is satisfied because (25) is negative as $\psi \to -\infty$, is strictly increasing in $\psi \leq 0$ since $\frac{g(\psi)}{1-G(\psi)}$ is strictly increasing, and is strictly positive for all $\psi \geq 0$. To prove the latter, note that for all $\psi \geq 0$ we have the lower bound $\frac{g(\psi)}{1-G(\psi)} \geq \frac{\psi + \sqrt{\psi^2 + 8/\pi}}{2}$ (see Duembgen 2010). In addition

$$z_w(\psi) \propto 6\left(\frac{g(\psi)}{1 - G(\psi)}\right)^2 - 4\psi\frac{g(\psi)}{1 - G(\psi)} - 1 = 6\left(\frac{g(\psi)}{1 - G(\psi)}\right)^2 + 4\frac{g'(\psi)}{1 - G(\psi)} - 1. \tag{26}$$

Condition 2ii is satisfied. Firstly as $\psi \to -\infty$, (26) tends to -1. Secondly (26) is strictly increasing in $\psi < -1$, because $\frac{g(\psi)}{1-G(\psi)}$ and $g'(\psi) > 0$ are both strictly increasing. Thirdly (26) is strictly positive for all $\psi \in [-1, 0]$. This can be proved by noting that on this interval, we have the lower bound $g(\psi) \geq \left(1 - \frac{\psi^2}{2}\right)/\sqrt{2\pi}$, and the upper bound $1 - G(\psi) \leq \frac{1}{2} - xg(0)$. Fourthly (26) is also strictly positive for all $\psi > 0$. This can be proved by noting that $\frac{g(\psi)}{1-G(\psi)}$ strictly increasing implies $2\left(\frac{g(\psi)}{1-G(\psi)}\right)^2 > 2\left(\frac{g(0)}{1-G(0)}\right)^2 > 1$, and also $4\left[\left(\frac{g(\psi)}{1-G(\psi)}\right)^2 + \frac{g'(\psi)}{1-G(\psi)}\right] > 0$.

3. Weibull: $G(\psi) = 1 - e^{-\psi^\alpha}$ on $[0, \infty)$ where $\alpha > 1$. Note that $\underset{\sim}{q} = -\infty$, $\tilde{q} = \infty$, $\sigma(\psi) = 1 - \left(\frac{\alpha-1}{\alpha\psi^\alpha}\right)$ and $\lim_{\psi \to a} \sigma(\psi) = -\infty$. Moreover

$$z_v(\psi) \propto (\alpha - 1)(\psi^\alpha - 1) + \alpha\psi^{2\alpha} \text{ and } z_w(\psi) \propto 2\alpha^2\psi^{2\alpha} + 3\alpha(\alpha - 1)\psi^\alpha - (\alpha - 1) \tag{27}$$

Conditions 1ii and 2ii are both satisfied, since both expressions in (27) are strictly negative as $\psi \to 0$, strictly increasing in ψ and strictly positive as $\psi \to \infty$.

4. Type I Extreme Value (Max version): $G(\psi) = e^{-e^{-\psi}}$ on $(-\infty, \infty)$. Note $\underset{\sim}{q} = -\infty$, $\tilde{q} = \infty$, $\sigma(\psi) = (e^{\psi} - 1)(e^{e^{-\psi}} - 1)$ and $\lim_{\psi \to a} \sigma(\psi) = -\infty$. Numerical simulations show that Conditions 1i and 2i are both satisfied.

5. Logistic: $G(\psi) = \frac{e^{\psi}}{1+e^{\psi}}$ on $(-\infty, \infty)$. Note that $\underset{\sim}{q} = -\infty$, $\tilde{q} = \infty$, $\sigma(\psi) = 1 - e^{-\psi}$ and $\lim_{\psi \to a} \sigma(\psi) = -\infty$. Condition 1ii is satisfied because $z_v(\psi) \propto e^{2\psi} - 1$, which is single-crossing from negative to positive at $\psi = 0$. However Condition 2ii is not satisfied since $z_w(\psi) \propto 2 + 2e^{-\psi}$, which is strictly positive everywhere.[19]

6. Type I Extreme Value (Min version): $G(\psi) = 1 - e^{-e^{\psi}}$ on $(-\infty, \infty)$. Note that $\underset{\sim}{q} = -\infty$, $\tilde{q} = \infty$, $\sigma(\psi) = 1 - e^{-\psi}$ and $\lim_{\psi \to a} \sigma(\psi) = -\infty$. Condition 1ii is satisfied because $z_v(\psi) \propto e^{-\psi}(1 - e^{-\psi}) - 1$, which is single-crossing from negative to positive at $\psi = \ln\left(\frac{-1+\sqrt{5}}{2}\right)$. However Condition 2ii is not satisfied since $z_w(\psi) \propto 2 + 3e^{-\psi}$, which is strictly positive everywhere.

Section B: Remaining Proofs

Proof of Lemma 6 and Proposition 6. We prove Lemma 6 and Proposition 6 together, in several steps.

1) Given that beliefs are price-independent, $E(q|r)$ fully determines prices. Hence $\mathbf{y}^*$ is necessary and sufficient to write down expected buyer surplus, total welfare, and profit (before penalties are deducted). Lemma 6i then follows (we return to 6ii later).

2) Buyer surplus. Firstly, buyer surplus is not maximized if any report $r = q_{i<n}$ is sent by more than one type. To see why, consider a new triangular matrix with $y'_{i,i} = \sum_{j=1}^{n-1} y^*_{i,j}$ and $y'_{i,n} = y^*_{i,n}$ for all $i < n$. Strict convexity of $v^*(q) \in (\underset{\sim}{q}, \tilde{q})$ implies that buyer surplus is strictly higher, by Jensen's inequality. Secondly, buyer surplus is not maximized if $E(q|r = q_n) > \tilde{q}$ and $y^*_{i,n} < 1$ for some $i < n$. This is because the derivative of expected buyer surplus with respect to $y^*_{i,n}$ is $x_i \left[v^*(\tilde{q}) - v^*(q_i)\right] > 0$. Thirdly, buyer surplus is not maximized if $E(q|r = q_n) = \tilde{q}$, and there exists some $j < k$ such that $y^*_{k,n} < 1$ but $y^*_{j,n} > 0$. To see this, note that $\left.\frac{\partial y^*_{j,n}}{\partial y^*_{k,n}}\right|_{E(q|r=q_n)=\tilde{q}} = -\frac{x_k(\tilde{q}-q_k)}{x_j(\tilde{q}-q_j)}$. The derivative of $E(v)$ with respect to $y^*_{k,n}$, whilst adjusting $y^*_{j,n}$ to ensure $E(q|r = q_n) = \tilde{q}$, is proportional to

$$(\tilde{q} - q_j)\left[v^*(\tilde{q}) - v^*(q_k)\right] - (\tilde{q} - q_k)\left[v^*(\tilde{q}) - v^*(q_j)\right],$$

which is strictly positive since $v^*(q)$ is strictly convex. Proposition 6i then follows.

3) Profit. Since $\pi^*(q)$ is convex, and strictly so for $q \in (\underset{\sim}{q}, \tilde{q})$, a similar approach to the first part of the previous step shows that expected profit (before penalties are deducted) is maximized by $y^*_{i,i} = 1$ for all i. Hence expected profit once penalties are deducted, is also maximized by

[19] Consequently a welfare-maximizing policymaker always optimally induces $y^* = 0$. This is also true for the next distribution.

$y^*_{i,i} = 1$ for all i, and Proposition 6ii follows.

4) Total welfare. Firstly, total welfare is not maximized if any report $r = q_i$ for $i < n$ is sent by more than one type, and the proof is similar to that for buyer surplus. Secondly, if $E(q|r = q_n) > \tilde{q}$ and there exists some $i < n$ with $y^*_{i,n} < 1$, total welfare is increasing in $y^*_{i,n}$ if and only if $q_i \geq L^*$. To see this, the derivative of $E(TW)$ with respect to $y^*_{i,n}$ is $v^*(\tilde{q}) + a + q_i - v^*(q_i) - \pi^*(q_i)$, which is positive if and only if q_i exceeds a threshold (which we call L^*). Thirdly, total welfare is not maximized if $E(q|r = q_n) = \tilde{q}$, and there exists some $j < k$ such that $y^*_{k,n} < 1$ but $y^*_{j,n} > 0$. The proof closely follows the same arguments for buyer surplus. Proposition 6iii then follows.

5) Implementation. Note that the maximum gain from false advertising is $\bar{\phi} = \pi^*(q_n) - \pi^*(q_1)$. First, set $\phi(q_i, q_j) = \bar{\phi}$ for all $j \notin \{q_i, q_n\}$ so that in any equilibrium, each firm either reports truthfully or reports $r = q_n$. Second, for any type i for whom $y^*_{i,i} = 1$, also set $\phi(q_i, q_n) = \bar{\phi}$. Third, for any type i for whom $y^*_{i,n} = 1$, set $\phi(q_i, q_n) = 0$. Fourth, let $q^e_n = (\sum_{j=1}^n x_j y^*_{j,n} q_j)/(\sum_{j=1}^n x_j y^*_{j,n})$. For any type i for whom $y^*_{i,i} = 1 - y^*_{i,n}$ and $y^*_{i,n} \in (0,1)$ (there is at most one such i) set $\phi(q_i, q_n) = \pi^*(q^e_n) - \pi^*(q_i)$. Fifth, it is easy to see there is a unique equilibrium outcome in which $\mathbf{y}^*$ is played, and so Lemma 6ii follows. □

Proof of Lemma 7. The proof closely follows that of Proposition 1.

a) As usual let $q^e_H = E(q|r = H)$ and $y^* = \Pr(r(L) = H)$. The second restriction implies that following $r = H$ the firm charges $p^*(q^e_H; H)$. Bayes' rule implies that following $r = H$ and $p = p^*(q^e_H; H)$ the firm is believed to have high quality with probability $(1-x)/(1-x+xy^*)$.

b) Suppose $r = L$ is on-path. Firstly if a firm reports $r = L$ its price must maximize profit given buyer beliefs. Secondly buyer beliefs must satisfy Bayes' rule following $r = L$ and any on-path price(s). Hence given the first restriction, a firm that reports $r = L$ must charge $p^*(L; L)$, and be believed to have low quality with probability 1.

c) Necessary conditions for optimality of the low type's behavior: Firstly given $y^* = 0$, reporting $r = L$ is weakly dominant only if $\phi \geq \phi'_0$. Secondly given $y^* = 1$, reporting $r = H$ is weakly dominant only if $\phi \leq \phi'_1$. Thirdly given $y^* \in (0,1)$, the low type is indifferent between $r = L$ and $r = H$ iff (16) holds. Note that for $c(H) - c(L)$ small, $\phi'_1 < \phi'_0$, and that (16) has a unique solution $y^* \in (0,1)$ if and only if $\phi \in (\phi'_1, \phi'_0)$.

d) The conditions given in the previous step are also sufficient for optimality of the low type's behavior, given appropriate off-path beliefs such as those in the lemma.

e) Clearly the high type strictly prefers to report $r = H$ and charge $p^*(q^e_H; H)$ for appropriate off-path beliefs, such as those in the lemma. □

Proof of Lemma 8. We can simply repeat all the steps used in the proof of Proposition 1. The only difference is that in the second step, each firm's price maximizes its profits given buyer beliefs and its conjecture about the other firm's price. Hence following a report $r = i$ for $i \in \{L, H\}$, the firms play Nash equilibrium prices $p_I^*(q_I, E(q_E|r = i))$ and $p_E^*(E(q_E|r = i), q_I)$. $\square$

Proof of Proposition 7. Under full information:

$$v^*(q_I, q_E) = \begin{cases} -\frac{5t}{4} + \frac{q_I+q_E}{2} + \frac{(q_E-q_I)^2}{36t} & \text{if } q_E \in \left(\underset{\sim}{q_E}, \tilde{q}_E\right) \\ q_I + \frac{t}{2} & \text{if } q_E \geq \tilde{q}_E \end{cases}$$

Expected buyer surplus is $E(v) = x(1 - y^*)v^*(q_I, L) + (1 - x + xy^*)v^*(q_I, q_H^e)$. Given $L \geq \underset{\sim}{q_E}$, $v^*(q_I, q_E)$ has the same shape as $v^*(q)$ in the monopoly problem under Condition 1i with $\hat{q}_v = \tilde{q}_E$. Hence the proposition is proved in a similar way to Proposition 2, just with $\tilde{q}_E$ replacing $q^*(L)$ and $\hat{q}_v$. $\square$

Proof of Proposition 8. Under full information, $w^*(q_I, q_E) = v^*(q_I, q_E) + \pi_I^*(q_I, q_E) + \pi_E^*(q_I, q_E)$ equals:

$$w^*(q_I, q_E) = \begin{cases} -\frac{t}{4} + \frac{q_I+q_E}{2} + \frac{5(q_E-q_I)^2}{36t} & \text{if } q_E \in \left(\underset{\sim}{q_E}, \tilde{q}_E\right) \\ q_E - \frac{t}{2} & \text{if } q_E \geq \tilde{q}_E \end{cases}$$

Expected total welfare is $E(w) = x(1 - y^*)w^*(q_I, L) + (1 - x + xy^*)w^*(q_I, q_H^e)$. When $L \in [\underset{\sim}{q_E}, \tilde{q}_E)$ direct computation reveals that a) $\partial E(w)/\partial y^* < 0$ when $q_H^e \leq \tilde{q}_E$, and b) for $q_H^e > \tilde{q}_E$, $\partial E(w)/\partial y^* < 0$ if and only if $L < L_E^* = q_I + 3t/5$. Hence the claim can be proved using a similar approach as in Proposition 4. $\square$

References

[1] Duembgen L. (2010) "Bounding Standard Gaussian Tail Probabilities" arXiv:1012.2063

[2] Fabinger M. and G. Weyl (2015) "A Tractable Approach to Pass-Through Patterns" Working paper

[3] Weyl G. and J. Tirole (2012) "Market Power Screens Willingness-to-Pay" Quarterly Journal of Economics, 127, 1971-2003