Centre for Economic and Financial Research at New Economic School

Vague lies and lax standards of proof: On the law and economics of advice

Mikhail Drugov
Marta Troya-Martinez

Vague lies and lax standards of proof:
On the law and economics of advice[*]

Mikhail Drugov[†]and Marta Troya-Martinez[‡]

June 21, 2018

Abstract

This paper analyses a persuasion game where a seller provides (un)biased and (im)precise advice and may be fined by an authority for misleading the buyers. In the equilibrium, biasing the advice and making it noisier are complements. The advice becomes both more biased and less precise with a stricter standard of proof employed by the authority, a larger share of credulous consumers, and a higher buyers' heterogeneity. The optimal policy of the authority is characterized in terms of a standard of proof and resources devoted to the investigation.

Keywords: Advice, Persuasion, Legal Procedure, Consumer Protection.

JEL Classification: D18, D8, K4, L1.

1   Introduction

Sales advice is pervasive in transactions where consumers cannot assess the product's quality or its utility prior to the purchase because of its complexity or number of experience attributes. Consumer electronics, insurance, banking and medical care contracts are important examples. In persuading the consumer, the seller can tell an outright lie about the desirability of the product, for instance, by exaggerating the quality of a financial product with the "hope to gain high commissions, or to achieve their sales targets for certain product". But the seller can also choose to give vague advice such as when a financial advisor does "not clearly define or explain the benchmarks for 'low-risk', 'mid-risk', or 'high-risk'. As a result, the client may not have accurately reflected his or her risk preferences" (see EC (2011), pp. 49-50).

Giving a misleading sales pitch, although profitable, may come at a cost. Customers complain because the product does not meet their expectations. These complaints damage the seller's reputation, draw the attention of consumer associations and may even trigger an enforcement action by a regulator or result in litigation. Indeed, the types of products and services mentioned above consistently top the list of consumers complaints worldwide, which raises the policy relevance of the problem.[1]

This paper explores how a seller strategically uses both outright lies and vagueness, when an authority can investigate and sanction the seller for misleading the consumers. It shows that lies and vagueness are complements in the equilibrium, that is, a bigger lie is also more vague. It also sheds light on how, given the seller's behavior and consumers' reaction, the authority should decide on the standard of proof as well as on the resources devoted to the investigation.

In the model, buyers are interested in buying a product that has some attributes whose usefulness will only become fully known through the use (Nelson (1970)). The seller knows the product's features but does not know how well they fit with the buyer's needs. As a result, the buyers' valuations for the product, determined by the quality of the match between the product characteristics and buyers' idiosyncratic preferences, are unknown at the point of sale.

The seller gives advice about the product. It generates an informative signal equal to the sum of the true match quality and an error term, both normally distributed.

The error term represents frictions in the communication as well as the experience features of the product. The seller secretly chooses its mean ("bias") and publicly its variance ("noise"). The posterior valuation of the product is the buyer's willingness to pay after the sales pitch. Some consumers are rational and correctly update their beliefs while others are credulous with the posterior equal to the signal realization.[2] The posterior valuation is shared in a fixed proportion between the seller and the buyer.

The bias unambiguously increases the perceived quality of the product and thus the seller's revenues. Because the (biased) signal is, on average, better than the true match quality, the seller wants rational buyers to pay more attention to the (biased) advice rather than to the prior. As a result, the seller would like to accompany a larger bias with a smaller noise. These incentives are akin to the ones in Johnson and Myatt (2006) for niche markets.[3]

However, the bias does not come for free. Misled buyers learn through use the true match quality and complain, which triggers an action by an authority that might be a consumer protection authority, a sectoral regulator or the court depending on the product and the country. The authority investigates the seller by surveying a random sample of customers or sending mystery shoppers. Based on this information, it estimates the bias and determines whether there is enough evidence that the seller has misled consumers by biasing his signal. More precisely, the authority presumes the innocence of the seller. That is, its null hypothesis is that there has been no bias. It then tests whether this null hypothesis can be rejected in favor of the alternative hypothesis of a positive bias. In doing so, it uses a significance level which is the standard of proof. If the seller is found guilty, he has to pay a fine that depends on the estimated bias. A larger bias always increases the expected fine. Noise affects the expected fine through two channels: it decreases the probability that the seller is found guilty but increases the fine if the seller is found guilty. The total effect is U-shaped. For a given bias, the costs are minimized at some intermediate level of noise.

We show that when the choice of the amount of information is endogenously

---

[2]Our model does not need credulous consumers to generate the main results but they make misleading by the seller harmful and allow for additional interesting comparative statics.

[3]Building on Lewis and Sappington (1994), Johnson and Myatt (2006) show the following. In a niche market, the price is above the prior expected valuation and, hence, in the absence of any information the buyer will not buy the product. The seller wants to provide as much (precise) information as possible since he benefits from the heterogeneity in the posterior valuations. The opposite happens in a mass market where the price is below the prior expected valuation.

linked to the future fine, two new important results emerge.[4] First, unlike Johnson and Myatt (2006), extreme policies regarding information disclosure are no longer optimal. Instead, the seller only discloses partial information about the product. Second, bias and noise are complements, that is, the seller "hides" a larger bias with a larger noise. For instance, in a less onerous punishment regime with a stricter standard of proof, the seller uses a larger bias. As a result, there is more need for noise. This is despite the fact that rational buyers pay less attention to the signal when the noise is larger and, therefore, are less easily swayed by bias.

We then characterize the optimal policy of the authority, assuming that some of the sellers are honest and never bias their advice. The policy of the authority consists of a standard of proof and the amount of resources devoted to the investigation which is reflected by the number of sampled consumers. The authority then minimizes a combination of the consumers' harm, the fines paid by honest sellers and the costs of investigation. A lower standard of proof reduces the sellers' incentives to bias their advice, but increases the likelihood of imposing a fine on honest sellers (type I error). The larger the scope of the investigation, the more precise the authority's test should be. This decreases both the bias and the type I error, but costs more in terms of resources. We find that the optimal standard of proof is lower and more resources are devoted to the investigation in a market where consumers are more heterogenous or there are more credulous consumers. A higher share of honest sellers leads to a higher standard of proof but the effect on the scope of investigation is non-monotonic.

Biasing the advice can be seen as a form of false advertising which arises when a low quality firm advertises itself as being high quality, that is, when the equilibrium of the signalling game is (semi-)pooling. While it was discussed by Nelson (1974), it has received attention only recently (see Renault (2016) for the latest survey of the advertising literature). In Corts (2013), Piccolo, Tedeschi and Ursino (2015) and Rhodes and Wilson (2018) consumers are rational and the firm uses advertising to signal its quality; a possible fine makes false advertising costly. In Glaeser and Ujhelyi (2010) and Hattori and Higashida (2012) consumers are credulous and false advertising boosts the demand and hence reduces the quantity distortion coming from the imperfect competition. Formally, however, our model is quite different since it

---

[4]The literature on strategic communication with lying costs as, for example, Kartik (2009), assumes some exogenous lying costs that put some discipline on the sender. Inderst and Ottaviani (2013) analyze contract cancellation and product return policies in a market where a more informed seller advises the buyer. Since advice in their model takes the form "to buy" or "not to buy", there is no room for the seller to use noise, only bias. We micro-found the lying costs as well as explicitly separate the content of information and its precision.

is a signal-jamming model à la Holmström (1999) with the specific feature that the seller also chooses the variance of the signal observed by the buyers.

This paper is also related to the literature on litigation and contributes to it in three important ways.[5] First, we microfound the signal obtained by the court about the behavior of the defendant (seller) while typically the literature assumes some unspecified process by which the defendant's action generates a signal. Second, we allow the defendant to change the quality of the signal (variance) about the harmful action (bias). Finally, we let the policy of the court (authority) change the equilibrium communication between the seller and the buyers rather than just the individual decision problem as in most papers.

The rest of the paper is organized as follows. Section 2 introduces the model. Section 3 solves for the equilibrium bias and noise. Section 4 characterizes the optimal policy of the authority. Section 5 considers the scenario when the price is fixed. Section 6 concludes. The proofs are contained in the Appendix.

2 The model

Buyers approach the seller in order to get information about his product and buy it. The seller does not know their product valuations while the buyers do not know the quality of the product and its features. Thus, at the beginning of the interaction, the match quality of the transaction, θ, is unknown both to the seller and the buyer.[6] Both the seller ("he") and the buyers ("she" when in singular) know that θ is distributed as $\mathcal{N}(\mu, \sigma^2)$. The average surplus from the interaction is thus μ while σ^2 reflects buyers' heterogeneity. Buyers learn θ after buying the product. The production costs are normalized to zero. The number of buyers is normalized to be of measure 1.

The communication process and buyers' valuation The seller provides information to buyers by revealing product characteristics, giving advice, advertising,

---

[5]See for instance, Rubinfeld and Sappington (1987), Miceli (1990), Davis (1994), Schrag and Scotchmer (1994), Lewis and Poitevin (1997), Demougin and Fluet (2008) and Kaplow (2011), to name just a few.

[6]This specification is also used in the advertising literature where it is assumed that the reservation price (or "match value") is unknown to the firm and to the consumer. See, for instance, Anderson and Renault (2006, 2009) and Johnson and Myatt (2006). This is of course also the setting used in the persuasion literature (e.g., Kamenica and Gentzkow (2011)), where the sender and the receiver have the same prior and the sender commits to the signal technology before knowing the realization of the state.

etc. In doing so, he can distort the communication strategically: he can exaggerate some positive features of the product and also be vague about them. More precisely, the seller provides information that generates an informative but possibly biased and noisy signal S which takes the following form:

$$S = \theta + \varepsilon,$$

where ε is the distortion introduced by the seller. It is distributed normally, $\varepsilon \rightsquigarrow \mathcal{N}(\beta, \eta^2)$, and both moments are controlled by the seller. We refer to β as *bias* and to η as *noise*. The signal is therefore distributed as $\mathcal{N}(\mu + \beta, \sigma^2 + \eta^2)$; denote its cdf and pdf by G and g, respectively. We assume that the noise is bounded from above since the seller is usually obliged to provide at least some information about the product.[7] We also assume that there is some minimum level of noise due to the product experience features. Otherwise, the seller could perfectly reveal the match quality which, by assumption, he does not know.

There are two types of buyers: rational and credulous, with shares $1 - c$ and c, respectively. Credulous buyers do not understand that the seller might provide biased and imprecise information and blindly believe the seller's signal. These buyers think that their valuation is equal to the realization of the signal.[8]

Rational buyers, instead, correctly interpret the signal. In particular, they do not observe the bias but have conjecture $\tilde{\beta}$ about it (which has to be correct in the equilibrium). However, they observe the noise, since they can evaluate how precise the seller's explanations are, how many details he provides, whether there is a trial period, etc. See the previous version Drugov and Troya-Martinez (2012) for the case of unobservable noise which yields qualitatively the same results. This also explains why we assume that the buyer's conjecture does not depend on the observed noise

---

[7] For instance, in the UK, "A Key Features Document is required to be provided for life policies, personal pensions schemes, stakeholder pensions schemes, investment trust savings schemes and cash deposit ISAs." (http://www.fca.org.uk/firms/being-regulated/meeting-your-obligations/cobs/disclosure).

[8] See Glaeser and Ujhelyi (2010) and Hattori and Higashida (2012) for models of credulous consumers and false advertising and Kartik, Ottaviani and Squintani (2007) for credulous receivers in a cheap talk model.

There is an extensive empirical evidence about credulous buyers. For instance, De Franco, Lu and Vasvari (2007) and Malmendier and Shanthikumar (2014) find that individual investors take the analysts' recommendations at face value while institutional investors adjust for the analysts' credibility. OFT (2011) highlights the evidence that "people with less education seem to place more trust in advisers" (p.57). More examples and evidence can be found in Gabaix and Laibson (2006), Inderst and Ottaviani (2009) and DellaVigna and Gentzkow (2010).

Alternatively, credulous consumers have an (improper) uniform prior over θ or they are in fact rational buyers that have a lapse with probability c.

which is an assumption used in the literature, see, for example, Judd and Riordan (1994).

A rational buyer who received realization s of the signal has the expected match quality equal to

$$E\left[\theta \mid s, \eta, \tilde{\beta}\right] = \mu + \left(s - \mu - \tilde{\beta}\right)\frac{\sigma^2}{\sigma^2 + \eta^2}. \tag{1}$$

A rational buyer takes the signal into account with the weight proportional to the prior variance. Expression (1) can also be written as $\left(\frac{\mu}{\sigma^2} + \frac{s-\tilde{\beta}}{\eta^2}\right)/\left(\frac{1}{\sigma^2} + \frac{1}{\eta^2}\right)$, that is, as the weighted average of the ex-ante quality of match and the ex-post signal realization (corrected for the bias), where the weights are precisions of the prior and the signal.

We assume that the seller charges the whole valuation of the buyer, though the results hold also for any fixed fraction.[9] That is, the seller sees the effect that the realization of the signal, s, has had on the buyer and hence, he knows that a rational buyer is ready to pay up to $E\left[\theta \mid s, \eta, \tilde{\beta}\right]$ while a credulous one is ready to pay up to s.[10] Therefore, the seller's revenues are equal to the expected valuation of the buyer:

$$\begin{aligned} R\left(\beta, \eta\right) &= (1 - c)\int_{-\infty}^{+\infty} E\left[\theta \mid s, \eta, \tilde{\beta}\right]g(s)ds + c\int_{-\infty}^{+\infty} sg(s)ds \\ &= \mu + (1 - c)\left(\beta - \tilde{\beta}\right)\frac{\sigma^2}{\sigma^2 + \eta^2} + c\beta. \end{aligned} \tag{2}$$

The seller extracts from the buyer the prior expected quality of match μ plus how much the buyer is misled into thinking that the product is better than it is. Note as well that the revenues decrease with the noise η if $\beta > \tilde{\beta}$ since rational buyers pay less attention to the signal, i.e., place a smaller weight on it.

---

[9]It is mainly done for technical reasons since we can then integrate with respect to the whole support of the signal distribution. However, there is also evidence that the use of bargaining and its consequent price discrimination has recently increased in the retail sector in areas such as travel agents, car dealers, credit card companies and electrical retailers (see The Sunday Times (2008), The Telegraph (2008) and The Economist (2009)). Section 5 considers a fixed price.

[10]This entails some negative payments which can be made negligible by making the average quality μ high enough. As we will see later, μ does not affect equilibrium bias and noise (Corollary 1). It also means that the seller always sells to the consumer. This is in the spirit of the career concerns models (Holmström (1999)), where a worker always gets a wage. The only reason it is done both there and in this paper is technical as it allows to integrate over the whole support of the distribution. In a different model, Spiegler (2006) also allows prices to be negative for technical reasons.

Estimated bias and seller's costs   When the buyers buy the product and start using it, they discover the true match quality. For some buyers, it will be lower than expected, and they will complain to a public body that we call the "authority" throughout the paper. We assume that this authority can inflict a punishment on the seller. Depending on the nature of the authority, the punishment may be publishing a negative report, ordering to withdraw a certain advertisement, prohibiting a certain commercial practice, refusing to grant a licence or imposing a fine on the seller. For the sake of simplicity, we take the last meaning and treat the punishment as a monetary fine.

Since bias is a deliberate and conscious way to mislead consumers, it is illegal and a fine is imposed on the seller if evidence of a positive bias is established.[11] The noise, however, is not punished since due to the experience nature of the product, there is some minimum noise that may be seller specific. Furthermore, the seller may provide information that is vague and open to interpretation, so that the buyers themselves make personal, independent, errors of interpretation. Poor information at the point of purchase may also be due to incompetent sales staff who have trouble giving clear advice. If these incompetent sales staff are nonetheless objective (i.e., they do not use bias), the buyers will not be misled on average.

The authority conducts an investigation by taking a random sample of N buyers in order to estimate the bias introduced by the seller. It can also send "mystery shoppers" to the seller who will report their experience afterwards. Competition and consumer protection authorities from the US and the European countries routinely commission mystery shopping exercises as part of their market studies.[12]

Each mystery shopper (or buyer) $i = 1, ..., N$ reports the signal s_i observed from the seller and quality θ_i.[13] Having received their reports, the authority uses a statis-

---

[11] In the European Union, the Unfair Commercial Practices Directive 2005/29/EC defines a commercial practice as misleading "...if it contains false information and is therefore untruthful or in any way, including overall presentation, deceives or is likely to deceive the average consumer..." (Article 6). Such practices are more generally called unfair and "...shall be prohibited" (Article 5).

See also the discussion about the regulator's need to establish "intentionality" of the violations by insurance agents in Brown and Minor (2014), p. 7.

[12] See FTC (2009), OFT (2010) and EC (2011) for recent examples in the markets of children entertainment, retail investment advice and debt management. Other tests are often conducted by courts and consumer bodies. For instance, "copy tests" are used to determine whether an advert is misleading. If enough consumers are misled, the consumer protection authority may order the advert to be withdrawn. The evidence may also come from a class action. See Issacharoff (1999) for a discussion of class actions and consumer protection.

[13] For instance, the authority could provide the shopper with a list of relevant attributes of the product. Then, for each attribute, the shopper needs to report what he/she expected from the advice and what he/she really found. This could be done in relatively coarse terms (i.e. better,

tical test to determine if the seller has used some bias. It computes the error terms $\varepsilon_i = s_i - \theta_i$ and estimates the bias as

$$\widehat{\beta} = \frac{1}{N} \sum_{i=1}^{N} \varepsilon_i \tag{3}$$

Since $\varepsilon_i \rightsquigarrow \mathcal{N}\left(\beta, \eta^2\right)$, this estimator is distributed as $\mathcal{N}\left(\beta, \frac{\eta^2}{N}\right)$.

There is presumption of innocence, that is, by default the seller is assumed not to have introduced any bias, unless enough evidence is provided.[14] To determine how convincing the evidence about the use of bias should be, the authority uses a standard hypothesis test where the null hypothesis of no bias, $H_0 : \beta = 0$, is assessed against the alternative $H_1 : \beta > 0$. The authority constructs the statistics $\frac{\widehat{\beta}}{\eta/\sqrt{N}}$ which, under the null H_0, is distributed as $\mathcal{N}(0,1)$. Denote z_α the threshold such that H_0 is rejected and hence, the seller is found guilty of biasing if and only if $\frac{\widehat{\beta}}{\eta/\sqrt{N}} \geq z_\alpha$, where α is the significance level of the test (i.e., the probability of incorrectly rejecting the null hypothesis). A natural interpretation of z_α is the "standard of proof". With a higher z_α (lower α), it is more difficult to reject H_0 and, therefore, the authority needs more evidence to convict the seller. For instance, if α is 5%, then $z_\alpha \approx 1.64$, and if α is 10%, then $z_\alpha \approx 1.28$.

It is widely accepted that penalties imposed on firms should be proportionate to the harm (the consumer detriment) caused by the firm's violations. Accordingly, if the seller is found guilty, a fine is imposed which is an increasing function of the estimated bias $\widehat{\beta}$. In this section, we take the fine to be equal to $\widehat{\beta}$ and consider punitive damages and damages that depend on the share of credulous consumers $dc\widehat{\beta}$ in the end of Section 4. Therefore, the authority uses a two-stage enforcement procedure that first determines whether bias has been used and, if the firm is found guilty, then determines the penalty. This is the standard procedure in practice.[15]

---

worse or as expected) and then the authority could aggregate the answers to obtain a measure of s_i and θ_i.

[14]The presumption of innocence is natural in court. When it is a competition authority that punishes the seller the presumption of innocence is explained by the fact that the competition authority may need to defend its position in court if the seller decides to appeal.

[15]For example, Sunstein et al. (2002) say (p. 8): "...the jury must first determine what is called the *liability* of the defendant. The question here is whether the jury finds the defendant to be liable - legally responsible - for the harm suffered by the plaintiff. If the jury finds the defendant liable, the jury must secondly determine the level of damages that will compensate the plaintiff for the harm suffered by the actions of the defendant..." (italics in the original).

The seller's costs are the expected fine:

$$C(\beta, \eta) = \frac{\eta}{\sqrt{N}} \int_{z_\alpha}^{+\infty} z \, dH(z) \tag{4a}$$

$$= \beta \left(1 - \Phi \left(z_\alpha - \frac{\beta \sqrt{N}}{\eta} \right) \right) + \frac{\eta}{\sqrt{N}} \phi \left(z_\alpha - \frac{\beta \sqrt{N}}{\eta} \right) \tag{4b}$$

where $z \rightsquigarrow \mathcal{N} \left(\beta \frac{\sqrt{N}}{\eta}, 1 \right)$ and H is its cdf (it is the distribution of $\frac{\widehat{\beta}}{\eta/\sqrt{N}}$); and Φ and ϕ are the cdf and pdf of the standard normal random variable, respectively. The first term in (4b) corresponds to how much the seller pays on average multiplied by the probability of being found guilty. The second term corrects for the selection bias as the truncation selects higher values of $\widehat{\beta}$. The crucial property of this cost function is that it is U-shaped with respect to the noise. A higher η makes it less likely for the seller to be found guilty (since the distribution of z shifts to the left, the integral in (4a) decreases), but increases the chances of a large fine if he is found guilty (this integral is multiplied by a larger number).

Lemma 1 *Expected fine (4) (i) increases with bias β, (ii) decreases with noise η when it is smaller than $\beta \sqrt{N} z_\alpha$ and increases otherwise.*

3 Equilibrium

In this section, we derive the equilibrium of the game between the seller and the buyers for a given policy of the authority, that is, for some standard of proof z_α and sample size N. In Section 4 we find the optimal policy of the authority (z_α^*, N^*).

The seller maximizes his profits $\Pi(\beta, \eta)$ which are equal to the revenues (2) minus the costs (4). An equilibrium is a pair (β^*, η^*) such that: (i) it maximizes seller's profits $\Pi(\beta, \eta)$ given rational buyer's conjecture $\tilde{\beta}$ and, (ii) the rational buyer's conjecture about the bias is correct, $\tilde{\beta} = \beta^*$. The next proposition derives the equilibrium in a closed form. To guarantee the existence of an interior equilibrium we make the following two assumptions. Denote $\lambda \equiv 1 - \Phi \left(z_\alpha - \frac{1}{z_\alpha} \right) + z_\alpha \phi \left(z_\alpha - \frac{1}{z_\alpha} \right)$.

Assumption 1 $\lambda > c$.

In the equilibrium, the marginal revenue from increasing the bias is equal to c since rational consumers correctly anticipate the bias, $\tilde{\beta} = \beta^*$, see (2). The marginal

cost of the bias is the derivative of $C(\beta, \eta)$ with respect to the bias, see (9). In the equilibrium it is equal to λ as we show below. Hence, Assumption 1 ensures that the seller does not want to increase the bias infinitely. This condition, $\lambda > c$, is relatively mild. If $\alpha = 5\%$ (i.e., $z_\alpha = 1.64$), then $\lambda \approx 0.54$. Thus, even if half of consumers are credulous, the seller is still constrained by the possible fine.

Assumption 2 $z_\alpha \phi \left(z_\alpha - \frac{1}{z_\alpha} \right) > (1 - c) \lambda (1 - \lambda)$.

Assumption 2 ensures the second-order condition, see the proof of Proposition 1 for details. It is satisfied for $\alpha > 0.75\%$ (i.e., $z_\alpha < 2.436$) and any c. Alternatively, it is satisfied for any α and $c > 0.07$.

Proposition 1 *Under Assumptions 1 and 2, the equilibrium bias and noise are*

$$\beta^* = \frac{\sigma}{\sqrt{N}} \frac{1}{z_\alpha} \sqrt{\frac{1 - \lambda}{\lambda - c}} \quad and \quad \eta^* = \sigma \sqrt{\frac{1 - \lambda}{\lambda - c}} \tag{5}$$

In equilibrium, the noise does not affect the revenues since rational buyers are not misled and pay the prior expected match quality μ while the credulous consumers overpay by β^*. Thus, for any bias, the optimal noise minimizes the costs for that bias. Since the cost function is U-shaped with respect to the noise (Lemma 1) this yields an interior solution $\frac{\beta \sqrt{N}}{\eta} = \frac{1}{z_\alpha}$. The bias increases the revenues by changing the signal distribution in the first-order stochastic dominance sense; the marginal revenue of the bias is the combination of weights of the signal in the rational and credulous buyer's posteriors, $(1 - c) \frac{\sigma^2}{\sigma^2 + \eta^2} + c$. The marginal cost of the bias becomes λ once the equilibrium ratio $\frac{\beta \sqrt{N}}{\eta} = \frac{1}{z_\alpha}$ is plugged in into (9); a closed-form solution is then easily found.

We now turn to the comparative statics of the equilibrium. We say that the equilibrium bias and noise are *complements in the equilibrium* if the sign of comparative statics is (weakly) the same with respect to any parameter.

Corollary 1 *The equilibrium bias and noise (5) are complements in the equilibrium. They are unaffected by the average quality μ and increase with the variance of the prior σ^2, the standard of proof z_α and the share of credulous consumers c. A larger sample N decreases the equilibrium bias but does not affect the noise.*

The complementarity between the noise and the bias is quite unexpected. Indeed, a straightforward intuition from observing the seller's revenues (2) is that whenever

the bias exceeds the buyer's conjecture the seller should minimize the noise. The reason is that this leads the buyer to assign a higher weight to the signal in her posterior. Hence, if, for instance, there is no fine but the bias is exogenously bound from above, the seller selects the highest bias and the minimum possible noise. Similarly, Johnson and Myatt (2006) find that the optimal noise is extreme. In our model, the complementarity comes from the U-shaped effect of the noise on the expected fine and the fact that (rational) buyers are not misled in the equilibrium.

The average quality μ does not affect the bias and the noise. Since the seller sells for any realization of the signal and the buyer's valuation is linear in quality, the seller's expected revenues are additive in μ. His marginal incentives then do not depend on μ.

When the prior is less precise, that is, σ^2 is larger, it becomes more profitable to mislead the consumer because more attention is paid to the signal. Hence, the bias is used more. The costs are unaffected by σ^2 per se. However, as the bias increases, a higher noise should be used to bring the "standardized" bias $\frac{\beta\sqrt{N}}{\eta}$ down to its optimal level $\frac{1}{z_\alpha}$. Also, the noise is proportional to σ^2 since it is the ratio of the two that determines the weights in the buyer's posterior (1). The parameter σ^2 reflects buyers' heterogeneity. Hence, when buyers are more heterogenous, we should expect larger bias and noise.

A stricter (i.e., higher) standard of proof z_α makes it more difficult to convict the seller so the bias becomes cheaper. However, it also increases the cost minimizing noise $\eta = z_\alpha \beta \sqrt{N}$, see Lemma 1(ii). With more noise, the buyers pay less attention to the signal and thus are less easily swayed by bias, making the use of bias less profitable. Overall, we find that this last effect is always dominated by the cost reduction, and as a result, the bias increases with a larger z_α.

The equilibrium bias increases with the share of credulous consumers since marginal returns to bias are higher. Since the cost-minimizing noise is increasing in the bias, the equilibrium noise then also increases.

When the authority increases the sample size N, the seller should counteract this increase in the precision of the bias estimation by increasing the noise (i.e., the cost minimizing noise shifts to the right). A larger noise decreases the marginal revenues of the bias (through a lower weight on the advice in the buyer's updating) and the bias is used less as a result. Since his advice is less biased, the seller also needs less noise to "hide" his misleading practice. It turns out that the noise decreases by

12

enough to compensate the original increase which leads to the quite surprising result that the equilibrium noise does not depend on N.

Going back to our initial example of the retail investment services in Section 1, EC (2011) (pp. 66-67) found that both instruments were used together: "Shoppers were generally provided with the risk level of the proposed investments (i.e. high, mid, or low investment risk). However, the explanation of risk level appeared to be rather vague – most advisors tend to highlight the 'low-risk' nature of the investment, but did not give further explanation to qualify the risk level or had provided descriptions that were too general. Furthermore, while the explanation of risk by the advisor depends to some extent on the feedback he receives from the consumer, there were instances where the advisor's descriptions of risk could be misleading."

4 Optimal Policy

In this Section, we explore how the authority should optimally choose its policy instruments $\{z_\alpha^*, N^*\}$. In order to have a non-trivial trade-off for the authority we introduce honest sellers into the model presented in Section 2. There is a growing evidence that in many markets both honest and "strategic" sellers coexist. For example, Egan, Matvos and Seru (2017), analyzing the universe of financial advisers in the US in 2005-2015, find that "...firms and advisers with clean records coexist with firms and advisers that persistently engage in misconduct" (p. 3). Another case in point is the mortgage market, see Griffin and Maturana (2016) and Piskorski, Seru and Witkin (2015). Moreover, even if checking for past misconduct is possible, many consumers fail to do so. As a result, Egan, Matvos and Seru (2017) find that "some firms specialize in misconduct and attract unsophisticated customers, and others cater to more sophisticated customers, and specialize in honesty..." (p. 4).

Hence, suppose that some sellers, of share h, are honest and do not bias their advice. Nonetheless their advice may be vague because there is some natural noise in communication and reducing it is costly. In other words, as strategic sellers, honest ones also have some minimum noise in the communication. Clearly, honest sellers will not increase the noise above the minimum. Denote it as η_h^i for a honest seller i and $\eta_h \equiv E\left(\eta_h^i\right)$ the average noise of honest sellers.[16] The credulous consumers do

---

[16]We take noise η_h^i as exogenous. Allowing honest sellers to reduce it at some cost would mitigate the negative effects of lower z_α and N on them but would not change the results qualitatively, in particular, Proposition 2 below.

not know whether a seller is honest or not. Rational consumers know it and compute their posterior valuation correctly. The authority has to rely on "hard evidence" to convict the seller and has to use the same statistical test as before.[17] Honest sellers are sometimes fined by mistake (type I error) and, modifying (4b) for the case $\beta = 0$, they are fined in expectation by $C\left(0, \eta_h^i\right) = \frac{\eta_h^i}{\sqrt{N}}\phi\left(z_\alpha\right)$. On average, honest sellers unduly pay $\frac{\eta_h}{\sqrt{N}}\phi\left(z_\alpha\right)$.

We also assume that the authority bears the cost of investigation which is an increasing function of the number of mystery shoppers. Let $I\left(N\right)$ be the investigation cost, with $I' > 0$ and $I'' \geq 0$.

We are now ready to define the loss function that the authority minimizes through the choice of its policy instruments z_α and N:

$$\min_{z_\alpha, N} L = \left(1 - h\right)c\frac{\sigma}{\sqrt{N}}\frac{1}{z_\alpha}\sqrt{\frac{1-\lambda}{\lambda - c}} + h\frac{\eta_h}{\sqrt{N}}\phi\left(z_\alpha\right) + I\left(N\right). \tag{6}$$

The first term corresponds to the loss in consumer welfare $\left(\beta^*\right)$ experienced by credulous consumers misled by strategic sellers. The second term is the fine unduly imposed on honest sellers and the third term is the cost of investigation. When choosing its policy instruments, the authority faces the following trade-offs. First, relaxing the standard of proof (lower z_α^*) reduces the equilibrium bias, and hence the harm on the credulous consumers. However, fines are more likely to be unfairly imposed on honest sellers. The second trade-off concerns the amount of resources devoted to an investigation. On one hand, a larger N^* reduces both the bias and the type I error; but on the other hand, it is more costly.

Proposition 2 *There is an interior optimal policy $\left(z_\alpha^*, N^*\right)$ which solves the authority problem (6).*

The optimal standard of proof z_α^* is higher for a smaller share of credulous consumers c, a lower variance of the prior σ^2, a higher average noise used by honest sellers η_h and a larger share of honest sellers h.

---

[17]Note that the authority does not use a Bayesian test. The test itself does not take into account the prior information about the share of honest sellers nor about the noise used by them. There is a long debate in the literature about the appropriateness of the "naked statistical evidence" (Kaye (1980)) (and the incidental "blue bus case" and "gatecrasher paradox" and exclusionary rules for evidence) and its optimality, see Nesson (1985), Schrag and Scotchmer (1994), Lewis and Poitevin (1997), Posner (1999) and Demougin and Fluet (2008), among others.

If rational consumers do not know the type of the seller, in the equilibrium they debias the signal by $\left(1 - h\right)\beta$ and the dishonest seller misleads them. The noise is used not only to decrease the expected fine but also to increase the revenues. The equilibrium noise is then lower than β^* in (5) but the exact characterization is difficult (see (10a)-(10b) with $\tilde{\beta} = \left(1 - h\right)\beta$).

The optimal size of investigation N^* is higher for a larger c, a higher σ^2 and a higher η_h. There is $\overline{h} \in (0,1)$ such that N^* increases with h for $h < \overline{h}$ and decreases with h for $h > \overline{h}$.

When choosing the standard of proof z_α, the authority trades off the bias of strategic sellers (by which the credulous consumers are misled) against the fine imposed on honest sellers, that is, type I error. A higher standard of proof z_α increases the bias but decreases the fine erroneously imposed on honest sellers. When the variance of the prior σ^2 increases, the bias also increases and hence, the authority adjusts its policy by lowering z_α^*. Therefore, the model predicts a higher buyers' heterogeneity to be associated with a more relaxed standard of proof. A larger share of credulous consumers c increases the bias and the number of consumers misled. As a result, the optimal z_α^* decreases. When the nature of the product is such that conveying precise information is difficult (i.e. η_h is large), the expected fine imposed on honest sellers is large and the authority counteracts this by increasing the standard of proof z_α^*. Finally, a higher share of honest sellers h makes the cost of a type I error relatively more important and, hence, the optimal z_α^* increases.

The optimal size of investigation N^* depends on the loss endured by the authority, which reflects the importance of the case. A higher share of credulous consumers c, the variance of the prior σ^2 and noise used by honest sellers η_h all make the loss larger (for any z_α) and the authority chooses to investigate the case better. Interestingly, the effect of a higher share of honest sellers h is ambiguous. When most sellers are dishonest (low h), the authority cares most about the bias and sets z_α^* to make it low. This means that honest sellers pay high fines on average. Increasing h, that is, replacing one dishonest seller by an honest one, increases the loss of the authority. It then responds by increasing the precision of the test, that is, N^*. When most sellers are honest instead (high h), z_α^* is high so that honest sellers do not pay high fines while dishonest sellers choose a high bias. Replacing one dishonest seller by an honest one now decreases the loss of the authority and the optimal N^* goes down.

The optimal standard of proof z_α^* can be restricted by the following upper bound. Denote by $z_{\max}$ such z_α that $\lambda = c$. The authority will never set the optimal standard of proof above $z_{\max}$ since then the seller will introduce infinite bias (see Assumption 1 and its discussion). Hence, $z_\alpha^* < z_{\max}$. For example, if a quarter of consumers are credulous, then $z_{\max} \approx 2.16$ which corresponds to $\alpha \approx 1.5\%$.

As is clear from (6), the authority does not care about the noise per se. If, for

instance, consumers are risk averse or there is a fixed price as in Section 5, then noise also affects welfare, and the authority should take it into account. The optimal standard of proof will then become lower: from Corollary 1 both equilibrium bias and noise increase with the standard of proof.

Discussion We now briefly discuss other directions in which the analysis of the optimal policy of the authority may be enriched.

We have assumed that the authority cares only about the monetary consequences of its policy: the loss suffered by credulous consumers and the fines imposed on honest sellers. It does not care per se about the fact of punishing honest sellers (type I error) and not punishing strategic sellers (type II error). It is often assumed that the court minimizes a combination of the type I and type II errors (Gennaioli and Shleifer (2007), Alesina and La Ferrara (2014)). In our model, the probability of type I error is $1 - \Phi(z_\alpha)$ (that is, α by definition) and the probability of type II error is equal to $\Phi\left(z_\alpha - \frac{1}{z_\alpha}\right)$ in equilibrium.[18] Introducing these two terms into the loss function (6) yields

$$L = (1-h)\left[c\frac{\sigma}{\sqrt{N}}\frac{1}{z_\alpha}\sqrt{\frac{1-\lambda}{\lambda-c}} + \Phi\left(z_\alpha - \frac{1}{z_\alpha}\right)\right] + h\left[\frac{\eta_h}{\sqrt{N}}\phi(z_\alpha) + 1 - \Phi(z_\alpha)\right] + I(N).$$

A larger z_α increases the type II error and decreases the type I error. Thus, it increases the part of the loss coming from strategic sellers and decreases the one coming from honest sellers, as before.

Another important instrument that the authority might consider is the size of the damages. We have assumed that the fine imposed on a seller who has been found guilty is equal to the estimated bias $\widehat{\beta}$. Suppose now that it is multiplied by c, the share of credulous consumers which are harmed in the equilibrium, and parameter d. If $d > 1$, this is the case of punitive (or exemplary) damages often used in common law countries. In the US, both the frequency and the magnitude of punitive damages verdicts have increased dramatically in recent years (Sunstein et al., 2002). If $d < 1$, it may reflect the fact that some cases are not followed due to the limited resources of the authority or are terminated because of procedural errors. The costs in (4) are then multiplied by dc, so equilibrium bias and noise (5) become $\beta_d^* = \frac{\sigma}{\sqrt{N}}\frac{1}{z_\alpha}\sqrt{\frac{1-cd\lambda}{c(d\lambda-1)}}$

---

[18]A guilty seller is acquitted when $\frac{\widehat{\beta}}{\eta/\sqrt{N}} < z_\alpha$, where $\widehat{\beta} \rightsquigarrow \mathcal{N}\left(\beta, \frac{\eta^2}{N}\right)$. In equilibrium, this happens with probability $\Phi\left(z_\alpha - \frac{1}{z_\alpha}\right)$.

and $\eta_d^* = \sigma\sqrt{\frac{1-cd\lambda}{c(d\lambda-1)}}$. As expected, a higher d leads to the lower equilibrium bias (and noise). However, a higher d also increases the fine imposed on honest sellers and, hence, the authority trades off the deterrence effect of higher fines against larger erroneous fines endured by honest sellers.

Finally, the authority may want to expend resources to educate consumers, thereby reducing the share of credulous consumers c. For example, in several countries government bodies promote financial literacy such as The Financial Literacy and Education Commission in the US, The Money Advice Service in the UK and The Australian Securities and Investments Commission in Australia. When the share of strategic sellers (lower h) or the variance of the prior σ are larger, the loss in consumer surplus by credulous consumers (first term in (6)) increases and the authority will then spend more resources on educating them.

5   Fixed price

According to OECD (2010), there are two types of consumer detriment: "Consumers will incur a loss in economic welfare if they are misled into making purchases of goods and services which they would not otherwise have made or if they pay more for the purchases than they would if they had been better informed" (p. 52). We addressed the second welfare loss in the main model in Section 2.

In this section, we address the first consumer detriment in a scenario with a fixed price and, for the ease of exposition, all buyers being rational. In many instances the price of a product is fixed at the selling stage and the buyer decides whether to buy it or not at that given price. For example, the seller may have committed to the price by publicly advertising it. Alternatively, the price may be fixed by the marketing division that operates independently from the selling division. Regulated prices are another good illustration. For instance, Brown and Minor (2014) study misconduct at the selling stage by life insurance agents who "... cannot adjust the prices faced by individual customers - this practice, called 'rebating', is illegal in most jurisdictions" (p. 5). Assume that price p is exogenously fixed at a pre-advice stage. The buyer decides to buy if her posterior valuation of the product (1) is above p. The seller's revenues are then

$$R\left(\beta,\eta\right) = p\Pr\left\{\frac{\left(s-\tilde{\beta}\right)\sigma^2 + \mu\eta^2}{\sigma^2+\eta^2} \geq p\right\} = p\left(1 - G\left(p+\tilde{\beta}+\left(p-\mu\right)\frac{\eta^2}{\sigma^2}\right)\right) \quad (7)$$

The seller's costs are unchanged and given by (4).[19]

Consider first the case where the price is exogenously fixed at the mean level μ.

Proposition 3 *When $p = \mu > 0$, the equilibrium bias and noise are*[20]

$$\beta_p^* = \frac{\sqrt{\max\left\{\frac{\mu^2}{2\pi\lambda^2} - \sigma^2, 0\right\}}}{z_\alpha \sqrt{N}} \quad and \quad \eta_p^* = \sqrt{\max\left\{\frac{\mu^2}{2\pi\lambda^2} - \sigma^2, 0\right\}}. \tag{8}$$

The bias-to-noise ratio, $\frac{\beta_p^*}{\eta_p^*}$, is the same as in the price discrimination setting and equals to $\frac{1}{z_\alpha \sqrt{N}}$.

Since the price is fixed, the goal of the seller is to increase the probability of selling, which means convincing the marginal buyer (the one whose posterior is equal to the price). The bias unambiguously increases the probability of selling and its marginal effect on the revenues is the price times the probability of encountering the marginal buyer $g(\mu + \beta) = \frac{1}{\sqrt{\sigma^2 + \eta^2}}\phi(0)$. The noise decreases the weight of the signal in the buyer's posterior valuation. It increases the probability of selling if the buyer overestimates the bias and decreases it otherwise. In the equilibrium, however, the noise does not affect the seller's revenues since the marginal buyer is the one with valuation equal to the prior mean μ and she correctly anticipates the bias; in other words, the probability of selling is equal to $\frac{1}{2}$. Thus, as in the model of Section 2, the noise can be seen as minimizing the costs for any given bias which pins down the bias-to-noise ratio, $\frac{\beta\sqrt{N}}{\eta} = \frac{1}{z_\alpha}$. A closed-form solution can then easily be found.

The comparative statics results are summarized in the following corollary.

Corollary 2 *When $p = \mu > 0$, the equilibrium bias and noise (8) are complements in the equilibrium. They increase with the average quality μ and the standard of proof z_α and decrease with the variance of the prior σ^2.[21] A larger sample N decreases the equilibrium bias but does not affect the noise.*

---

[19] Alternatively, the fine could be proportional to the actual harm to consumers which is the difference between the price and the true match quality for those who bought the product due to the bias. However, since this harm is monotonic in the bias, such a cost function is still increasing in bias and U-shaped in noise.

[20] A sufficient (but not necessary) condition for the second order conditions to hold is $z_\alpha \geq 1.09$ (i.e., $\alpha \leq 13.79\%$).

[21] A sufficient (but not necessary) condition for $\frac{\partial \beta_p^*}{\partial z_\alpha} > 0$ is $z_\alpha \geq 1.29$ (i.e., α is less than approximately 10%). See the proof for the exact condition.

Unlike the model of Section 2, the average quality now affects the bias and noise. A higher μ increases both the posterior expected quality and the price, so the probability of selling remains unchanged. However, an increase in μ also increases the revenue per unit sold through the increase in p. The increase in the profitability of the sale induces the seller to use more bias, as this increases the probability of selling. Following the increase in the bias, the noise also increases to minimize the costs of the larger bias.

The comparative statics with respect to the variance of the prior, σ^2, have the opposite sign from that in Section 3 (see Corollary 1). The intuition is the following: while σ^2 still enters the buyer's posterior valuation, it does not affect her decision to buy. Indeed, since the price is equal to the mean valuation, the buyer buys if the signal exceeds the mean (after debiasing) and does not buy otherwise. Thus, the only effect of a higher σ^2 is to decrease the probability of the marginal buyer. Then, the marginal revenues of the bias are smaller, and the seller uses it less, and correspondingly he introduces a lower noise. In Johnson and Myatt (2006) the relationship between the buyers' heterogeneity and the informativeness of the signal is the same: the benefits of giving precise information are higher if the buyers differ largely in their tastes and, therefore, more idiosyncratic products are complemented by detailed advertising and marketing activities.

When the standard of proof, z_α, increases, the seller can afford to use greater noise as the probability of ending up convicted and paying a large fine decreases. A stricter standard of proof also decreases the marginal cost of biasing the advice. However, as in Section 3, greater noise also decreases the marginal revenue of the bias although through a different mechanism. Similarly to σ^2, greater noise decreases the probability of the marginal consumer and hence makes biasing the advice less profitable. However, if z_α is large enough, the decrease of the marginal revenues is smaller than the one of the marginal costs and the bias increases.

The intuition for the comparative statics of the sample size N is the same as in Section 3. However, the mechanism behind the decrease in the marginal revenues of the bias following an increase in the noise differs. Now it is due to a decrease in the probability of the marginal buyer rather than a decrease in the attention paid to the advice.

Finally, if μ is high enough, then both the equilibrium bias and the equilibrium noise are higher than those determined in Section 3. A higher $\mu\,(=p)$ increases the marginal benefit of biasing the advice while it is constant in the price discrimination

setting.

Consider now the case $p \neq \mu$. While we cannot obtain closed-form solutions, we still can shed light on the bias-to-noise ratio in the equilibrium.

Proposition 4 *As compared to the level when the price is equal to the average valuation μ, the equilibrium bias-to-noise ratio is higher (lower) when the price is higher (lower).*

When the price is different from μ, the noise changes the identity of the marginal buyer since the posterior shifts towards the prior mean μ. In particular, a higher noise brings the buyer's posterior valuation closer to the mean which decreases the seller's revenues when $p > \mu$ since the marginal consumer has the valuation above the mean and increases them when $p < \mu$ since the marginal consumer has the valuation below the mean. As a result, when $p > \mu$ the bias-to-noise ratio will be larger than the one in Section 3 and smaller when $p < \mu$. Actually, the analysis of the effects of the noise on the seller's revenues becomes similar to the one of Johnson and Myatt (2006). The case $p > \mu$ corresponds to their "niche market", where the seller only serves the high valuation consumers and provides them with a very precise information, and the case $p < \mu$ corresponds to their "mass market" where the seller serves most consumers and gives them little information. However, contrary to their setup where noise is costless, in our model the noise also affects the seller's costs and, therefore, extreme noise is not optimal.

6 Conclusion

In this paper, we investigated a seller's incentives to provide (un)informative and (un)biased advice, the resulting equilibrium communication and the optimal policy of the authority in charge of protecting consumer welfare in this market. We found that biasing the advice and making it more noisy are complements: the seller tells either an exact truth or a vague lie. For example, a higher standard of proof employed by the authority and a higher share of credulous consumers make the advice given by the seller more biased and less precise. A higher buyers' heterogeneity has the same effect in the price discrimination setup and the opposite one when the price is fixed at the mean valuation.

We then analyzed the optimal policy of the authority. A higher standard of proof makes the fine less likely and, therefore, the equilibrium bias increases, leading

to a higher welfare loss for credulous consumers. On the other hand, it reduces the type I error, that is, the probability that an honest seller is found guilty. The optimal standard of proof trades off these two effects and is higher if there are more honest sellers, if they communicate with a higher noise, when there are less credulous consumers and if the demand is more homogenous. A higher quality of investigation, proxied by the size of the sample used by the authority, decreases both the loss of credulous consumers and the type I error but is costly. It is higher when the case is more important, in particular, when there are more credulous consumers, the demand is more heterogenous or honest sellers communicate more noisily.

In the previous version Drugov and Troya-Martinez (2012) we considered a number of extensions of the model of Section 2. We looked at the case of the unobserved noise which may come from the fine print that is typically attached to complex contracts. As found by OFT (2011), these clauses are rarely read in the moment of the purchase and they can be more or less precise. Thus, at the moment of observing the signal the buyer cannot assess how precise it is. The persuasion literature only considers the case where the buyer (or more generally, the receiver) knows the precision of the information, the only exception being Troya-Martinez (2016). While the analysis becomes technically very hard, qualitatively Proposition 1 still holds. We also looked at the case of informational externalities between two sellers selling related products. There is a disciplining effect, and equilibrium bias and noise - they are still complementary - decrease. Surprisingly, if there are also common shocks in the communication technology (that is, the error terms are correlated), the disciplining effect of the second seller can be reversed.

Finally, contributing to the literature on law and economics, we provided a microfounded model of the signal process generated by an equilibrium continuous harmful action with endogenous informativeness which allows a detailed study of different instruments available to courts. We hope that this framework will be used in the future to enrich existing models of litigation.

Appendix

Proof of Lemma 1. The first derivative of (4b) with respect to β is

$$\frac{\partial C(\beta, \eta)}{\partial \beta} = 1 - \Phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right) + z_\alpha \phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right) \tag{9}$$

and it is always positive. The first derivative of (4b) with respect to η is

$$\frac{\partial C\left(\beta,\eta\right)}{\partial \eta} = \frac{1}{\sqrt{N}}\phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right)\left[1 - \frac{\beta\sqrt{N}}{\eta}z_\alpha\right]$$

and it is negative for $\eta < \beta\sqrt{N}z_\alpha$ and positive for $\eta > \beta\sqrt{N}z_\alpha$. ∎

Proof of Proposition 1. The first-order conditions of the seller's problem with respect to β and η, respectively, are

$$(1-c)\frac{\sigma^2}{\sigma^2 + \eta^2} + c - \left(1 - \Phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right) + z_\alpha\phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right)\right) = 0 \quad (10a)$$

$$-(1-c)\left(\beta - \tilde{\beta}\right)\frac{2\eta\sigma^2}{(\sigma^2+\eta^2)^2} - \frac{1}{\sqrt{N}}\phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right)\left[1 - \frac{\beta\sqrt{N}}{\eta}z_\alpha\right] = 0 \quad (10b)$$

In the equilibrium $\beta = \tilde{\beta}$ and from (10b) $\frac{\beta\sqrt{N}}{\eta} = \frac{1}{z_\alpha}$. Plug this into (10a) to obtain η^* and then β^*.

To check the second-order conditions, differentiate (10a) with respect to β to obtain $\frac{\partial^2\Pi}{\partial\beta^2}$ and (10b) with respect to η and β to obtain $\frac{\partial^2\Pi}{\partial\eta^2}$ and $\frac{\partial^2\Pi}{\partial\eta\partial\beta}$, respectively (denote $\phi_C = \phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right)$)

$$\frac{\partial^2\Pi}{\partial\beta^2} = -\frac{\sqrt{N}}{\eta}\phi_C\left(1 - \frac{\beta\sqrt{N}}{\eta}z_\alpha + z_\alpha^2\right)$$

$$\frac{\partial^2\Pi}{\partial\eta^2} = -(1-c)\left(\beta - \tilde{\beta}\right)2\sigma^2\frac{\sigma^2 - 3\eta^2}{(\sigma^2+\eta^2)^3} - \frac{\beta^2\sqrt{N}}{\eta^3}\phi_C\left(1 - \frac{\beta\sqrt{N}}{\eta}z_\alpha + z_\alpha^2\right)$$

$$\frac{\partial^2\Pi}{\partial\eta\partial\beta} = -(1-c)\frac{2\eta\sigma^2}{(\sigma^2+\eta^2)^2} + \frac{\beta\sqrt{N}}{\eta^2}\phi_C\left(1 - \frac{\beta\sqrt{N}}{\eta}z_\alpha + z_\alpha^2\right)$$

In the equilibrium $\beta = \beta^* = \tilde{\beta}$ and $\frac{\beta^*\sqrt{N}}{\eta^*} = \frac{1}{z_\alpha}$ so these derivatives become ($\phi_C^* = \phi\left(z_\alpha - \frac{1}{z_\alpha}\right)$)

$$\frac{\partial^2\Pi}{\partial\beta^2} = -\frac{\sqrt{N}}{\eta^*}\phi_C^* z_\alpha^2 < 0$$

$$\frac{\partial^2\Pi}{\partial\eta^2} = -\frac{\beta^{*2}\sqrt{N}}{\eta^{*3}}\phi_C^* z_\alpha^2 < 0$$

$$\frac{\partial^2\Pi}{\partial\eta\partial\beta} = -(1-c)\frac{2\eta^*\sigma^2}{(\sigma^2+\eta^{*2})^2} + \frac{\beta^*\sqrt{N}}{\eta^{*2}}\phi_C^* z_\alpha^2$$

Check that the determinant of the Hessian is positive:

$$
\begin{aligned}
\frac{\partial^2 \Pi}{\partial \beta^2}\frac{\partial^2 \Pi}{\partial \eta^2} - \left(\frac{\partial^2 \Pi}{\partial \eta \partial \beta}\right)^2 \;=\;& \frac{\beta^{*2} N}{\eta^{*4}}\phi_C^{*2} z_\alpha^4 - \left((1-c)\frac{2\eta^* \sigma^2}{(\sigma^2+\eta^{*2})^2}\right)^2 \\
& + (1-c)\frac{4\eta^* \sigma^2}{(\sigma^2+\eta^{*2})^2}\frac{\beta^* \sqrt{N}}{\eta^2}\phi_C^* z_\alpha^2 - \frac{\beta^{*2} N}{\eta^{*4}}\phi_C^{*2} z_\alpha^4 \\
=\;& (1-c)\frac{4\sigma^2}{(\sigma^2+\eta^{*2})^2}\left(\frac{\beta^*\sqrt{N}}{\eta^{*2}}\phi_C^* z_\alpha^2 - (1-c)\frac{\eta^*\sigma^2}{(\sigma^2+\eta^{*2})^2}\right) \\
=\;& (1-c)\frac{4\sigma^2}{\eta^*(\sigma^2+\eta^{*2})^2}\left(\phi_C^* z_\alpha - (1-c)\lambda(1-\lambda)\right)
\end{aligned}
\tag{11}
$$

The curve $\phi_C^* z_\alpha - (1-c)\lambda(1-\lambda) = 0$ has a negative slope in (z_α, c) space. Hence, a lower z_α or a higher c help to make the determinant positive. In particular, it is positive if $z_\alpha < 2.436$, i.e., $\alpha > 0.75\%$, and any c and it is positive if $c > 0.07$ and any z_α.

Finally, note that these are the second-order conditions at the extremum. In our candidate equilibrium, i.e., if the seller's bias and buyer's conjecture is β^* and the noise is η^*, the seller earns $\mu + c\beta^*$ minus the costs. Potentially, when c is small enough he might prefer to deviate in the following way: stop the communication by introducing infinite noise, in which case the rational buyers disregard the signal and still pay μ on average, and introduce a negative bias in order to decrease the costs. However, the assumption on the provision of some minimal information, that is, an upper bound on noise, rules out this deviation. ■

Proof of Proposition 2. **Interior optimal policy.** Let us first show that the optimal policy optimal policy (z_α^*, N^*) is interior.

Start with z_α^*. Since $\lambda(0) = 1$, $\lambda(+\infty) = 0$ and $\lambda(z)$ is strictly decreasing in z, there is a unique solution to $\lambda(z) = c$; denote it $z_{\max}$. We first show that the optimal standard of proof $z_\alpha^* \in (0, z_{\max})$. Take the first derivative of L with respect to z_α:

$$
\frac{\partial L}{\partial z_\alpha} = \frac{(1-h)c\sigma}{\sqrt{N}}\frac{1}{\sqrt{(1-\lambda)(\lambda - c)}}\left[\frac{z_\alpha}{2}\phi(z_\alpha - \frac{1}{z_\alpha})\frac{1-c}{\lambda - c} - \frac{1-\lambda}{z_\alpha^2}\right] - h\frac{\eta_h}{\sqrt{N}}z_\alpha \phi(z_\alpha)
\tag{12}
$$

When $z_\alpha \to z_{\max}$, $\frac{1}{\lambda - c} \to +\infty$ and, therefore, optimal $z_\alpha^* < z_{\max}$.

The case $z_\alpha \to 0$ is more involved. We need the facts that $\frac{\partial \phi(z_\alpha - \frac{1}{z_\alpha})}{\partial z_\alpha} = -z_\alpha\left(1 - \frac{1}{z_\alpha^4}\right)\phi(z_\alpha - \frac{1}{z_\alpha})$ and $\frac{\partial \lambda}{\partial z_\alpha} = -z_\alpha^2 \phi(z_\alpha - \frac{1}{z_\alpha})$ and the following lemma with its corollary.

Lemma 2 $\lim\limits_{z_\alpha \to 0} \frac{\phi(z_\alpha - \frac{1}{z_\alpha})}{z_\alpha^x} = 0.$

Proof. $\phi(z_\alpha - \frac{1}{z_\alpha}) = \frac{1}{\sqrt{2\pi}}e^{-\frac{1}{2}\left(z_\alpha - \frac{1}{z_\alpha}\right)^2}$. Thus, $\lim\limits_{z_\alpha \to 0}\phi(z_\alpha - \frac{1}{z_\alpha}) = 0$ and the result follows for $x \leq 0$.

Consider now $x > 0$. Substituting $y = \frac{1}{z_\alpha}$ yields

$$\lim_{z_\alpha \to 0} \frac{\phi(z_\alpha - \frac{1}{z_\alpha})}{z_\alpha^x} = \lim_{y \to +\infty} y^x \phi(\frac{1}{y} - y) = \frac{1}{\sqrt{2\pi}} \lim_{y \to +\infty} \frac{y^x}{e^{\frac{1}{2}\left(\frac{1}{y} - y\right)^2}}.$$

Use l'Hôpital's rule to transform the last expression into

$$\frac{1}{\sqrt{2\pi}} \lim_{y \to +\infty} \frac{xy^{x-1}}{y\left(1 - \frac{1}{y^4}\right) e^{\frac{1}{2}\left(\frac{1}{y} - y\right)^2}} = \frac{1}{\sqrt{2\pi}} x \lim_{y \to +\infty} \frac{1}{\left(1 - \frac{1}{y^4}\right)} \lim_{y \to +\infty} \frac{y^{x-2}}{e^{\frac{1}{2}\left(\frac{1}{y} - y\right)^2}}$$

$$= \frac{1}{\sqrt{2\pi}} x \lim_{y \to +\infty} \frac{y^{x-2}}{e^{\frac{1}{2}\left(\frac{1}{y} - y\right)^2}}$$

Repeating l'Hôpital's rule until the numerator inside the limit is y to a negative power leads to the result. ∎

Corollary 3 $\lim\limits_{z_\alpha \to 0} \frac{z_\alpha^x \phi^n(z_\alpha - \frac{1}{z_\alpha})}{1 - \lambda} = \lim\limits_{z_\alpha \to 0} \frac{1 - \lambda}{z_\alpha^x} = 0$ *for any x and $n > 1$.*

Proof. Using l'Hôpital's rule for both limits

$$\lim_{z_\alpha \to 0} \frac{z_\alpha^x \phi^n\left(z_\alpha - \frac{1}{z_\alpha}\right)}{1 - \lambda} = \lim_{z_\alpha \to 0} \frac{xz_\alpha^{x-1} \phi^n\left(z_\alpha - \frac{1}{z_\alpha}\right) - z_\alpha^x n z_\alpha \left(1 - \frac{1}{z_\alpha^4}\right) \phi^n\left(z_\alpha - \frac{1}{z_\alpha}\right)}{z_\alpha^2 \phi\left(z_\alpha - \frac{1}{z_\alpha}\right)}$$

$$= \lim_{z_\alpha \to 0} z_\alpha^{x-3} \left(x - n z_\alpha^2 \left(1 - \frac{1}{z_\alpha^4}\right)\right) \phi^{n-1}\left(z_\alpha - \frac{1}{z_\alpha}\right) = 0$$

and

$$\lim_{z_\alpha \to 0} \frac{1 - \lambda}{z_\alpha^x} = \lim_{z_\alpha \to 0} \frac{z_\alpha^2 \phi\left(z_\alpha - \frac{1}{z_\alpha}\right)}{x z_\alpha^{x-1}} = 0,$$

where the last equalities in each case follow from Lemma 2. ∎

The first term in (12) is zero since $\lim\limits_{z_\alpha \to 0} \frac{z_\alpha \phi(z_\alpha - \frac{1}{z_\alpha})}{\sqrt{1 - \lambda}} = \lim\limits_{z_\alpha \to 0} \frac{\sqrt{1 - \lambda}}{z_\alpha^2} = 0$ by Corollary 3. Hence, $\lim\limits_{z_\alpha \to 0} \frac{\partial L}{\partial z_\alpha} = 0$. Then, compute the second derivative

$$\frac{\partial^2 L}{\partial z_\alpha^2} = \frac{(1-h)c\sigma}{\sqrt{N}} \left(\frac{1-c}{2} \frac{\phi(z_\alpha - \frac{1}{z_\alpha})}{\sqrt{(1-\lambda)(\lambda - c)^3}} \left(1 - z_\alpha^2 + \frac{1}{z_\alpha^2} - \frac{z_\alpha^3(4\lambda - 3 - c)}{2(1-\lambda)(\lambda - c)} \phi(z_\alpha - \frac{1}{z_\alpha})\right) \right.$$
$$\left. + \frac{2}{z_\alpha^3} \sqrt{\frac{1-\lambda}{\lambda - c}} - (1-c) \frac{\phi(z_\alpha - \frac{1}{z_\alpha})}{2\sqrt{(1-\lambda)(\lambda - c)^3}} \right) \tag{13}$$
$$- h \frac{\eta_h}{\sqrt{N}} \phi(z_\alpha) (1 - z_\alpha^2)$$

Corollary 3 applies to all the terms in the first two lines of (13). Hence, $\lim\limits_{z_\alpha \to 0} \frac{\partial^2 L}{\partial z_\alpha^2} = -h \frac{\eta_h}{\sqrt{N}} \phi(0) < 0$. Since $\lim\limits_{z_\alpha \to 0} \frac{\partial L}{\partial z_\alpha} = 0$, $\frac{\partial L}{\partial z_\alpha} < 0$ in a neighborhood of zero and $z_\alpha^* > 0$.

Now, we show that N^* is interior. Denote by L^*_{-N} the value of the two first terms in the loss function (6) at z_α^* multiplied by $\sqrt{N}$ (note that z_α^* does not depend on

N since it cancels out in (12)). The first-order condition with respect to N can be written as

$$\frac{1}{2} L^*_{-N} = I'(N) N^{\frac{3}{2}}.$$

The left-hand side is positive and constant (in terms of N), while the right hand side is strictly increasing from zero to infinity since $I'(N) > 0$ and $I'' \geq 0$.[22] Thus, there is a unique and interior solution N^*.

Comparative statics of z^*_α. Taking the total derivative of the first-order condition $\frac{\partial L}{\partial z_\alpha} = 0$ at (z^*_α, N^*) with respect to some parameter p, we get

$$\frac{\partial^2 L}{\partial z_\alpha \partial p} + \frac{\partial^2 L}{\partial z_\alpha^2} \frac{\partial z^*_\alpha}{\partial p} + \frac{\partial^2 L}{\partial z_\alpha \partial N} \frac{\partial N^*}{\partial p} = \frac{\partial^2 L}{\partial z_\alpha \partial p} + \frac{\partial^2 L}{\partial z_\alpha^2} \frac{\partial z^*_\alpha}{\partial p} = 0,$$

where the first equality follows from $\frac{\partial^2 L}{\partial z_\alpha \partial N} = 0$ at z^*_α since $\frac{\partial L}{\partial z_\alpha}$ is multiplicative in N, see (12).

Then, $\frac{\partial z^*_\alpha}{\partial p}$ has the opposite sign of $\frac{\partial^2 L}{\partial z_\alpha \partial p}$ since $\frac{\partial^2 L}{\partial z_\alpha^2}$ is positive.

Using (12), it is easy to see that $\frac{\partial^2 L}{\partial z_\alpha \partial \sigma} > 0$, $\frac{\partial^2 L}{\partial z_\alpha \partial \eta_h} < 0$ and $\frac{\partial^2 L}{\partial z_\alpha \partial h} < 0$. For the effect of c, note that both $\frac{c}{\sqrt{1-\lambda}}$ and $\frac{1-c}{\lambda - c}$ increase in c and, hence, $\frac{\partial^2 L}{\partial z_\alpha \partial c} > 0$.

Comparative statics of N^*. Analogously to the case of z^*_α, $\frac{\partial N^*}{\partial p}$ has the opposite sign of $\frac{\partial^2 L}{\partial N \partial p}$. The sign of $\frac{\partial^2 L}{\partial N \partial p}$ is opposite to the one of $\frac{\partial L^*_{-N}}{\partial p}$. It is easy to see that $\frac{\partial L^*_{-N}}{\partial c} > 0$, $\frac{\partial L^*_{-N}}{\partial \sigma} > 0$ and $\frac{\partial L^*_{-N}}{\partial \eta_h} > 0$. The effect of h is given by

$$\frac{\partial L^*_{-N}}{\partial h} = -c\sigma \frac{1}{z_\alpha} \sqrt{\frac{1-\lambda}{\lambda - c}} + \eta_h \phi(z_\alpha).$$

When h is low enough, $\frac{\partial L^*_{-N}}{\partial h} > 0$ since the authority cares most about decreasing the bias and sets low z_α letting the type I error to be high. When h is high enough the opposite happens and $\frac{\partial L^*_{-N}}{\partial h} < 0$. Let us take the second derivative

$$\frac{\partial^2 L^*_{-N}}{\partial h^2} = \left(-\frac{\partial}{\partial z_\alpha} \left(c\sigma \frac{1}{z_\alpha} \sqrt{\frac{1-\lambda}{\lambda - c}} \right) + \frac{\partial}{\partial z_\alpha} \eta_h \phi(z_\alpha) \right) \frac{\partial z^*_\alpha}{\partial h}.$$

Both terms in brackets are negative while $\frac{\partial z^*_\alpha}{\partial h}$ is positive. L^*_{-N} is then concave in h and reaches the maximum at some $\overline{h} \in (0,1)$, so $\frac{\partial L^*_{-N}}{\partial h} \gtrless 0$ when $h \lessgtr \overline{h}$. ∎

Proof of Proposition 3. The first-order conditions of the seller's problem to maximize revenues (7) minus costs (4) with respect to β and η, respectively, are

---

[22]Condition $I'' \geq 0$ is sufficient but not necessary. What we really need is that $I'(N) N^{\frac{3}{2}}$ is increasing, that is, $I(N)$ can be concave but not "too much".

(denote $\Phi_C = \Phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right)$ and $\phi_C = \phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right)$)

$$\mu g\left(\mu + \tilde{\beta}\right) - (1 - \Phi_C + z_\alpha \phi_C) = 0 \qquad (14a)$$

$$-\mu\eta g\left(\mu + \tilde{\beta}\right)\frac{\beta - \tilde{\beta}}{\sigma^2 + \eta^2} - \frac{1}{\sqrt{N}}\phi_C\left(1 - \frac{\beta\sqrt{N}}{\eta}z_\alpha\right) = 0 \qquad (14b)$$

In the equilibrium, $\beta = \tilde{\beta}$. From (14b) $\frac{\beta\sqrt{N}}{\eta} = \frac{1}{z_\alpha}$ as in Section 3. Plug this into (14a) to obtain η_p^* and then β_p^*.

To check the second-order conditions, we differentiate (14a) with respect to β and η to obtain $\frac{\partial^2 \Pi}{\partial \beta^2}$ and $\frac{\partial^2 \Pi}{\partial \eta \partial \beta}$, respectively and (14b) with respect to η to obtain $\frac{\partial^2 \Pi}{\partial \eta^2}$:

$$\frac{\partial^2 \Pi}{\partial \beta^2} = \mu g(\mu + \tilde{\beta})\frac{\tilde{\beta} - \beta}{\sigma^2 + \eta^2} - \frac{\sqrt{N}}{\eta}\phi_C\left(1 + z_\alpha\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right)\right)$$

$$\frac{\partial^2 \Pi}{\partial \eta^2} = -\mu g(\mu + \tilde{\beta})\frac{\beta - \tilde{\beta}}{\sigma^2 + \eta^2}\left[1 + \eta^2\left(\frac{\left(\tilde{\beta} - \beta\right)^2 - 3\left(\sigma^2 + \eta^2\right)}{\left(\sigma^2 + \eta^2\right)^2}\right)\right]$$
$$-\frac{\beta^2}{\eta^2}\frac{\sqrt{N}}{\eta}\phi_C\left[1 + z_\alpha\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right)\right]$$

$$\frac{\partial^2 \Pi}{\partial \eta \partial \beta} = \mu\eta g(\mu + \tilde{\beta})\frac{\left(\tilde{\beta} - \beta\right)^2 - \left(\sigma^2 + \eta^2\right)}{\left(\sigma^2 + \eta^2\right)^2} + \frac{\beta\sqrt{N}}{\eta^2}\phi_C\left(1 + z_\alpha\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right)\right)$$

In equilibrium, $\beta = \beta_p^* = \tilde{\beta}$ and $\frac{\beta_p^*\sqrt{N}}{\eta_p^*} = \frac{1}{z_\alpha}$ so these derivatives become ($\phi_C^* = \phi\left(z_\alpha - \frac{1}{z_\alpha}\right)$)

$$\frac{\partial^2 \Pi}{\partial \beta^2} = -\frac{z_\alpha^2 \sqrt{N}}{\eta_p^*}\phi_C^* < 0$$

$$\frac{\partial^2 \Pi}{\partial \eta^2} = -\frac{z_\alpha \beta_p^*}{\eta_p^{*2}}\phi_C^* < 0$$

$$\frac{\partial^2 \Pi}{\partial \eta \partial \beta} = -\mu\eta_p^* g(\mu + \beta_p^*)\frac{1}{\sigma^2 + \eta_p^{*2}} + \frac{z_\alpha}{\eta_p^*}\phi_C^*$$

The determinant of the Hessian, $\frac{\partial^2 \Pi}{\partial \beta^2}\frac{\partial^2 \Pi}{\partial \eta^2} - \left(\frac{\partial^2 \Pi}{\partial \eta \partial \beta}\right)^2$, is equal to

$$
\frac{z_\alpha^2 \sqrt{N}}{\eta_p^*}\frac{z_\alpha \beta_p^*}{\eta_p^{*2}}\phi_C^{*2} - \left(\frac{z_\alpha}{\eta_p^*}\phi_C^* - \mu \eta_p^* g(\mu + \beta_p^*)\frac{1}{\sigma^2 + \eta_p^{*2}}\right)^2
$$

$$
= 2\frac{\mu g(\mu + \beta_p^*)}{\sigma^2 + \eta_p^{*2}}z_\alpha \phi_C^* - \left(\frac{\mu \eta_p^* g(\mu + \beta_p^*)}{\sigma^2 + \eta_p^{*2}}\right)^2
$$

$$
= \frac{\mu g(\mu + \beta_p^*)}{\sigma^2 + \eta_p^{*2}}\left[2z_\alpha \phi_C^* - \frac{\eta_p^{*2}\lambda}{\sigma^2 + \eta_p^{*2}}\right]
$$

$$
= \frac{g(\mu + \beta_p^*)}{\mu\left(\sigma^2 + \eta_p^{*2}\right)}\left(2\pi \lambda^3 \sigma^2 - \mu^2\left[1 - \Phi_C^* - z_\alpha \phi_C^*\right]\right)
$$

In Figure 1 we plot the parameter range for which it is positive and, therefore, the second-order conditions are satisfied. In particular, they are satisfied if $z_\alpha \geq 1.09$ (i.e., $\alpha \leq 13.8\%$), in which case $1 - \Phi_C^* - z_\alpha \phi_C^* < 0$. ∎

Proof of Corollary 2. Consider the effect of z_α. Since $\frac{\partial \lambda}{\partial z_\alpha} = -z_\alpha^2 \phi_C^* < 0$, $\frac{\partial \eta_p^*}{\partial z_\alpha} > 0$. The effect on the equilibrium bias is more involved:

$$
\frac{\partial \beta_p^*}{\partial z_\alpha} = \frac{\sqrt{\frac{\mu^2}{2\pi\lambda^2} - \sigma^2}}{z_\alpha^2 \sqrt{N}}\left[\frac{\mu^2 z_\alpha^3 \phi_C^*}{\lambda\left(\mu^2 - 2\pi\lambda^2\sigma^2\right)} - 1\right].
$$

Since $\frac{\mu^2}{\sigma^2 2\pi} > \lambda^2$ for the bias to be positive, $\frac{\partial \beta}{\partial z_\alpha}$ is positive if:

$$
2\pi\lambda^3\sigma^2 > \mu^2\left[\lambda - z_\alpha^3 \phi_C^*\right].
$$

In Figure 1 we plot the parameter range for which this inequality holds. In particular, it holds if $z_\alpha \geq 1.29$ (i.e., $\alpha \leq 9.9\%$) in which case $\lambda - z_\alpha^3 \phi_C^* < 0$.

The remaining comparative statics are straightforward. ∎

Proof of Proposition 4. The first-order conditions of the seller's problem with respect to β and η, respectively, are

$$
pg\left(\bar{s}\right) - \left(1 - \Phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right) + z_\alpha \phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right)\right) = 0 \quad (16a)
$$

$$
-p\eta g(\bar{s})\left[\frac{\beta - \tilde{\beta}}{\sigma^2 + \eta^2} + \frac{p - \mu}{\sigma^2}\right] - \frac{1}{\sqrt{N}}\phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right)\left[1 - \frac{\beta\sqrt{N}}{\eta}z_\alpha\right] = 0 \quad (16b)
$$

where $\bar{s} = p + \tilde{\beta} + (p - \mu)\frac{\eta^2}{\sigma^2}$. In the equilibrium, $\beta = \tilde{\beta}$. If $p > (<)\mu$ then $\frac{1}{\sqrt{N}z_\alpha} < (>)\frac{\beta}{\eta}$ in order for (16b) to hold. ∎
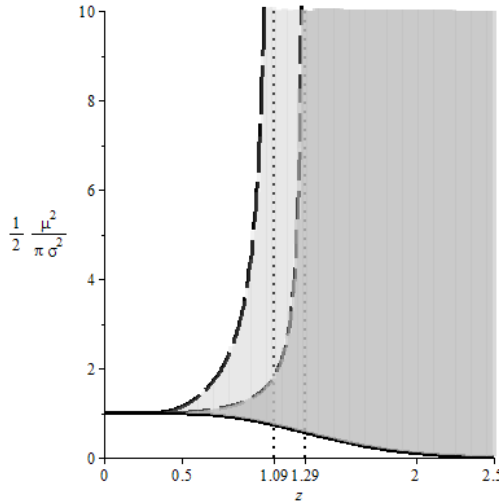
Figure 1: In both shaded areas the second-order conditions are satisfied and β_p^* and η_p^* are positive, that is, $\frac{\mu^2}{2\pi\lambda^2} - \sigma^2 > 0$. In the dark shaded area $\frac{\partial \beta_p^*}{\partial z_\alpha}$ is positive.

References

Alesina, Alberto, and Eliana La Ferrara. 2014. "A Test of Racial Bias in Capital Sentencing." *American Economic Review*, 104(11): 3397–3433.

Anderson, Simon, P., and Régis Renault. 2006. "Advertising Content." *American Economic Review*, 96(1): 93–113.

Anderson, Simon P., and Régis Renault. 2009. "Comparative Advertising: Disclosing Horizontal Match Information." *RAND Journal of Economics*, 40(3): 558–581.

Brown, Jennifer, and Dylan B. Minor. 2014. "Selecting the Best? Spillover and Shadows in Elimination Tournaments." *Management Science*, 60(12): 3087–3102.

Corts, Kenneth S. 2013. "Prohibitions on False and Unsubstantiated Claims: Inducing the Acquisition and Revelation of Information through Competition Policy." *Journal of Law and Economics*, 56(2): 453–486.

Davis, Michael L. 1994. "The Value of Truth and the Optimal Standard of Proof in Legal Disputes." *Journal of Law, Economics, and Organization*, 10(2): 343–359.

De Franco, Gus, Hai Lu, and Florin P. Vasvari. 2007. "Wealth Transfer Effects of Analysts' Misleading Behavior." *Journal of Accounting Research*, 45(1): 71–110.

DellaVigna, Stefano, and Matthew Gentzkow. 2010. "Persuasion: Empirical Evidence." *Annual Review of Economics*, 2(1): 643–669.

Demougin, Dominique, and Claude Fluet. 2008. "Rules of Proof, Courts, and Incentives." *RAND Journal of Economics*, 39(1): 20–40.

Drugov, Mikhail, and Marta Troya-Martinez. 2012. "Vague Lies: How to Advise Consumers When They Complain." CEPR Discussion Paper

No. 9201, http://www.cepr.org/active/publications/discussion_papers/dp.php?dpno=9201.

EC. 2011. "Consumer Market Study on Advice within the Area of Retail Investment Services - Final Report." European Commission, www.ec.europa.eu/consumers/rights/docs/investment_advice_study_en.pdf.

Egan, Mark, Gregor Matvos, and Amit Seru. 2017. "The Market for Financial Adviser Misconduct." *Journal of Political Economy, forthcoming.*

FTC. 2009. "Marketing Violent Entertainment to Children: A Sixth Follow-up Review of Industry Practices in the Motion Picture, Music Recording & Electronic Game Industries." Federal Trade Commission, www.ftc.gov/os/2009/12/P994511violententertainment.pdf.

Gabaix, Xavier, and David Laibson. 2006. "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets." *Quarterly Journal of Economics*, 121(2): 505–540.

Gennaioli, Nicola, and Andrei Shleifer. 2007. "The Evolution of Common Law." *Journal of Political Economy*, 115(1): 43–68.

Glaeser, Edward L., and Gergely Ujhelyi. 2010. "Regulating misinformation." *Journal of Public Economics*, 94(3-4): 247–257.

Griffin, John M., and Gonzalo Maturana. 2016. "Who Facilitated Misreporting in Securitized Loans?" *Review of Financial Studies*, 29(2): 384–419.

Hattori, Keisuke, and Keisaku Higashida. 2012. "Misleading advertising in duopoly." *Canadian Journal of Economics*, 45(3): 1154–1187.

Holmström, Bengt. 1999. "Managerial Incentive Problems: A Dynamic Perspective." *Review of Economic Studies*, 66(1): 169–182.

Inderst, Roman, and Marco Ottaviani. 2009. "Misselling through Agents." *American Economic Review*, 99(3): 883–908.

Inderst, Roman, and Marco Ottaviani. 2013. "Sales Talk, Cancellation Terms and the Role of Consumer Protection." *Review of Economic Studies*, 80(3): 1002–1026.

Issacharoff, Samuel. 1999. "Group Litigation of Consumer Claims: Lessons from the U.S. Experience." *Texas International Law Journal*, 34: 135–150.

Johnson, Justin P., and David P. Myatt. 2006. "On the Simple Economics of Advertising, Marketing, and Product Design." *American Economic Review*, 96(3): 756–784.

Judd, Kenneth L., and Michael H. Riordan. 1994. "Price and Quality in a New Product Monopoly." *Review of Economic Studies*, 61(4): 773–789.

Kamenica, Emir, and Matthew Gentzkow. 2011. "Bayesian Persuasion." *American Economic Review*, 101(6): 2590–2615.

Kaplow, Louis. 2011. "On the Optimal Burden of Proof." *Journal of Political Economy*, 119(6): 1104–1140.

Kartik, Navin. 2009. "Strategic Communication with Lying Costs." *Review of Economic Studies*, 76(4): 1359–1395.

Kartik, Navin, Marco Ottaviani, and Francesco Squintani. 2007. "Credulity, lies, and costly talk." *Journal of Economic Theory*, 134(1): 93–116.

Kaye, David. 1980. "Naked Statistical Evidence." *Yale Law Journal*, 89(3): 601–611.

Lewis, Tracy R., and David E. M. Sappington. 1994. "Supplying Information to Facilitate Price Discrimination." *International Economic Review*, 35(2): 309–327.

Lewis, Tracy R., and Michel Poitevin. 1997. "Disclosure of Information in Regulatory Proceedings." *Journal of Law, Economics, & Organization*, 13(1): 50–73.

Malmendier, Ulrike, and Devin Shanthikumar. 2014. "Do Security Analysts Speak in Two Tongues?" *Review of Financial Studies*, 27(5): 1287–1322.

Miceli, Thomas J. 1990. "Optimal Prosecution of Defendants Whose Guilt Is Uncertain." *Journal of Law, Economics, and Organization*, 6(1): 189–201.

Nelson, Phillip. 1970. "Information and Consumer Behavior." *Journal of Political Economy*, 78(2): 311–329.

Nelson, Phillip. 1974. "Advertising as Information." *Journal of Political Economy*, 82(4): 729–754.

Nesson, Charles. 1985. "The Evidence or the Event? On Judicial Proof and the Acceptability of Verdicts." *Harvard Law Review*, 98(7): 1357–1392.

OECD. 2010. "Consumer Policy Toolkit." The Organisation for Economic Co-operation and Development, http://www.oecd.org/sti/consumer/consumerpolicytoolkit.htm.

OFT. 2010. "Debt Management Guidance Compliance Review: Mystery Shopping." Office of Fair Trading, http://webarchive.nationalarchives.gov.uk/20140402142426/http://www.oft.gov.uk/shared_oft/business_leaflets/credit_licences/OFT1265.pdf.

OFT. 2011. "Consumer Contracts." Office of Fair Trading, http://webarchive.nationalarchives.gov.uk/20140402142426/http://www.oft.gov.uk/shared_oft/market-studies/consumercontracts/oft1312.pdf.

Piccolo, Salvatore, Piero Tedeschi, and Giovanni Ursino. 2015. "How limiting deceptive practices harms consumers." *RAND Journal of Economics*, 46(3): 611–624.

Piskorski, Tomasz, Amit Seru, and James Witkin. 2015. "Asset Quality Misrepresentation by Financial Intermediaries: Evidence from the RMBS Market." *Journal of Finance*, 70(6): 2635–2678.

Posner, Richard A. 1999. "An Economic Approach to the Law of Evidence." *Stanford Law Review*, 51(6): 1477–1546.

Renault, Régis. 2016. "Advertising in Markets." *Handbook of Media Economics*, , ed. Simon P. Anderson, Joel Waldfogel and Strömberg David, Chapter 4, 121–204. Elsevier.

Rhodes, Andrew, and Chris M. Wilson. 2018. "False advertising." *RAND Journal of Economics*, 49(2): 348–369.

Rubinfeld, Daniel L., and David E. M. Sappington. 1987. "Efficient Awards and Standards of Proof in Judicial Proceedings." *RAND Journal of Economics*, 18(2): 308–315.

Schrag, Joel, and Suzanne Scotchmer. 1994. "Crime and Prejudice: The Use of Character Evidence in Criminal Trials." *Journal of Law, Economics, & Organization*, 10(2): 319–342.

Spiegler, Ran. 2006. "Competition over agents with boundedly rational expectations." *Theoretical Economics*, 1(2): 207–231.

Sunstein, Cass, Reid Hastie, John Payne, David Schkade, and Kip Viscusi. 2002. *Punitive Damages: How Juries Decide.* Chicago:University Of Chicago Press.

The Economist. 2009. "Let's Make a Deal." 5 February, www.economist.com/node/13059804.

The Sunday Times. 2008. "Shops Cave in on Hagglers in 'Souk Britain'." 17 August, www.thesundaytimes.co.uk/sto/news/uk_news/article112528.ece.

The Telegraph. 2008. "Haggle your Way to a Bargain." 4 November, www.telegraph.co.uk/finance/personalfinance/consumertips/3377202/Haggle-your-way-to-a-bargain.html.

Troya-Martinez, Marta. 2016. "Vagueness and information-sharing." *Games and Economic Behavior*, 100: 301–320.