Bias, Noise and Consumer Protection*

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

May 5, 2011

Abstract

This paper analyzes the incentives of a seller to provide (un)biased and (im)precise advice about a complex product such as insurance, banking and telecommunication services. Misleading the buyers by biasing the advice upwards increases the revenues but also the expected fine imposed by the authority. Making the advice less precise does not affect the revenues in equilibrium but interferes in the authority's inference and affects the expected fine in a non-monotonic way. In particular, making the advice less precise makes it harder to convict the seller but increases the expected fine when the seller is found guilty. We find that biasing the advice and making it noisier are complements; in particular, a higher buyers' heterogeneity, a stricter standard of proof employed by the authority and a larger share of credulous consumers make the advice more biased and less precise.

JEL Classification: D18, D83, K40, L14.

1 Introduction

A buyer with limited or no experience walks into a shop in order to buy a mobile phone and sign up for a calling plan. There are many models of mobiles and different calling

*This paper has been previously circulated under the title "Bias, Noise and Litigation". We are grateful to Vicente Cuñat, Péter Eső, Natalia Fabra, Rosa Ferrer, Maria Goltsman, Martin Hellwig, Paul Klemperer, Marc Möller, David Myatt, Bent Nielsen and participants of the workshop "Advances in Industrial Organisation" (Vienna), IIOC 2010 (Vancouver), CRESSE 2010 (Chania), JEI 2010 (Madrid) and the seminars at the Max Planck Institute for Research on Collective Goods (Bonn) and IIES (Stockholm) for many insightful comments. Any remaining errors are our own.

[†]Department of Economics, University Carlos III de Madrid; mdrugov@eco.uc3m.es.

[‡]Department of Economics, University of Oxford; marta.troyamartinez@economics.ox.ac.uk.

plans and, moreover, the buyer is not sure about her needs, which will only become fully known through the use of the mobile. The seller, who may not know her needs either (and hence the product's quality of match), not only sets a price, but also provides some sales advice to the prospective buyer with the goal of convincing her. In doing so, he can choose to be more or less precise and can also simply cheat the buyer by exaggerating the goodness of the product. Sending an imprecise and/or misleading sales pitch comes, however, at a cost of customers' complaints that will damage seller's reputation, draw the attention of a consumer association, trigger an action by the competition authority or even result in a litigation.

Some other goods also share these features. For example, consumer electronics, insurance, banking and medical care contracts are all characterized by the fact that neither sellers nor buyers themselves fully know the ex-ante buyers' quality of match with the product, sellers' advice is widely used and buyers learn through experience. Also, these very products systematically top in the list of consumers complaints worldwide which brings about the policy relevance of the problem. For instance, the Federal Trade Commission in the United States has recently released a report listing top complaints consumers filed with the agency in 2009.[1] Among the top 15 places it is possible to find health care, internet services, credit cards, advance-fee loans, credit protection, banks and lenders in general and computer equipment and software. Similarly, in the United Kingdom, the Office of Fair Trading has just reported (through Consumer Direct, its telephone and online service provider of information on consumer rights) that among the top 10 complaints, it is possible to find mobile phones (service agreements), mobile phones (hardware), telephone services (land line), lap-tops, notebooks and tablet PCs.[2]

We explore the seller's incentives to provide (un)biased and (im)precise advice and the resulting equilibrium communication in the following model. The seller advises the buyer about a certain product trying to increase the buyer's perceived valuation of it. The buyer's valuation for the product is determined by the quality of the match between the product characteristics (unknown to the buyer) and buyer's taste (unknown to the seller). The buyer and the seller share the same prior about quality of the match, that

[1]See: http://www.ftc.gov/opa/2010/02/2009fraud.shtm (accessed on the 26 December 2010).
[2]See: http://www.consumerdirect.gov.uk/news/press_releases/national/2010/2009top10 (accessed on the 26 December 2010).

is, none of them has any private information.[3] The seller gives an advice (i.e., sends a signal) to the buyer about the product. The signal is the sum of the true match quality and an error term that represents frictions in the communication between the seller and the buyer.[4] Both the match quality and the error term are distributed normally and the seller secretly chooses the mean and publicly the variance of the error term.[5] We call them "bias" and "noise", respectively.

Upon receiving the signal, the buyer updates her beliefs about the match quality using her conjecture about the bias introduced by the seller (which has to be correct in the equilibrium). This posterior valuation of the product is the surplus generated if trade occurs and that the seller and the buyer share in a fixed proportion. This reflects the popular use of negotiation in selling banking or mobile phone contracts. While the bias unambiguously increases the perceived quality of the product and thus the seller's revenues, the effect of the noise is more subtle and depends on the particular model used.[6] We adopt a specification where the noise does not affect seller's revenues in the equilibrium, that is, when the buyer correctly anticipates the bias.

Misleading the buyers is costly since eventually, through use, they realize the true quality of the match and complain if the signal they received was much larger. These complaints trigger an action by an authority which might be a consumer protection authority (like Office of Fair Trading), regulator (like Financial Services Authority) or the court depending on the product and the country. The authority then investigates the seller 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 to conclude that the seller has biased his signal. More precisely, the authority presumes the innocence of the seller, that is, its null hypothesis is that there has been no bias.[7] 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 some given significance level which can be interpreted

[3]An alternative interpretation is that the seller is facing the whole demand curve rather than an individual buyer as, for example, in Justin P. Johnson & David P. Myatt (2006).

[4]The error term can also represent the fact that the true match quality will be learned only through experience.

[5]We also consider the unobservable choice of variance in Section 4.6.

[6]See, for example, Tracy R. Lewis & David E. M. Sappington (1994) and Johnson & Myatt (2006) where the seller always prefers to provide an extreme (i.e., maximum or minimum) amount of information.

[7]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.

as a 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 costs, i.e., the expected fine. The noise affects the costs through two channels: it decreases the probability that the seller is found guilty but increases the fine when the seller is found guilty. The total effect is U-shaped.[8]

We derive closed form solutions for the equilibrium bias and noise and find that they are complements. When the authority uses a stricter standard of proof or the match quality is more heterogenous, the seller both biases the signal more and makes it more noisy. Introduction of punitive damages makes the seller send a less biased and more precise signal.

In the baseline model the buyers are fully rational and are not misled in the equilibrium which makes the intervention by the authority somewhat questionable. We introduce a proportion of credulous consumers who blindly follow the seller's advice. We find that the seller then sends a more biased and less precise signal. Other extensions that we consider include non-observability of noise and informational externalities between two sellers.

1.1 Related literature

From the theory perspective, this model is related to the literature on career concerns (see Bengt Holmström (1999) and Mathias Dewatripont, Ian Jewitt & Jean Tirole (1999)). In a two-period version of these models, the worker's expected second-period performance is his ability since he will not make any effort. The ability is not observed and firms estimate it based on the first-period worker's performance.[9] The worker then exerts some effort in the first period in order to improve his performance and trick firms into thinking that he has a higher ability. In the equilibrium the firms are not tricked; instead, they anticipate the worker's effort and estimate the worker's ability correctly. Our model is a signal-jamming model as the career concerns models. However, our seller not only jams the signal by manipulating its mean (what we call "bias"), but also by changing its variance

[8] In many situations, especially with repeated purchases, reputational concerns also put a limit on the seller's bias. Our model is thus better suited for one-off or infrequent purchases.

[9] The output usually equals to the sum of the effort and the ability. See Dewatripont, Jewitt & Tirole (1999) for a more general setup.

(noise) and, hence, its information content.

The choice of variability as a strategic variable has been considered in a large variety of setups.[10] The general conclusion of these papers is that the players that are at disadvantage tend to choose more risky strategies than those who are in a favorable position, that is, they "gamble for resurrection". In our model the seller's expected revenues are additive in the expected quality and hence his marginal incentives do not depend on it. A different model is the one of Ran Spiegler (2006) where firms send noisy, unbiased and costless signals of their prices as obfuscation strategies in order to soften competition. In our model, there is no competition, the signal is biased (and noisy) and costly.

Johnson & Myatt (2006) consider how much information a monopolist would want to provide to its potential customers. Using the fact that information about the product rotates the demand curve, they show that the monopolist profits are convex in the information. This generalizes the result of Lewis & Sappington (1994) who showed that the monopolist supplies either maximal or minimal information. In both papers, the choice of the amount of information is costless for the seller, while in our setup this choice is endogenously linked with future fines and we also allow for the possibility of biasing the signal. Furthermore, both papers only consider the case when the amount of information is observable while we also analyze the unobservable case.[11]

The idea that introducing noise might make the bias less costly is also present in Wei Li (2010) where the sender shares the blame for the wrong message with the noisy and possibly biased intermediary. In Andreas Blume & Olivier Board (2010) noisy messages are useful in mitigating the conflict between the sender and the receiver in a cheap-talk model.

As we discussed in the beginning, this paper is motivated by the markets where

[10] For instance, Axel Anderson & Luís M. B. Cabral (2007) analyze a model of R&D races where there are two contestants who are allowed to choose the variance of their stochastic R&D technology. Ilia Tsetlin, Anil Gaba & Robert L. Winkler (2004) consider the choice of variability of the performance distribution in a multi-round contest. Matthias Kräkel, Petra Nieken & Judith Przemeck (2008) undertake a similar analysis in the auction context.

[11] To get an intuition of why observability matters, imagine that the product is bad on average, that is, in the absence of any information the buyer will not buy it (at a given price). To convince the buyer to buy, the seller has to generate a high realisation of the signal and wants the buyer to attach a high weight to this signal in her updating. If the noise is observable, adding noise to the signal will make the buyer attach a lower weight to the signal realization so the optimal strategy is actually to reduce the noise as much as possible (see Johnson & Myatt (2006) for the proof). However, if the noise is not observable, the buyer will not adjust the weight given to the signal to account for the noise and, therefore, the optimal strategy of the seller is to increase the noise of the signal.

consumers often complain and, therefore, is related to the literature on consumer protection.[12] Roman Inderst & Marco Ottaviani (2009), for instance, analyze contract cancellation and product return policies in a market where a more informed seller advises the buyer. Cancellation or return are costly for the seller and, therefore, the advice cannot be too misleading. 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.

The revenue part of our model can be seen as a cheap-talk model. The interests of the buyer and the seller are opposed since the seller always wants to charge a higher price while the buyer always wants to pay a lower price. Then, there is no information transmission in the equilibrium as is well known starting from the work of Vincent P. Crawford & Joel Sobel (1982).[13] Then, some exogenous (unmodelled) costs are usually added such as effort costs in career concerns models, refund costs in Inderst & Ottaviani (2009), or unspecified lying costs in Navin Kartik (2009). In our model the lying costs, which are the expected fine imposed on the seller, are microfounded.

The rest of the paper is organized as follows. Section 2 introduces the model. Section 3 finds the equilibrium and derives comparative statics results. Section 4 contains various extensions and modifications of the baseline model of Section 2. Section 5 concludes.

2 The model

Buyers approach the seller in order to get informed about his product and buy it. The seller does not know their valuations for the product 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. However, they know that θ is distributed as $\mathcal{N}(\mu, \sigma^2)$.[14] The production costs are normalized to zero.

[12] See John Vickers (2004) and Mark Armstrong (2008) for an overview on this literature.

[13] See, however, Archishman Chakraborty & Rick Harbaugh (2010) where information transmission is possible in a multidimensional cheap-talk model despite opposed preferences.

[14] As standard in career concerns models, we adopt a setting where the two sides start with the same prior to avoid issues of signalling. See Holmström (1999) and Dewatripont, Jewitt & Tirole (1999).

2.1 The communication process

The seller reveals product characteristics and gives advise about the shopping decision to the buyer. In doing so, he can distort the communication strategically. The seller can exaggerate some positive features of the product and he can also be vague about them. More precisely, the seller sends 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.[15] We refer to β as *bias* and to η^2 as *noise*. The signal is therefore distributed as $\mathcal{N}(\mu + \beta, \sigma^2 + \eta^2)$; denote its pdf by g. The buyer does not observe the bias since she cannot know if a certain feature is exaggerated without actually buying the product and trying it. She does observe the noise, however, since she can evaluate how precise the seller's explanations are, how much into details he goes, whether there is a trial period, etc.[16]

2.2 Buyers' valuation and seller's revenues

The better the quality of match of the product, θ, is the more the buyer values it. For simplicity, the buyer's valuation of the product is linear in the match quality θ and normalized to be equal to it. Since this quality is only known through use, her valuation at the moment of purchase is the expected match quality given a particular realisation s of the signal.

When the buyer observes a realisation s of the signal and the noise η^2, she updates

[15]The might be some natural noise in the communication process. We discuss this possibility in Section 4.5.

[16]We discuss the case of unobservable noise in Section 4.6.

the expected quality in the standard way using a conjecture about the seller's bias $\widetilde{\beta}$:[17]

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

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

Instead of fixing the price before sending the signal, the seller charges the buyer some (fixed) fraction of her valuation which is normalized to 1. That is, the seller sees the effect that the realisation of his signal, s, has had on the buyer and hence, he knows that the buyer is ready to pay up to $E\left[\theta \mid s, \eta^2, \widetilde{\beta}\right]$ for the product. This entails some negative payments which can be made negligible by making the average quality μ high enough.[18] It also means that the seller always sells to the consumer. This is in the spirit of the career concerns models, 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.[19]

Therefore, the seller's revenues are equal to the expected valuation of the buyer:

$$R(\beta, \eta) = \int_{-\infty}^{+\infty} E\left[\theta \mid s, \eta, \widetilde{\beta}\right] g(s) ds = \mu + \left(\beta - \widetilde{\beta}\right) \frac{\sigma^2}{\sigma^2 + \eta^2}. \tag{2}$$

Note that 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, adjusted for the weight. In equilibrium, when the buyer's conjecture about the bias used is correct, the seller will only be able to extract the prior expected match quality μ. The revenues decrease with the noise η (if $\beta > \widetilde{\beta}$) since the buyers pay less attention to the signal, i.e., place a smaller weight on it.

[17] We assume that the buyer's conjecture does not depend on the observed noise. This assumption is used in the literature, see, for example, Kenneth L. Judd & Michael H. Riordan (1994). Also, this makes the model with unobservable noise (Section 4.6) similar to the model of this section. It would be interesting, though, to investigate other possibilities since different equilibria are then obtained as is shown by Sandro Shelegia (2010).

[18] As we will see later, μ does not affect equilibrium bias and noise (Corollary 1).

[19] In a different model, Spiegler (2006) also allows negative prices for technical reasons.

2.3 Estimated bias and seller's costs

When the buyers buy the product and start using it, they discover the true match quality. Some buyers, especially those with a large gap between the advice (signal) received from the seller and the true quality, will complain to a public body that we call "authority" throughout the paper. This might be a consumer association, consumer protection agency, a competition authority, a sectoral regulator or the court, depending on the good or service in question and the country. 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 or imposing a fine on the seller. For the sake of simplicity, we adopt the last meaning and treat the punishment as a monetary fine.

The bias, as a deliberate and conscious way to mislead consumers, is illegal and a fine is imposed on the seller if any positive bias is discovered. 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).[20]

The noise, however, is not punished since it may come from some (unmodelled) external shocks in the communication process between the seller and the buyers. The seller may provide information that is vague and open to interpretation, so that the buyers themselves make personal, independent, errors of interpretation. Even if the seller does not bias the signal at all, there will be always some unlucky buyers that will receive a very high signal realisation as compared to their true match quality (unless the communication is absolutely noiseless). For instance, a bad quality of the information at the point of purchase may be due to incompetent sales staff who have trouble giving a clear advice. If these incompetent sales staff are nonetheless objective (i.e., they do not use

[20] In the common law tradition, misrepresentation is a contract law concept which means a false statement of fact made by one party to another party, which has the effect of inducing that party into the contract. For example, under certain circumstances, false statements or promises made by a seller of goods regarding the quality or nature of the product that the seller has may constitute misrepresentation. In English law, misrepresentation is regulated generally by the United Kingdom Misrepresentation Act 1967 and, in the consumer protection area, by Trade Descriptions Act 1968 and Consumer Protection Act 1987.

bias), the buyers will not be misled on average.

The authority conducts an investigation by taking a random sample of size N of buyers in order to estimate the bias introduced by the seller. It can also send "mystery shoppers" to the seller each of them reporting afterwards their experience. In the US and in the UK the competition and consumer protection authorities (FTC and OFT, respectively) routinely commission mystery shopping exercises as part of their market studies.[21]

Each mystery shopper (or buyer) $i = 1..N$ reports signal s_i received from the seller and quality θ_i. In Section 4.3 we consider the case when only signals are observed. Having received their reports, the authority uses a statistical test to determine if the seller is guilty of having introduced the bias.[22] It computes the error terms $\varepsilon_i = s_i - \theta_i$ and estimates the bias used as:

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

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

There is the presumption of innocence, that is, by default the seller is assumed not to have introduced any bias, unless enough evidence is provided. In order 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 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).[23] A natural interpretation

[21]See http://www.ftc.gov/os/2009/12/P994511violententertainment.pdf and http://www.oft.gov.uk/shared_oft/business_leaflets/credit_licences/OFT1265.pdf for recent examples (accessed on the 25 December 2010). 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 Samuel Issacharoff (1999) for a discussion of class actions and consumer protection.

[22]Scholars in law and economics have long been modelling court decision making as probabilistic, see A. Mitchell Polinsky & Steven Shavell (2000) for a survey. In legal literature it is also "... now generally accepted that since all the evidence is probabilistic ... evidence should not be excluded merely because its accuracy can be expressed in explicitly probabilistic terms..." (Richard A. Posner (2004), p. 370). See also Thomas J. Miceli (1990) and Michael L. Davis (1994) for describing statistical testing in relationship with court decision making.

[23]By definition, the probability that the standard normal random variable takes a value above z_α is equal to α.

of z_α is the "standard of proof". A higher z_α (lower α) means that it is more difficult to reject the null, and therefore, the authority needs more evidence to be convinced. For instance, if α is 5%, then $z_\alpha \approx 1.64$, and if α is 10%, then $z_\alpha \approx 1.28$.[24]

If the seller is found guilty, he is imposed a fine 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 $d\widehat{\beta}$ in Section 4.2. The seller's costs are the expected fine:

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

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

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 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 (the integral in (4a) decreases since z has a lower mean), but increases the chances of a really large fine once he is found guilty (the integral in (4a) is multiplied by a higher number). This is the next lemma. See also Figure 1.

Lemma 1 *The seller's costs (4b) are increasing in β and U-shaped with respect to η.*

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

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

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

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

[24] We assume that $z_\alpha < 2.436$, that is, $\alpha > 0.75\%$. This is needed for the second-order conditions to hold in Proposition 1.

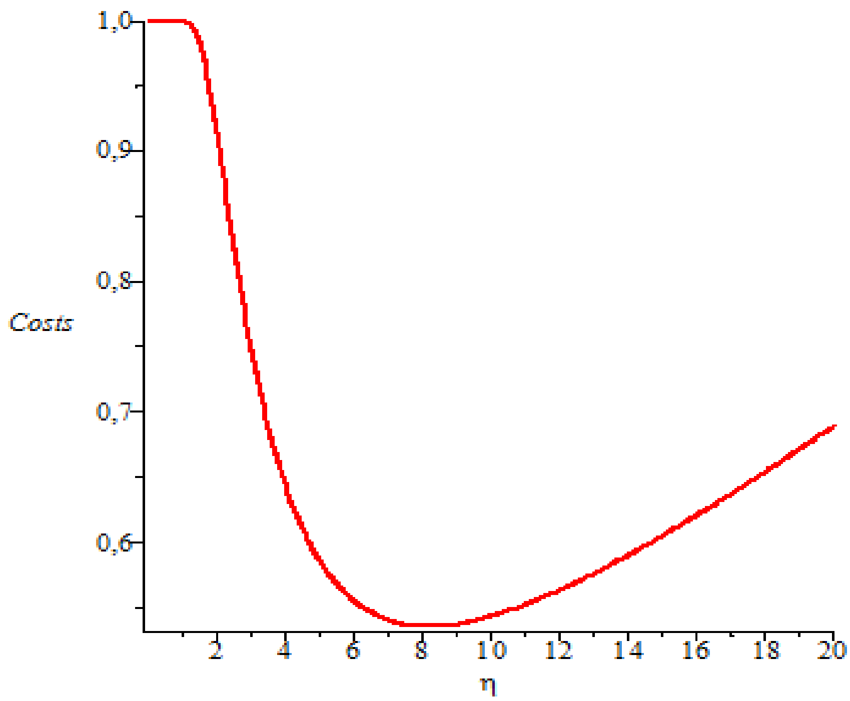

Figure 1: Seller's costs as a function of noise η; $\beta = 1, N = 25, z_\alpha = 1.64$.

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

3 Equilibrium

The seller maximizes his profits which are equal to the revenues (2) minus the costs (4b):

$$\Pi(\beta,\eta) = \mu + \left(\beta - \widetilde{\beta}\right)\frac{\sigma^2}{\sigma^2+\eta^2} - \beta\left(1 - \Phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right)\right) - \frac{\eta}{\sqrt{N}}\phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right).$$

The equilibrium is a pair (β^*, η^*) that maximizes seller's profits $\Pi(\beta,\eta)$ given buyer's conjecture β^*. Next proposition derives the equilibrium in a closed form.

Proposition 1 *The equilibrium bias and noise are*

$$\beta^* = \frac{\sigma}{\sqrt{N}}\frac{1}{z_\alpha}\sqrt{\kappa}, \tag{5a}$$

$$\eta^* = \sigma\sqrt{\kappa}, \tag{5b}$$

where $\kappa = \frac{\Phi\left(z_\alpha - \frac{1}{z_\alpha}\right) - z_\alpha\phi\left(z_\alpha - \frac{1}{z_\alpha}\right)}{1 - \Phi\left(z_\alpha - \frac{1}{z_\alpha}\right) + z_\alpha\phi\left(z_\alpha - \frac{1}{z_\alpha}\right)}$.

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

$$\frac{\sigma^2}{\sigma^2+\eta^2} - \left(1 - \Phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right) + z_\alpha \phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right)\right) = 0 \quad \text{(6a)}$$

$$-\left(\beta - \widetilde{\beta}\right)\frac{2\eta\sigma^2}{\left(\sigma^2+\eta^2\right)^2} - \frac{1}{\sqrt{N}}\phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right)\left[1 - \frac{\beta\sqrt{N}}{\eta}z_\alpha\right] = 0 \quad \text{(6b)}$$

In the equilibrium $\beta = \widetilde{\beta}$ and from (6b) $\frac{\beta\sqrt{N}}{\eta} = \frac{1}{z_\alpha}$. Plug this into (6a) to obtain β^* and then η^*. The second-order conditions are checked in the Appendix. ■

The bias increases the revenues by changing the signal distribution in the first-order stochastic dominance sense. Since the noise is observable, not only it affects the signal distribution but also the weight the buyer attaches to the signal in her posterior. The noise does not affect the revenues in the equilibrium since the buyer is not misled and pays the prior expected match quality μ.[25] Thus, given some bias, the noise is chosen to minimize the costs for that bias. Since the costs are U-shaped with respect to the noise as Lemma 1 showed, it results in an interior solution.

The comparative statics results are summarized in the next corollary.

Corollary 1 *The comparative statics of the equilibrium bias and noise (5a-5b) is the following:*

- $\frac{\partial\beta^*}{\partial\mu} = \frac{\partial\eta^*}{\partial\mu} = 0$;
- $\frac{\partial\beta^*}{\partial\sigma^2}, \frac{\partial\eta^*}{\partial\sigma^2} > 0$;
- $\frac{\partial\beta^*}{\partial z_\alpha}, \frac{\partial\eta^*}{\partial z_\alpha} > 0$;
- $\frac{\partial\beta^*}{\partial N} < 0$ *and* $\frac{\partial\eta^*}{\partial N} = 0$.

Given that the equilibrium bias and noise never change in opposite directions following a change in a parameter, we can say that they are *complements.*

[25]Outside the equilibrium, a higher noise decreases the seller's benefits from lying, if $\beta > \widetilde{\beta}$, as the buyer gives a lower weight to the signal.

The average quality μ does not affect the bias and the noise. Since the seller sells for any realisation of the signal and the buyer's valuation is linear in the 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 variance of the prior σ^2, which is the buyers' heterogeneity, can be thought of as affected by the seller at an earlier stage through the product design. In Johnson & Myatt (2006) the relationship between the buyers' heterogeneity and the informativeness of the signal is the opposite: 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. Therefore, if litigation is possible and the revenues are not affected by the noise, we should expect more heterogeneity in the product design to come together with larger bias and noise. Note that heterogeneity in the product design is a salient feature of the markets we described in the introduction.

A higher standard of proof z_α means that it is more difficult to convict the seller, so the bias becomes cheaper and is used more. The noise increases for two reasons. First, a higher bias calls for a higher noise to mitigate its effect. Second, even if the bias did not change, a higher z_α would lead to a higher noise since the costs are minimized when $\frac{\beta\sqrt{N}}{\eta} = \frac{1}{z_\alpha}$.

When the authority increases the sample size N, its estimation becomes more precise and the bias is more likely to be detected. Then, the bias is used less. The fact that the equilibrium noise does not depend on N is quite surprising since the first intuition would be that it should increase with N, that is, the seller should counteract an increase in the precision of the bias estimation. The reason is that the minimum costs with respect to the noise are still obtained when $\frac{\beta\sqrt{N}}{\eta} = \frac{1}{z_\alpha}$. The noise is found from (6a) $\frac{\sigma^2}{\sigma^2+\eta^2} = 1 - \Phi\left(z_\alpha - \frac{1}{z_\alpha}\right) + z_\alpha\phi\left(z_\alpha - \frac{1}{z_\alpha}\right)$, where the right hand side is the equilibrium marginal costs of bias; the equilibrium noise then does not depend on N.

Let us conclude this section with the following observation. In our model, setting a sufficiently high sample size N or a sufficiently low standard of proof z_α (or a sufficiently high punitive damages d that we investigate in Section 4.2) drives equilibrium noise and bias to zero. We would not, however, overestimate practical implications of this result for two reasons. First, this might not be the case if there is some natural and irreducible noise in communication between the buyer and the seller (see Section 4.5), as there will be then some lower bound on the equilibrium noise. Second, a richer model (which is beyond the scope of this paper and is an interesting topic for future work) is needed for a study of the optimal policy. Such a model will have physical costs of sending N mystery shoppers and welfare costs of punishing innocent sellers. Then, allowing for positive noise and bias in the equilibrium might be optimal.

4 Extensions and modifications

4.1 Why bias is bad: credulous buyers

In the baseline model, the buyers are rational and anticipate correctly the bias introduced by the seller. Then, the bias (and the noise) do not harm the buyers in the equilibrium: each of them pays on average μ which is the average match quality. It is not then clear why the seller should be prosecuted for introducing the bias.

However, there is extensive empirical evidence about the existence of credulous buyers. For instance, brokers are usually paid by the consumers but also receive compensation from the lenders - which influence the broker to offer the consumer loan terms or products that are not in the consumer's interest. According to the FTC (2008), "many consumers purportedly view mortgage brokers as trusted advisors who shop for the best loan for the consumer" (p. 16).

In this Section, we let some buyers to be credulous and blindly believe the seller's signal. More precisely, a share of credulous buyers c does not understand that the seller might send a biased and imprecise signal. Thus, these consumers blindly follow the signal.

The seller's expected revenues are then

$$R(\beta,\eta) = \mu + (1-c)\left(\beta - \widetilde{\beta}\right)\frac{\sigma^2}{\sigma^2+\eta^2} + c\beta.$$

Proposition 2 *When there is share c of credulous consumers, the equilibrium bias and noise are*

$$\begin{aligned} \beta_c^* &= \frac{\sigma}{\sqrt{N}}\frac{1}{z_\alpha}\sqrt{\kappa_c}, \\ \eta_c^* &= \sigma\sqrt{\kappa_c}, \end{aligned}$$

where $\kappa_c = \frac{\Phi\left(z_\alpha - \frac{1}{z_\alpha}\right) - z_\alpha\phi\left(z_\alpha - \frac{1}{z_\alpha}\right)}{1-\Phi\left(z_\alpha - \frac{1}{z_\alpha}\right) + z_\alpha\phi\left(z_\alpha - \frac{1}{z_\alpha}\right) - c}$.

They both increase with c.

As expected, equilibrium bias increases with the share of credulous consumers since marginal returns to bias are higher. The equilibrium noise also increases since it is complementary to the bias.

4.2 Punitive damages

Punitive (or exemplary) damages are used often in common law countries. In the US, both the frequency and the magnitude of punitive damages verdicts has increased dramatically in recent years (Cass Sunstein, Reid Hastie, John Payne, David Schkade & Kip Viscusi (2002)). We can model them by assuming that the seller is fined by the amount $d\widehat{\beta}$ when he is found guilty, $d \geq 1$.[26] The cost function (4b) is then multiplied by d. Following the same steps as in the baseline model, we obtain the equilibrium bias and noise.

Proposition 3 *When there are punitive damages d, the equilibrium bias and noise are*

$$\begin{aligned} \beta_d^* &= \frac{\sigma}{\sqrt{N}}\frac{1}{z_\alpha}\sqrt{\max\{\kappa_d, 0\}}, \\ \eta_d^* &= \sigma\sqrt{\max\{\kappa_d, 0\}}, \end{aligned}$$

where $\kappa_d = \frac{1-d\left[1-\Phi\left(z_\alpha - \frac{1}{z_\alpha}\right) + z_\alpha\phi\left(z_\alpha - \frac{1}{z_\alpha}\right)\right]}{d\left[1-\Phi\left(z_\alpha - \frac{1}{z_\alpha}\right) + z_\alpha\phi\left(z_\alpha - \frac{1}{z_\alpha}\right)\right]}$.

They both decrease with d.

[26] We abstract from the fact that in practice it is not very clear how the amount of punitive damages awarded depends on the harm made. Sunstein et al. (2002) say that there "...is inability to explain ... various punitive damages verdicts on a rational basis" (p. 2).

Proof. Similar to the proof of Proposition 1. ■

When $d = 1$, we obtain our baseline case results (5a-5b). When the punishment becomes more severe, that is, d increases, there is less bias and less noise. If d is high enough, the seller is completely deterred from misleading the buyers.

4.3 Only signals are observed[27]

If it takes a considerable amount of time to learn the quality of the match or there is a lot of subjectivity involved, then, when the authority sends mystery shoppers, they observe only signals s_i, $i = 1, ..., N$. The estimated bias is

$$\widehat{\beta} = \frac{1}{N}\sum_{i=1}^{N} s_i - \mu.$$

The authority tests $H_0 : \beta = 0$ against $H_1 : \beta > 0$ and punishes the seller only when H_0 is rejected. Since $s_i \rightsquigarrow \mathcal{N}(\mu + \beta, \sigma^2 + \eta^2)$, the estimator $\widehat{\beta}$ is distributed as $\widehat{\beta} \rightsquigarrow \mathcal{N}\left(\beta, \frac{\sigma^2+\eta^2}{N}\right)$. It has a higher variance, $\frac{\sigma^2+\eta^2}{N}$, than the one in Section 2.3, $\frac{\eta^2}{N}$, since the authority does not observe the match quality. The costs are

$$C(\beta, \eta) = \beta\left(1 - \Phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\sqrt{\sigma^2+\eta^2}}\right)\right) + \frac{\sqrt{\sigma^2+\eta^2}}{\sqrt{N}}\phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\sqrt{\sigma^2+\eta^2}}\right).$$

It turns out that the seller's problem is not concave with this cost function unless there are high enough punitive damages d as in Section 4.2.

Proposition 4 *If only signals are observed (and there are punitive damages d), the equilibrium bias and noise are*

$$\begin{aligned} \beta_s^* &= \frac{\sigma}{\sqrt{N}}\frac{1}{z_\alpha}\sqrt{\max\{\kappa_d, 0\} + 1}, \\ \eta_s^* &= \sigma\sqrt{\max\{\kappa_d, 0\}}, \end{aligned}$$

The equilibrium bias is higher, $\beta_s^* > \beta_d^*$, and the noise is the same, $\eta_s^* = \eta_d^*$, as in the

[27]The third possibility, namely, that only quality of the match is observed (or can be used in court) does not help here: since all consumers buy for all signal realizations, the ex post distribution of quality is the same as the prior one. This would be different if consumers buy only above a certain threshold as in Section 4.7.

case when both signals and match qualities are observed.

Proof. Similar to the proof of Proposition 1. ■

The equilibrium bias increases since it is now more difficult to convict the seller. The equilibrium noise does not change. This perhaps counterintuitive result has the same explanation as the one of the equilibrium noise not depending on the sample size N in Corollary 1. The minimum costs with respect to the noise are obtained when $\frac{\beta\sqrt{N}}{\sqrt{\sigma^2+\eta^2}} = \frac{1}{z_\alpha}$. The noise is then found from $\frac{\sigma^2}{\sigma^2+\eta^2} = d\left(1-\Phi\left(z_\alpha - \frac{1}{z_\alpha}\right) + z_\alpha\phi\left(z_\alpha - \frac{1}{z_\alpha}\right)\right)$, where the right hand side is the equilibrium marginal costs of bias that do not depend on σ^2.

4.4 One-stage procedure used by the authority

The procedure used by the authority in the baseline model of Section 2 is probably best interpreted as a two-stage procedure. At the first stage, the authority determines if the seller is guilty, and if it is the case, the authority fixes the punishment at the second stage. For this reason, the fine depends only on the estimated bias but not on the noise.

If the authority were to determine both if the seller is guilty and the size of the fine at the same stage, it then might account for the noise when fixing the fine, that is, to fine the seller by $\frac{\widehat{\beta}}{\eta/\sqrt{N}}$. The costs are then

$$C(\beta,\eta) = \int_{z_\alpha\frac{\eta}{\sqrt{N}}}^{+\infty} \widehat{\beta}\frac{\sqrt{N}}{\eta} f\left(\widehat{\beta}\right) d\widehat{\beta} = \beta\frac{\sqrt{N}}{\eta}\left(1-\Phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right)\right) + \phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right).$$

The costs now depend only on the ratio $\frac{\beta}{\eta}$. A higher noise decreases both the probability of conviction and the fine whenever it is imposed. Therefore, the seller will always introduce maximum possible noise. This observation leads to the next proposition.

Proposition 5 *If the authority adjusts the fine for the noise, there is no information transmission in the equilibrium.*

Proof. The first derivative of the profit function with respect to the noise is

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

In the equilibrium the buyers are not misled, $\beta = \widetilde{\beta}$, and this derivative is positive. The seller introduces then infinite noise, $\eta = +\infty$, and the buyers ignore the seller's signal. ■

4.5 Natural noise in the communication

Suppose there is a natural noise in the communication υ^2. For example, the communication is always imperfect unless the seller makes it better, possibly, at some effort costs. Then, the error term ε is distributed as $\varepsilon \rightsquigarrow \mathcal{N}(\beta, \upsilon^2 + \eta^2)$. Let us consider two polar cases. First, suppose that the effort costs are zero, that is, the seller can set η^2 equal to $-\upsilon^2$ making communication perfect at no cost. The seller's problem with respect to noise can be thought of as choosing the total variance $\upsilon^2 + \eta^2$ and its equilibrium level will be the same as in baseline model (5b), that is, $\sqrt{\upsilon^2 + \eta^{*2}} = \sigma\sqrt{\kappa}$. The equilibrium bias will also be the same as in the baseline model (5a).

The opposite case is when the effort costs are so high that the seller never chooses a negative η^2. The seller's problem with respect to noise is choosing the total variance $\upsilon^2 + \eta^2$ with the constraint $\eta^2 \geq 0$. If the (unconstrained) equilibrium noise η^* (5b) is higher than υ^2, then this constraint is not binding and the equilibrium is not affected. Otherwise, the equilibrium becomes $\eta^* = 0$ (that is, the variance of the error term is at its lowest level υ^2) and the equilibrium bias is found from the first-order condition (6a) when $\eta = \upsilon$.

4.6 Unobservable noise

In some cases it is more reasonable to assume that the noise chosen by the seller is not observed by the buyer. For instance, the seller can increase the noise by failing to explain hidden costs that the buyer may incur just once he uses the service or some features that the buyer did not think about or expected to use significantly (such as a bank overdraft). The unobserved noise can also come from the small script that is typically attached to this type of contracts. These clauses are usually not read in the moment of the purchase and they can be more or less precise. In all these examples, the buyer cannot assess at the moment of receiving the signal how precise it is.

The buyer updates her beliefs in a way similar to (1):

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

where $\widetilde{\eta}$ is the buyer's conjecture about the noise used by the seller (that has to be correct in the equilibrium). The seller's revenues are equal to the expected valuation that the seller extract from the buyers:

$$R(\beta, \eta) = \int_{-\infty}^{+\infty} E\left[\theta \mid s, \widetilde{\eta}, \widetilde{\beta}\right] g(s) ds = \mu + \left(\beta - \widetilde{\beta}\right) \frac{\sigma^2}{\sigma^2 + \widetilde{\eta}^2}.$$

Note that the seller extracts from the buyer the expected quality of match plus how much the buyer is misled into thinking that the product is good. The only difference with (2) is that the buyer uses her conjecture about the noise $\widetilde{\eta}$ since the noise is not observable. The noise η does not affect the revenues then. In equilibrium, when the buyer's conjecture about the bias used is correct, the seller extracts the expected match quality μ as before.

If buyers do not observe the noise, the authority does not observe it either and has to estimate it as

$$\widehat{\eta^2} = \frac{1}{N-1} \sum_{i=1}^{N} \left(\varepsilon_i - \widehat{\beta}\right)^2,$$

where $\widehat{\beta}$ is the estimated bias (3). This estimator is distributed as $\widehat{\eta^2} \rightsquigarrow \chi^2(N-1)$. The authority constructs statistics $t = \frac{\widehat{\beta}}{\widehat{\eta}/\sqrt{N}}$ which, under the null hypothesis of seller's innocence $\beta = 0$ has $t-$distribution with $N-1$ degrees of freedom. The authority rejects H_0 if $\frac{\widehat{\beta}}{\widehat{\eta}/\sqrt{N}} \geq t_\alpha$, where α is the significance level of the test.[28] However, for a given bias β, the true distribution of $\frac{\widehat{\beta}}{\widehat{\eta}/\sqrt{N}}$ is non-central $t-$distribution with $N-1$ degrees of freedom and non-centrality parameter $ncp = \beta\frac{\sqrt{N}}{\eta}$ (denote its cdf as T_{N-1}^{ncp}).[29] The seller's costs are

[28]By definition, the probability that a random variable following $t-$distribution with $N-1$ degrees of freedom takes a value above t_α is equal to α.

[29]Indeed,

$$\Pr\left\{\frac{\widehat{\beta}}{\widehat{\eta}/\sqrt{N}} \geq t_\alpha\right\} = \Pr\left\{\frac{\frac{\widehat{\beta}-\beta}{\eta/\sqrt{N}} + \frac{\beta}{\eta/\sqrt{N}}}{\sqrt{\frac{N-1}{N-1}\frac{\widehat{\eta}}{\eta}}} \geq t_\alpha\right\} = \Pr\left\{\frac{\mathcal{N}(0,1) + \frac{\beta}{\eta/\sqrt{N}}}{\sqrt{\frac{\chi^2_{N-1}}{N-1}}} \geq t_\alpha\right\} = 1 - T_{N-1}^{ncp}(t_\alpha),$$

where the last equality is the definition of the non-central $t-$distribution.

$$C(\beta,\eta) = \frac{\eta}{\sqrt{N}} \int_{t_\alpha}^{+\infty} t dT_{N-1}^{ncp}(t). \tag{7}$$

The only difference with costs (4a) is that the normal distribution there is now replaced with non-central $t-$distribution. A higher β increases ncp and, therefore, the distribution T_{N-1}^{ncp} moves to the right and the integral in (7) goes up. A higher η has two effects, as before: it makes it less likely for the seller to be found guilty (the integral in (7) decreases since ncp is lower), but increases the chances of a really large fine once he is found guilty (the integral in (7) is multiplied by a higher number). It is easily shown that costs (7) are decreasing at $\eta = 0$ and increasing at $\eta \to +\infty$ and, hence, minimized for some interior η.

The equilibrium then has a similar structure to the one in Section 3 (in particular, equilibrium bias and noise are complements) and is found in a similar way. Now, in the equilibrium $\beta = \widetilde{\beta}$ and $\eta = \widetilde{\eta}$. The first-order conditions of the seller's profit maximization in the equilibrium are

$$\frac{\partial \Pi(\beta,\eta)}{\partial \beta} = \frac{\sigma^2}{\sigma^2+\eta^2} - \frac{\partial}{\partial ncp} \int_{t_\alpha}^{+\infty} t dT_{N-1}^{ncp}(t) = 0, \tag{8a}$$
$$\frac{\partial \Pi(\beta,\eta)}{\partial \eta} = -\frac{1}{\sqrt{N}} \int_{t_\alpha}^{+\infty} t dT_{N-1}^{ncp}(t) + \frac{\beta}{\eta}\frac{\partial}{\partial ncp} \int_{t_\alpha}^{+\infty} t dT_{N-1}^{ncp}(t) = 0. \tag{8b}$$

Condition (8b) depends only on $ncp = \beta\frac{\sqrt{N}}{\eta}$ (after multiplying it by $\sqrt{N}$) and, therefore, in the equilibrium the ratio $\frac{\beta}{\eta}$ is chosen to minimize the costs with respect to the noise. Then, the noise is found from (8a), in which the second term, the equilibrium marginal costs of bias $\frac{\partial}{\partial ncp}\int_{t_\alpha}^{+\infty} t dT_{N-1}^{ncp}(t)$, does not depend on bias and noise.

With observable noise, $\beta\frac{\sqrt{N}}{\eta}$ is also chosen to minimize the costs. Then, η is found from $\frac{\partial \Pi(\beta,\eta)}{\partial \beta} = 0$. With unobservable noise, it is the buyer's conjecture $\widetilde{\eta}$ that makes the condition $\frac{\partial \Pi(\beta,\eta)}{\partial \beta} = 0$ satisfied. Since the equilibrium $\widetilde{\eta}$ is equal to η, all the difference comes from the fact that the authority statistics has normal distribution in the observable case and non-central $t-$distribution here. It also means that, when N increases and costs (7) converge to the costs in the baseline case (4a) since the non-observability of the noise matters less and less for the authority's estimation, the equilibrium then also converges to the equilibrium in the baseline case (5a-5b).

Finally, when the noise is unobserved, there is no off-equilibrium events. The buyer observes only a signal realization and any realization is compatible with any choice of bias and noise by the seller. The fact that the equilibrium when the noise is unobservable looks similar to the case when the noise is observable is one of the main reasons why we focused on the constant conjecture $\widetilde{\beta}(\eta)$ in the baseline model (see fn. 17).

4.7 Pricing

In many cases the price of a product is fixed and the buyer decides whether to buy it or not at that given price. It is an interesting question (for future work) how the pricing strategy interacts with the choice of bias and noise. Assume for now that the there is an exogenous price p. The buyer decides to buy if his posterior valuation of the product (1) is higher than p. The seller's revenues are then

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

and the seller's costs are unchanged. Unfortunately, we cannot obtain a closed-form solution for the bias and the noise in this setup. We can, however, determine how the bias to noise ratio compares to the one found in Section 3.

Proposition 6 *When $p > (<)\mu$, then $\frac{\beta}{\eta} > (<)\frac{\beta^*}{\eta^*}$, where β^* and η^* are defined in (5a) and (5b).*

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

$$pg(\overline{s}) - \left(1-\Phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right)+z_\alpha\phi\left(z_\alpha-\frac{\beta\sqrt{N}}{\eta}\right)\right) = 0 \quad \text{(9a)}$$

$$-p\eta g(\overline{s})\left[\frac{\beta-\widetilde{\beta}}{\sigma^2+\eta^2}+\frac{p-\mu}{\sigma^2}\right]-\frac{1}{\sqrt{N}}\phi\left(z_\alpha-\frac{\beta\sqrt{N}}{\eta}\right)\left[1-\frac{\beta\sqrt{N}}{\eta}z_\alpha\right] = 0 \quad \text{(9b)}$$

where $\overline{s} = p+\widetilde{\beta}+(p-\mu)\frac{\eta^2}{\sigma^2}$. In the equilibrium, $\beta=\widetilde{\beta}$. If $p>(<)\mu$ then $\frac{1}{\sqrt{N}z_\alpha} < (>)\frac{\beta}{\eta}$ in order for (9b) to hold. ■

On the revenue side, the only effect of the bias is to sell to the marginal consumer, i.e., to the one who receives signal $s = \overline{s}$. The noise has two effects on the marginal revenues: like the bias, it affects the decision of the marginal consumer (in the direction that depends on how the marginal consumer is related to the mean of the signal) but, since the noise is observable, it also changes the identity of the marginal consumer. In the equilibrium, the noise decreases the seller's revenues when $p > \mu$ and increases them otherwise. As a result, when $p > \mu$ the bias-to-noise ratio will be larger than the one in Section 3 and smaller otherwise. Actually, the analysis of the effects of the noise on the seller's revenues becomes essentially the one of Johnson & Myatt (2006). The case $p > \mu$ corresponds to their "niche market" and $p < \mu$ corresponds to their "mass market". However, contrary to their setup where noise is costless, the noise also affects the seller's costs and, therefore, an extreme noise might not be optimal.

4.8 Informational externalities between sellers

Suppose that, while still buying from a given seller, the buyer also observes a signal from another seller. For example, the buyer may want to buy only a certain brand of a mobile phone, say, for aesthetic or compatibility reasons, but she can learn about some new technical features from the mobile phones of other brands as well. Sellers cannot distinguish between buyers who buy from them and those who only get informed.

There are two sellers, seller 1 and seller 2, each sending a signal with their respective bias and noise. The buyer has the match quality $\theta_1 \rightsquigarrow \mathcal{N}(\mu_1, \sigma_1^2)$ with seller 1 and $\theta_2 \rightsquigarrow \mathcal{N}(\mu_2, \sigma_2^2)$ with seller 2. The correlation between θ_1 and θ_2 is ρ_θ and the correlation between distortions ε_1 and ε_2 is ρ_ε. Seller i's signal s_i is distributed as $\mathcal{N}(\mu_i + \beta_i, \sigma_i^2 + \eta_i^2)$. Having observed signals from both sellers, s_1 and s_2, the buyer who buys from seller 1 computes her expected match quality in the following way (see Appendix for the derivation):

$$E\left[\theta_1 \mid s_1, s_2, \eta_1, \eta_2, \widetilde{\beta}_1, \widetilde{\beta}_2\right] = \mu_1 + \frac{1+\frac{\eta_2^2}{\sigma_2^2} - \rho_\theta\left(\rho_\theta + \rho_\varepsilon \frac{\eta_1\eta_2}{\sigma_1\sigma_2}\right)}{\left(1+\frac{\eta_2^2}{\sigma_2^2}\right)\left(1+\frac{\eta_1^2}{\sigma_1^2}\right) - \left(\rho_\theta + \rho_\varepsilon \frac{\eta_1\eta_2}{\sigma_1\sigma_2}\right)^2}\left(s_1 - \mu_1 - \widetilde{\beta}_1\right)$$
$$+ \frac{\rho_\theta \frac{\eta_1^2}{\sigma_1\sigma_2} - \rho_\varepsilon \frac{\eta_1\eta_2}{\sigma_2^2}}{\left(1+\frac{\eta_2^2}{\sigma_2^2}\right)\left(1+\frac{\eta_1^2}{\sigma_1^2}\right) - \left(\rho_\theta + \rho_\varepsilon \frac{\eta_1\eta_2}{\sigma_1\sigma_2}\right)^2}\left(s_2 - \mu_2 - \widetilde{\beta}_2\right). \quad (10)$$

The two sellers move simultaneously. In the equilibrium, when deciding on the bias and noise, each of them knows the bias and noise of the other but not the realisation of the other seller's signal. Also, the buyer has correct conjectures about the biases of both sellers. If seller 1 decides to deviate from the equilibrium to some bias β_1 (when seller 2 does not) his expected revenues are

$$\begin{aligned} R_1(\beta_1, \eta_1) &= \int_{-\infty}^{+\infty}\int_{-\infty}^{+\infty} E\left[\theta_1 \mid s_1, s_2, \eta_1, \eta_2, \widetilde{\beta}_1, \widetilde{\beta}_2\right] dG(s_1)\, dG(s_2) \\ &= \mu_1 + \frac{1 + \frac{\eta_2^2}{\sigma_2^2} - \rho_\theta\left(\rho_\theta + \rho_\varepsilon \frac{\eta_1\eta_2}{\sigma_1\sigma_2}\right)}{\left(1 + \frac{\eta_2^2}{\sigma_2^2}\right)\left(1 + \frac{\eta_1^2}{\sigma_1^2}\right) - \left(\rho_\theta + \rho_\varepsilon \frac{\eta_1\eta_2}{\sigma_1\sigma_2}\right)^2}\left(\beta_1 - \widetilde{\beta}_1\right). \end{aligned} \tag{11}$$

The costs are still given by (4b).[30]

Proposition 7 *With informational externalities between the two sellers, when correlation is imperfect, $|\rho_\theta| < 1$ and $|\rho_\varepsilon| < 1$, there is a unique symmetric (pure-strategy) equilibrium*

$$\begin{aligned} \beta_i^e &= \tfrac{1}{z_\alpha}\tfrac{\eta_i^e}{\sqrt{N}}, \\ \eta_i^e &= \tfrac{\sigma_i}{\sqrt{2}}\sqrt{(k-1)\tfrac{1-\rho_\theta\rho_\varepsilon}{1-\rho_\varepsilon^2} + \sqrt{(k-1)^2\left(\tfrac{1-\rho_\theta\rho_\varepsilon}{1-\rho_\varepsilon^2}\right)^2 + 4k\tfrac{1-\rho_\theta^2}{1-\rho_\varepsilon^2}}} \end{aligned} \tag{12}$$

where κ is defined in Proposition 1.

If the match qualities are perfectly correlated, $|\rho_\theta| = \pm 1$, there is always equilibrium $\beta_1^e = \beta_2^e = \eta_1^e = \eta_2^e = 0$. If and only if $\kappa > 1$, there is another equilibrium[31]

$$\begin{aligned} \beta_i^e &= \frac{1}{z_\alpha}\frac{\eta_i^e}{\sqrt{N}}, \\ \eta_i^e &= \sigma_i\sqrt{\frac{\kappa - 1}{1 \pm \rho_\varepsilon}}. \end{aligned}$$

If the shocks are perfectly correlated, $|\rho_\varepsilon| = \pm 1$, if and only if $\kappa < 1$, there is a unique

[30] Despite the correlation between ε_1 and ε_2, the estimator of the bias is still the sample average. See George Casella & Roger L. Berger (2002), p. 358, ex. 7.18.

[31] A higher κ means a higher standard of proof z_α (and a lower α); $\kappa = 1$ corresponds to $z_\alpha \approx 1.70$ (and $\alpha \approx 4.5\%$).

symmetric (pure-strategy) equilibrium

$$\begin{aligned}\beta_i^e &= \frac{1}{z_\alpha}\frac{\eta_i^e}{\sqrt{N}},\\ \eta_i^e &= \sigma_i\sqrt{\frac{\kappa}{1-\kappa}\left(1\pm\rho_\theta\right)}.\end{aligned}$$

Proof. See Appendix. ■

The effects of the correlations ρ_θ and ρ_ε are sometimes ambiguous when they are both quite high. If ρ_ε is close to zero, then the effect of ρ_θ is to decrease the equilibrium value of bias and noise. This comes from the fact that when ρ_θ increases, the buyer puts a lower weight on the signal of her seller and a higher one on the signal of the other seller, as can be seen from (10). Then, each seller has less incentives to bias his signal and, correspondingly, he introduces less noise.

When ρ_θ is close to zero, an increase in ρ_ε has an opposite effect on the buyer's posterior. Indeed, the buyer trusts the signal from her seller more since she knows something about the realization of the error term through the correlation between ε_1 and ε_2. Then, the seller has more incentives to bias his signal and add noise. It can be easily checked that β_i^e and η_i^e in (12) are higher than β^* and η^* in (5a-5b) for $\rho_\theta = 0$ and $\rho_\varepsilon \neq 0$. Thus, the presence of another seller who sells an unrelated good but with correlated shocks in the communication process decreases welfare. The seller cannot commit not to bias his signal more and, while the consumers are not misled in the equilibrium, he ends up paying a higher fine.[32]

When the match quality of one seller is perfectly correlated with the one of the other, $|\rho_\theta| = \pm 1$, the situation is slightly different. If one of the sellers does not use any noise, he perfectly reveals his match quality and, therefore, the match quality of the other seller. The latter then cannot mislead the buyer and chooses zero bias and noise. There might be another equilibrium where both sellers do use some bias and noise and which is the limit of the equilibrium with imperfect correlation (12).

[32] See Margaret A. Meyer & John Vickers (1997) for a dynamic model of implicit incentives where informational externalities between two agents also may reduce welfare.

5 Conclusion

In this paper we investigated seller's incentives to provide (un)biased and (un)informative advice and the resulting equilibrium communication. We found that the biasing the advice and making it more noisy are complements: the seller employs says either an exact truth or a vague lie. For example, a higher buyers' heterogeneity, a higher standard of proof employed by the authority and a higher share of credulous consumers make the signal sent by the seller to the buyers more biased and less precise.

An interesting direction for future work is to characterize the optimal policy of the authority. In the model, a policy is described by three parameters: the number of consumers sampled N, the standard of proof used when establishing if the seller is guilty z_α and the size of the fine in relation to the estimated bias d. We have taken these parameters as given and found that the comparative statics are intuitive: a higher N, a lower z_α and a higher d reduce the equilibrium bias. This poses the question of why the authority does not set them up at the level that would completely deter the seller from biasing the signal. The reason is that there are of course costs of a strict policy. Surveying a lot of consumers, that is, a high N, is costly in monetary terms; a lower standard of proof z_α makes conviction of innocent sellers more likely and higher punitive damages d make innocent sellers pay a higher fine. Extending the model to incorporate these costs may produce relevant and realistic recommendations on the optimal policy.

Appendix

Second-order conditions for Proposition 1. Differentiate (6a) with respect to β to obtain $\frac{\partial^2\Pi}{\partial\beta^2}$ and (6b) 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:

$$\begin{aligned}
\frac{\partial^2\Pi}{\partial\beta^2} &= -\frac{\sqrt{N}}{\eta}\phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right)\left(1 - \frac{\beta\sqrt{N}}{\eta}z_\alpha + z_\alpha^2\right)\\
\frac{\partial^2\Pi}{\partial\eta^2} &= -\left(\beta - \widetilde{\beta}\right)\frac{\sigma^2 - 3\eta^2}{\left(\sigma^2+\eta^2\right)^3} - \frac{\beta^2\sqrt{N}}{\eta^3}\phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right)\left(1 - \frac{\beta\sqrt{N}}{\eta}z_\alpha + z_\alpha^2\right)\\
\frac{\partial^2\Pi}{\partial\eta\partial\beta} &= -\frac{2\eta\sigma^2}{\left(\sigma^2+\eta^2\right)^2} + \frac{\beta\sqrt{N}}{\eta^2}\phi\left(z_\alpha - \frac{\beta\sqrt{N}}{\eta}\right)\left(1 - \frac{\beta\sqrt{N}}{\eta}z_\alpha + z_\alpha^2\right)
\end{aligned}$$

In the equilibrium $\beta = \beta^* = \widetilde{\beta}$ and $\frac{\beta^*\sqrt{N}}{\eta^*} = \frac{1}{z_\alpha}$ so these derivatives become

$$\begin{aligned}
\frac{\partial^2 \Pi}{\partial \beta^2} &= -\frac{\sqrt{N}}{\eta^*}\phi\left(z_\alpha - \frac{1}{z_\alpha}\right) z_\alpha^2 < 0 \\
\frac{\partial^2 \Pi}{\partial \eta^2} &= -\frac{\beta^{*2}\sqrt{N}}{\eta^{*3}}\phi\left(z_\alpha - \frac{1}{z_\alpha}\right) z_\alpha^2 < 0 \\
\frac{\partial^2 \Pi}{\partial \eta \partial \beta} &= -\frac{2\eta^*\sigma^2}{\left(\sigma^2 + \eta^{*2}\right)^2} + \frac{\beta^*\sqrt{N}}{\eta^{*2}}\phi\left(z_\alpha - \frac{1}{z_\alpha}\right) z_\alpha^2
\end{aligned}$$

Check that the determinant of the Hessian is positive:

$$\begin{aligned}
\frac{\partial^2 \Pi}{\partial \beta^2}\frac{\partial^2 \Pi}{\partial \eta^2} - \left(\frac{\partial^2 \Pi}{\partial \eta \partial \beta}\right)^2 &= \frac{\beta^{*2} N}{\eta^{*4}}\phi^2\left(z_\alpha - \frac{1}{z_\alpha}\right) z_\alpha^4 - \left(\frac{2\eta^*\sigma^2}{\left(\sigma^2 + \eta^{*2}\right)^2}\right)^2 \\
&\quad + \frac{4\eta^*\sigma^2}{\left(\sigma^2 + \eta^{*2}\right)^2}\frac{\beta^*\sqrt{N}}{\eta^2}\phi\left(z_\alpha - \frac{1}{z_\alpha}\right) z_\alpha^2 - \frac{\beta^{*2} N}{\eta^{*4}}\phi^2\left(z_\alpha - \frac{1}{z_\alpha}\right) z_\alpha^4 \\
&= \frac{4\sigma^2}{\left(\sigma^2 + \eta^{*2}\right)^2}\left(\frac{\beta^*\sqrt{N}}{\eta^{*2}}\phi\left(z_\alpha - \frac{1}{z_\alpha}\right) z_\alpha^2 - \frac{\eta^*\sigma^2}{\left(\sigma^2 + \eta^{*2}\right)^2}\right) \\
&= \frac{4\sigma^2}{\eta^*\left(\sigma^2 + \eta^{*2}\right)^2}\left(\phi\left(z_\alpha - \frac{1}{z_\alpha}\right) z_\alpha - \frac{\kappa}{\left(1+\kappa\right)^2}\right)
\end{aligned}$$

It is positive for $z_\alpha < 2.436$, i.e., $\alpha > 0.75\%$.

These are only local second-order conditions. In our candidate equilibrium, i.e., if the seller's bias and buyer's conjecture are (5a) and the noise is (5b), the seller earns μ minus the costs. Then, he might prefer to deviate in the following way: stop the communication by introducing infinite noise, in which case the buyer disregards the signal and the revenues are still μ, and introduce an infinite negative bias in order to drive the costs to zero.[33] Usually, however, the seller is obliged to provide some minimal amount of information about the product; in other words, there is an upper bound on η. While a small negative bias can be interpreted as "modesty", an infinite one is clearly unrealistic. We thus assume that an upper bound on noise and a lower bound on bias are such that this deviation is not profitable and (5a-5b) is an equilibrium. ■

Derivation of the buyer's posterior in Section 4.8. See Henry Theil (1971), ch. 4.7, for the details of the derivation of posterior in the multivariate normal case.

[33]For a negative or zero bias the costs are strictly increasing in the noise, see the proof of Lemma 1.

Let us write the covariance matrix of θ_1, s_1 and s_2. Using

$$\begin{aligned}
Cov(\theta_1, s_1) &= Cov(\theta_1, \theta_1 + \varepsilon_1) = Var(\theta_1) = \sigma_1^2 \\
Cov(\theta_1, s_2) &= Cov(\theta_1, \theta_2 + \varepsilon_2) = Cov(\theta_1, \theta_2) = \rho_\theta \sigma_1 \sigma_2 \\
Cov(s_1, s_2) &= Cov(\theta_1 + \varepsilon_1, \theta_2 + \varepsilon_2) = Cov(\theta_1, \theta_2) + Cov(\varepsilon_1, \varepsilon_2) = \rho_\theta \sigma_1 \sigma_2 + \rho_\varepsilon \eta_1 \eta_2
\end{aligned}$$

the matrix is

	θ_1	s_1	s_2
θ_1	σ_1^2	σ_1^2	$\rho_\theta\sigma_1\sigma_2$
s_1	σ_1^2	$\sigma_1^2 + \eta_1^2$	$\rho_\theta\sigma_1\sigma_2 + \rho_\varepsilon\eta_1\eta_2$
s_2	$\rho_\theta\sigma_1\sigma_2$	$\rho_\theta\sigma_1\sigma_2 + \rho_\varepsilon\eta_1\eta_2$	$\sigma_2^2 + \eta_2^2$

Denote the covariance matrix of s_1 and s_2 as Σ and find its inverse:

$$\Sigma^{-1} = \frac{1}{\det(\Sigma)} \begin{pmatrix} \sigma_2^2 + \eta_2^2 & -\rho_\theta\sigma_1\sigma_2 - \rho_\varepsilon\eta_1\eta_2 \\ -\rho_\theta\sigma_1\sigma_2 - \rho_\varepsilon\eta_1\eta_2 & \sigma_1^2 + \eta_1^2 \end{pmatrix},$$

where $\det(\Sigma) = (\sigma_2^2 + \eta_2^2)(\sigma_1^2 + \eta_1^2) - (\rho_\theta\sigma_1\sigma_2 + \rho_\varepsilon\eta_1\eta_2)^2$. The expectation of θ_1 conditional on s_1 and s_2, $E\left[\theta_1 \mid s_1, s_2, \eta_1, \eta_2, \widetilde{\beta}_1, \widetilde{\beta}_2\right]$, equals to

$$\begin{aligned}
& \mu_1 + \begin{pmatrix} \sigma_1^2 & \rho_\theta\sigma_1\sigma_2 \end{pmatrix} \Sigma^{-1} \begin{pmatrix} s_1 - \mu_1 - \widetilde{\beta}_1 \\ s_2 - \mu_2 - \widetilde{\beta}_2 \end{pmatrix} \\
= {} & \mu_1 + \frac{1}{\det(\Sigma)} \begin{pmatrix} \sigma_1^2(\sigma_2^2 + \eta_2^2) - \rho_\theta\sigma_1\sigma_2(\rho_\theta\sigma_1\sigma_2 + \rho_\varepsilon\eta_1\eta_2) \\ -\sigma_1^2(\rho_\theta\sigma_1\sigma_2 + \rho_\varepsilon\eta_1\eta_2) + \rho_\theta\sigma_1\sigma_2(\sigma_1^2 + \eta_1^2) \end{pmatrix}^{-1} \begin{pmatrix} s_1 - \mu_1 - \widetilde{\beta}_1 \\ s_2 - \mu_2 - \widetilde{\beta}_2 \end{pmatrix} \\
= {} & \mu_1 + \frac{\sigma_1^2(\sigma_2^2 + \eta_2^2) - \rho_\theta\sigma_1\sigma_2(\rho_\theta\sigma_1\sigma_2 + \rho_\varepsilon\eta_1\eta_2)}{(\sigma_2^2 + \eta_2^2)(\sigma_1^2 + \eta_1^2) - (\rho_\theta\sigma_1\sigma_2 + \rho_\varepsilon\eta_1\eta_2)^2} \left(s_1 - \mu_1 - \widetilde{\beta}_1\right) \\
& + \frac{-\sigma_1^2\rho_\varepsilon\eta_1\eta_2 + \rho_\theta\sigma_1\sigma_2\eta_1^2}{(\sigma_2^2 + \eta_2^2)(\sigma_1^2 + \eta_1^2) - (\rho_\theta\sigma_1\sigma_2 + \rho_\varepsilon\eta_1\eta_2)^2} \left(s_2 - \mu_2 - \widetilde{\beta}_2\right) \\
= {} & \mu_1 + \frac{1 + \frac{\eta_2^2}{\sigma_2^2} - \rho_\theta\left(\rho_\theta + \rho_\varepsilon\frac{\eta_1\eta_2}{\sigma_1\sigma_2}\right)}{\left(1 + \frac{\eta_2^2}{\sigma_2^2}\right)\left(1 + \frac{\eta_1^2}{\sigma_1^2}\right) - \left(\rho_\theta + \rho_\varepsilon\frac{\eta_1\eta_2}{\sigma_1\sigma_2}\right)^2} \left(s_1 - \mu_1 - \widetilde{\beta}_1\right) \\
& + \frac{\rho_\theta\frac{\eta_1^2}{\sigma_1\sigma_2} - \rho_\varepsilon\frac{\eta_1\eta_2}{\sigma_2^2}}{\left(1 + \frac{\eta_2^2}{\sigma_2^2}\right)\left(1 + \frac{\eta_1^2}{\sigma_1^2}\right) - \left(\rho_\theta + \rho_\varepsilon\frac{\eta_1\eta_2}{\sigma_1\sigma_2}\right)^2} \left(s_2 - \mu_2 - \widetilde{\beta}_2\right).
\end{aligned}$$

■

Proof of Proposition 7.

Seller 1 maximizes revenues (11) minus costs (4b) taking η_2 as given. In the equilibrium $\beta_1 = \widetilde{\beta}_1$ and, as in the proof of Proposition 1, $\frac{\beta_1\sqrt{N}}{\eta_1} = \frac{1}{z_\alpha}$. The first-order condition with respect to the bias becomes

$$\frac{1+\frac{\eta_2^2}{\sigma_2^2} - \rho_\theta\left(\rho_\theta + \rho_\varepsilon \frac{\eta_1\eta_2}{\sigma_1\sigma_2}\right)}{\left(1+\frac{\eta_2^2}{\sigma_2^2}\right)\left(1+\frac{\eta_1^2}{\sigma_1^2}\right) - \left(\rho_\theta + \rho_\varepsilon \frac{\eta_1\eta_2}{\sigma_1\sigma_2}\right)^2} = 1 - \Phi\left(z_\alpha - \frac{1}{z_\alpha}\right) + z_\alpha \phi\left(z_\alpha - \frac{1}{z_\alpha}\right). \tag{13}$$

Note that (13) is symmetric in "normalized" noises $\frac{\eta_1}{\sigma_1}$ and $\frac{\eta_2}{\sigma_2}$. To find symmetric equilibria, denote $\widetilde{\eta} = \frac{\eta_1}{\sigma_1} = \frac{\eta_2}{\sigma_2}$ and rewrite (13) as

$$\left(1-\rho_\varepsilon^2\right)\widetilde{\eta}^4 - (k-1)\left(1-\rho_\theta\rho_\varepsilon\right)\widetilde{\eta}^2 - k\left(1-\rho_\theta^2\right) = 0$$

which is a quadratic equation in $\widetilde{\eta}^2$ (unless $\rho_\varepsilon = \pm 1$). It has two roots of opposite signs and the positive root is

$$\widetilde{\eta}^2 = \frac{(k-1)\left(1-\rho_\theta\rho_\varepsilon\right) + \sqrt{(k-1)^2\left(1-\rho_\theta\rho_\varepsilon\right)^2 + 4k\left(1-\rho_\theta^2\right)\left(1-\rho_\varepsilon^2\right)}}{2\left(1-\rho_\varepsilon^2\right)}.$$

Considering the cases of $\rho_\varepsilon = \pm 1$ and $\rho_\theta = \pm 1$ is straightforward.

We could not prove the absence of asymmetric equilibria; however, we have not encountered them in many numerical examples that we plotted.[34]

The second-order conditions are difficult to check but they are satisfied, by continuity, at least for ρ_ε and ρ_θ close to zero. ■

References

Anderson, Axel, and Luís M. B. Cabral. 2007. "Go for broke or play it safe? Dynamic competition with choice of variance." *The RAND Journal of Economics*,

[34] To find asymmetric equilibria (or to prove that they do not exist), denote $\widetilde{\eta}_1 = \frac{\eta_1}{\sigma_1}$ and $\widetilde{\eta}_2 = \frac{\eta_2}{\sigma_2}$. The reaction curve $\widetilde{\eta}_1(\widetilde{\eta}_2)$ can be found from rewriting (13) as

$$\widetilde{\eta}_1^2\left(1+\widetilde{\eta}_2^2\left(1-\rho_\varepsilon^2\right)\right) + (\kappa-1)\widetilde{\eta}_1\widetilde{\eta}_2\rho_\theta\rho_\varepsilon - \kappa\left(1-\rho_\theta^2+\widetilde{\eta}_2^2\right) = 0.$$

Equilibria are its intersections with the (symmetric) reaction curve $\widetilde{\eta}_2(\widetilde{\eta}_1)$.

38(3): 593–609.

Armstrong, Mark. 2008. "Interactions between Competition and Consumer Policy." *Competition Policy International*, 4(1): 96–147.

Blume, Andreas, and Olivier Board. 2010. "Intentional Vagueness." http://www.pitt.edu/~ojboard/papers/vagueness.pdf.

Casella, George, and Roger L. Berger. 2002. *Statistical inference.* Belmont:Duxbury Press.

Chakraborty, Archishman, and Rick Harbaugh. 2010. "Persuasion by Cheap Talk." *American Economic Review*, 100(5): 2361–82.

Crawford, Vincent P., and Joel Sobel. 1982. "Strategic Information Transmission." *Econometrica*, 50(6): 1431–1451.

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.

Dewatripont, Mathias, Ian Jewitt, and Jean Tirole. 1999. "The Economics of Career Concerns, Part I: Comparing Information Structures." *The Review of Economic Studies*, 66(1): 183–198.

FTC. 2008. "Before the Board of Governors of the Federal Reserve System: In the Matter of Request for Comments on Truth in Lending, Proposed Rule Docket No. R-1305." Available at http://www.federalreserve.gov/SECRS/2008/April/20080424/R-1305/R-1305_1349_1.pdf (accessed on the 26 December 2010).

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

Inderst, Roman, and Marco Ottaviani. 2009. "Sales Talk, Cancellation Terms, and the Role of Consumer Protection." http://www.wiwi.uni-frankfurt.de/profs/inderst/Theory/refunds.pdf.

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." *The American Economic Review*, 96(3): 756–784.

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

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

Kräkel, Matthias, Petra Nieken, and Judith Przemeck. 2008. "Risk Taking in Winner-Take-All Competition." Bonn Econ Discussion Paper bgse7_2008.

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

Li, Wei. 2010. "Peddling Influence through Intermediaries." *The American Economic Review*, 100(3): 1136–1162.

Meyer, Margaret A., and John Vickers. 1997. "Performance Comparisons and Dynamic Incentives." *The Journal of Political Economy*, 105(3): 547–581.

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

Polinsky, A. Mitchell, and Steven Shavell. 2000. "The Economic Theory of Public Enforcement of Law." *Journal of Economic Literature*, 38(1): 45–76.

Posner, Richard A. 2004. *Frontiers of Legal Theory.* Cambridge:Harvard University Press.

Shelegia, Sandro. 2010. "Quality Choice, Observability and Price Signaling." http://homepage.univie.ac.at/sandro.shelegia/Personal/Research_files/quality choice.pdf.

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.

Theil, Henry. 1971. *Principles of Econometrics.* Amsterdam:North-Holland.

Tsetlin, Ilia, Anil Gaba, and Robert L. Winkler. 2004. "Strategic Choice of Variability in Multiround Contests and Contests with Handicaps." *Journal of Risk and Uncertainty*, 29(2): 143–158.

Vickers, John. 2004. "Economics for Consumer Policy." In *Proceedings of the British Academy.* Vol. 125. Oxford:Oxford University Press.