



ELSEVIER

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb



Noisy signaling: Theory and experiment ☆

Thomas de Haan, Theo Offerman, Randolph Sloof*

Amsterdam School of Economics, Roetersstraat 11, 1018 WB Amsterdam, The Netherlands

ARTICLE INFO

Article history:

Received 16 September 2009

Available online xxxx

JEL classification:

C91

D82

Keywords:

Signaling games

Noise

Separation

Experiments

ABSTRACT

We introduce noise in the signaling technology of an otherwise standard wasteful signaling model (Spence, 1973). We theoretically derive the properties of the equilibria under different levels of noise and we experimentally test how behavior changes with noise. We obtain three main insights. First, if the amount of noise increases, high types aiming for separation (must) increase their signaling expenditures. This theoretical prediction is confirmed in our experiment. Second, for intermediate and high levels of noise, a separating and pooling equilibrium co-exist. In the experiment, subjects tend to shift from coordinating on a separating outcome to a pooling one as noise increases. Third, a surprising theoretical insight is that a separating equilibrium ceases to exist for low levels of noise (and an unfavorable prior). Yet in the experiment subjects then do coordinate on separation. A simple attraction learning model incorporating belief learning, imitation and reinforcement, explains this stable non-equilibrium behavior.

© 2011 Elsevier Inc. All rights reserved.

1. Introduction

An important insight already dating back to Veblen (1899) is that people may be willing to invest in wasteful activities to get some of their unobserved qualities recognized. He argued that the nineteenth century 'nouveau riche' heavily invested in wasteful, highly visible ways of ('conspicuous') consumption to openly display their wealth. These lavish expenditures served no other function than to signal social status. In a seminal contribution, Spence (1973) formalized this idea by providing the first game theoretic analysis of a signaling game in which a job applicant's investment in education is seen as a costly signal of his unobservable ability type. In this model, education does not improve productivity at all, but allows higher ability types to separate from lower ability types because obtaining education is less costly for them. The observed expenditures on education can thus effectively serve as convincing evidence of unobserved ability. Since then, the Spence model has become one of the most important tools in modern economics and beyond.

Given its central place, it is important to know to what extent the conclusions from Spence's wasteful signaling model are robust. A particularly relevant question is whether the results remain valid when *noise* is introduced in the signaling technology, such that the signal effectively received is not identical to the signal chosen or intended. In that case, unlike in the Spence setup, the sender has no perfect control over the signal the receiver actually observes.¹ In Spence's education application, the signal production technology is inherently noisy, because the exams that make up a particular degree will

☆ The authors are grateful to CREED-programmer Jos Theelen for programming the experiment. The research of Theo Offerman has been made possible by a fellowship of the Royal Netherlands Academy of Arts and Sciences.

* Corresponding author.

E-mail address: r.sloof@uva.nl (R. Sloof).

¹ Another issue that already caught the attention of Spence concerns the assumed perfect negative correlation between signaling costs and type; higher productivity types always have lower costs of choosing a given level of education than lower types have. This allows that (in market equilibrium) education becomes a perfect point prediction of productivity. As Spence (1974, Chapter 6) shows, however, similar conclusions are obtained when this assumption is relaxed. As long as there is a negative correlation between costs of signaling and type, education is an effective (though imperfect) signal of productivity.

not always be equally difficult over the many occasions they are taken. Alternatively, there is a possibility that the job applicant accidentally underperforms at the exam, perhaps because he is bothered by a flu. For these reasons, it is very hard for receivers to precisely judge the actual (costs of the) signal received. In another prominent application, the receiver is not able to exactly identify the cost of a signal. Consider the case of advertising as a costly signal to potential consumers. If consumers infer the quality of a firm's product from its expenditure on advertising, high-quality firms have an incentive to distinguish themselves from low-quality sellers by spending money on otherwise useless advertising (Nelson 1970, 1974; Kihlstrom and Riordan, 1984; Milgrom and Roberts, 1986). This may explain why firms are willing to spend huge amounts, like for example the 3 million dollars firms paid on average for broadcasting a 30-second spot during the 2010 Super Bowl.² For TV watchers at the time it may have been very hard to precisely judge the actual costs of the signal, because the typical football fan is unaware of the precise amounts of money involved. Most of them will only have an imprecise notion that the costs must have been high.

In fact, the assumption that there is no noise at all in the signal seems to be too strong in most real world applications. In this paper we therefore investigate both theoretically and experimentally what happens when the assumption of no noise is dropped. Our results show that a noisy signaling game differs profoundly from a standard signaling game without noise, both in terms of theoretical predictions and in terms of experimental outcomes.

Like Spence (1973), we focus on a pure signaling game in which signaling is in itself entirely wasteful. In our game, a seller offers a product for sale that is either of high quality or of low quality. Nature first determines the quality of his product. Only the seller learns the actual quality. Then the seller chooses his level of signal costs, i.e. we equate the message he chooses with its costs. To these signal costs nature adds a random noise term. The buyer observes the resulting overall signal, but not the original signal costs, and decides whether or not to buy. Preferences are such that the buyer prefers to buy if and only if quality is high. Moreover, a sale is more valuable to a high-quality seller than to a low-quality seller. In this game signaling is completely wasteful, because both seller types would prefer to pool on zero signal costs if the buyer would ignore the seller's signal.

Intuitively one would a priori expect that, the higher the amount of noise, the greater pains high-quality types have to take to get their unobserved quality recognized. With more noise they will thus probably try harder to separate themselves by choosing higher levels of signal costs. At the same time high-quality types may be less inclined to separate when the noise in the signal becomes large, because noise makes signaling both more expensive and less informative. The latter holds because in the presence of noise, the actual signal the receiver observes does not provide conclusive evidence about the sender's type and the high-quality sender faces the risk of a 'bad' draw and thus being considered a low-quality type. The observation that separation becomes relatively less attractive when the noise in the signal increases suggests that noise works in favor of pooling. Overall, intuition seems to suggest that the introduction of noise first induces high-quality sellers to choose higher levels of signal costs, but when the noise becomes too large these sender types will stop aiming for separation and will pool with the low-quality sellers instead.

Our theoretical analysis shows that these common sense intuitions are only partly in line with equilibrium predictions. It indeed holds true that the level of signal costs the high-quality seller chooses in a separating equilibrium increases with noise over a wide range of noise levels. High levels of noise thus force high-quality sellers to choose high signal costs, just as in the advertisement example mentioned above.³ It is also the case that a separating equilibrium ceases to exist when the noise becomes large, while irrespective of the noise level, a pooling (on no signaling) equilibrium always exists. High-quality types will thus stop aiming for separation for high noise levels. Surprisingly, however, a (pure strategy) separating equilibrium also does not exist when just a small amount of noise is introduced.⁴ The intuition behind this a priori somewhat counterintuitive result runs as follows. In a separating equilibrium, the buyer buys if and only if she receives a signal higher than a cutoff. With a low level of noise, this cutoff is much lower than the one used in the case without noise. It therefore becomes attractive for the low-quality seller to jump from providing zero signal costs to his interior optimum. This undermines the logic of the separating equilibrium in which the low-quality seller should refrain from signaling.

The theoretical prediction that just a little bit of noise thwarts all attempts to invest in wasteful signaling activities challenges the robustness of the conclusions obtained from Spence's original (1973) contribution. It therefore becomes important to test this prediction empirically. Moreover, the existence of multiple equilibria is another reason why data are

² See <http://www.bnet.com/blog/advertising-business/rise-in-super-bowl-ad-prices-threatens-raw-deal-for-advertisers/4222>. These costs came on top of the costs of producing the commercial, like fees for actors, equipment and advertisement agencies.

³ Another example of high signaling costs is provided by the Yanomamö, a contemporary tribe of about twenty thousand Indians living in the Amazon rainforest on the border between Brazil and Venezuela. Yanomamö men sometimes risk their lives in their vigorous pursuit of a fierce image (Chagnon, 1992). Disputes about women occasionally culminate in a club fight, where two males take turns striking each other on the head with a club of eight to ten feet long. The men are very proud of their heads that are covered with deep scars. Some men have a tonsure shaved on the top of their heads and they rub red pigment on their scars to make sure that nobody misses them. Regularly, people get killed in club fights or other outbursts of violence. Having a fierce image pays off among the Yanomamö. Chagnon (1988) reports that men who killed had on average two and half times as many wives and three times as many children than men who did not. Note that in this example, the production technology of the signal is noisy, because e.g. the deepness of a scar may be affected by incidental factors such as the angle in which the club hit the head.

⁴ This theoretical result only applies in the case of an unfavorable prior, where the prior belief about quality is not sufficient to support a sale. In the experiment we focus on this more interesting case where information transmission is necessary to realize the efficiency gains from trade. In our theoretical analysis we consider the opposite case of a favorable prior as well.

important. Even when a separating equilibrium exists besides a pooling one, it is a priori unclear whether and when people will actually coordinate on a separating outcome. The actual impact of noise on signaling thus remain a priori uncertain. Field data of signaling games lack the control needed to investigate the effect of noise on signaling empirically. Ideally, one investigates actual signaling in games that only differ in the amount of noise in the signaling technology and are isomorphic in all other dimensions. This level of control can only be acquired in a laboratory experiment. We therefore test the noisy wasteful signaling model in the lab.

In our experimental design we vary the level of noise between our four treatments: from no noise to low noise (without a separating equilibrium) to intermediate noise (with a unique separating equilibrium based on an intermediate level of signal costs) to high noise (with a unique separating equilibrium based on a high level of signal costs). Our design allows us to address three main questions: (i) Do the signal costs chosen by those high-quality sellers that aim for separation increase with the level of noise? (ii) Are high-quality sellers less inclined to separate when the noise becomes high? and (iii) Do subjects refrain from separating for a low noise level that prevents separation in theory? For the treatments where a pooling and a separating equilibrium coexist, we also investigate the effect of noise on the proximity of actual play to either equilibrium.

We obtain the following experimental results. For no, low, and intermediate noise, subjects tend to separate. For high levels of noise some matching groups still play according to the logic of a separating equilibrium, but the majority switches to pooling on no signal costs. Conditional on aiming for separation, high-quality sellers' signal costs increase monotonically with the noise in the signal. Very high signal costs are occasionally observed when the noise in the signal is large.

We explain our data with an attraction learning model that allows a mixture of belief learning (best response), imitation and reinforcement. Based on this model the anomaly that subjects separate in the low noise treatment where according to theory separation cannot be supported in equilibrium, can be intuitively understood as follows. With only a low amount of noise buyers initially behave as if there is no noise and thus use a higher cutoff than prescribed by equilibrium play. It is then in the interest of the sellers to choose separating signal costs. Buyers feel no pressure to change their disequilibrium behavior, because the noise in the signal smoothes their expected payoffs and their actual choices are not noticeably worse than their best responses. Therefore, separation does not unravel and Spence's original result is "saved" for behavioral reasons. Simulations based on the estimated learning model suggest that separation for low noise would also have resulted if subjects had initially played much closer to the pooling equilibrium than they actually did. In addition, the simulations indicate that the separating result is not due to the limited number of periods in the experiment. Interestingly, stable non-equilibrium results like ours also appear in other signaling games. For instance, Cai and Wang (2006) and Wang et al. (2010) find that senders consistently overcommunicate compared to the most informative equilibrium in cheap-talk games.

Although a large literature on signaling games exists, noisy signaling has received little attention up till now. The path-breaking paper here is by Matthews and Mirman (1983), who include noisy signals in a (signaling) model of limit pricing. Within this context they show that the introduction of noise has substantial implications for the equilibrium predictions. Another important contribution is by Carlsson and Dasgupta (1997), who propose using vanishing noise as an equilibrium selection device in signaling games without noise. We consider the effect of noise in Spence's original pure signaling model in which signaling constitutes a pure social waste. In this noisy signaling game separating and pooling equilibria coexist, in contrast to the noisy signaling games of Matthews and Mirman, and Carlsson and Dasgupta that only allow for separating equilibria. An important result is that only in the setup of Spence the existence of a separating equilibrium depends on the amount of (non-vanishing) noise in the signal. We are therefore able to study an issue that these earlier papers did not address, viz. how the occurrence of different types of equilibria (i.e. separating vs. pooling) varies with the level of noise. In the next section we will elaborate further on the distinctive features of our theoretical analysis.

Previous experimental papers on signaling have searched for empirical equilibrium selection devices (in a noise-free context). Miller and Plott (1985) investigated signaling in a rich market institution, where sellers chose prices as well as costly quality increments to the product. In markets where the signaling costs were relatively low, market outcomes tended to converge to the separating equilibrium. Usually, high-quality sellers started with inefficiently high signaling costs before they converged to the minimum level that distinguished them from the low-quality types. Brandts and Holt (1992) studied the predictive power of belief-based refinements in a game that modeled workers' choices for education and employers' subsequent hiring decisions. In early sessions, they found that play converged to the intuitive pooling equilibrium. Having studied the dynamics in the sessions, they were able to alter the parameters such that play tended to converge to the unintuitive pooling equilibrium. Cooper et al. (1997a, 1997b) investigated a limit pricing game where low-cost monopolists had incentives to deter entry by high-cost monopolists. Subjects started at their "myopic optima", which allowed entrants to infer the monopolist's actual type and to act accordingly. This encouraged high-type monopolists to pool with the low-cost types. If no pooling equilibrium existed, initial attempts at pooling were shattered and play converged to a separating equilibrium. None of these papers considered the possibility that noise in the signaling technology might profoundly affect how people play signaling games.⁵

⁵ There are, however, experimental papers that investigated noisy communication in other games. Güth et al. (2006) investigate the effects of noisy leadership in a sequential duopoly game. Aoyagi and Fréchette (2009) study collusion in a repeated prisoners' dilemma game where the opponent's past actions are imperfectly revealed in a noisy public signal. Feltovich et al. (2002) consider a signaling game experiment in which the receiver (in contrast to our setup) does not perfectly observe the message chosen by the sender, but also receives some exogenous noisy information.

Table 1
Payoffs of seller and buyer over action–state pairs.

	Seller is of bad type ($1 - p$)	Seller is of good type (p)
Buy	$b - m, -y$	$g - m, x$
Refrain	$-m, 0$	$-m, 0$

Remarks: The first (second) number in each cell refers to the seller's (buyer's) payoffs. $g > b > 0$ and $x, y > 0$ are parameters of the model. $m \geq 0$ denotes the signal costs chosen by the seller.

Jeitschko and Normann (2009) also conduct experiments on noisy signaling games. Unlike us, their setup closely follows the one of Carlsson and Dasgupta (1997) (and Matthews and Mirman, 1983) where signaling is not entirely wasteful. In their setup, only separating equilibria exist (for all levels of noise) and the issue of how subjects' inclination to separate varies with different levels of noise is void. Instead, their main focus is on how prior beliefs affect play in both noisy and deterministic games.

The remainder of our paper is organized in the following way. Section 2 provides a detailed description of the game and the theoretical analysis. Section 3 describes the experimental design and procedures. Section 4 presents the experimental results and Section 5 concludes.

2. Theory

2.1. The noisy signaling game

We consider a simple signaling game between an informed seller and an uninformed buyer. The seller can be of two types, either good or bad; $p \equiv \Pr(\text{good})$, with $0 < p < 1$, denotes the buyer's prior belief that the seller is of the good type. The seller first chooses his message $m \in [0, \infty)$ at signal costs m . The buyer then observes a noisy signal $z \in \mathbb{R}$, i.e. we assume that she observes m with some additive noise:

$$z = m + \sigma \cdot \varepsilon.$$

Here ε is a random variable with distribution function $F : \mathbb{R} \rightarrow [0, 1]$ and $\sigma \geq 0$ reflects a scaling parameter to account for changes in the amount of noise. In line with Matthews and Mirman (1983) and Carlsson and Dasgupta (1997), we assume that the density f is continuous and strictly positive everywhere. Moreover, like them we also assume that the conditional density of z given m satisfies the strict monotone likelihood ratio property (MLRP).⁶ Intuitively, this means that higher signal costs m become more likely when the observed signal z increases.

Having observed signal z , the buyer decides whether to buy or refrain from buying the product. The latter yields her 0, irrespective of the seller's type. If the buyer decides to buy she obtains a payoff equal to $x > 0$ when the seller is of the good type and $-y < 0$ when he is of the bad type. The seller always bears the signal costs of his message choice m . Apart from that, the good (bad) type seller obtains a gross payoff of g (b) from a sale. We assume that $g > b > 0$, i.e. the sorting condition is satisfied. Table 1 summarizes these payoffs. Both seller and buyer are assumed to be risk-neutral.

Note that our setup is isomorphic to the original one of Spence (1973) when we divide the seller's payoffs by his type $t \in \{b, g\}$; i.e. the seller obtains $1 - \frac{m}{t}$ when the buyer buys and $-\frac{m}{t}$ if she does not. In this alternative specification the two seller types do not differ in their benefits of a sale, but rather in their (marginal) costs of producing message m (which corresponds to the level of education in Spence's original formulation). Because this is just a normalization, equilibrium predictions are exactly the same.

2.2. Equilibrium analysis

From Spence (1973) it is well known that without noise ($\sigma = 0$) there are many Perfect Bayesian equilibria. Among these are pooling equilibria in which both seller types choose $m = 0$ and separating equilibria in which the bad type chooses $m = 0$ and the good type chooses some level of signal costs $m_g \in [b, g]$. All these equilibria exist independent of prior belief p . Moreover, when the buyer would buy in the absence of additional information – i.e., $p > \beta^* \equiv \frac{y}{x+y}$ – also pooling on any $m \in (0, b]$ can be supported as equilibrium.

Introducing noise by letting $\sigma > 0$ narrows down the equilibrium set considerably. First, pooling on some $m > 0$ cannot occur any longer. If both seller types choose the same signal costs m , then the buyer's posterior belief necessarily equals

⁶ These assumptions facilitate the equilibrium analysis, because they imply that the buyer necessarily uses a cutoff strategy in any non-pooling equilibrium (cf. Appendix A). The important assumption here is that each possible signal comes from each type with strictly positive probability. As long as this assumption remains fulfilled, the analysis may be generalized to the case where signals have bounded support. In applications where negative signals are problematic, one may then choose to normalize the game such that only positive signals are possible. Alternatively, negative signals may be natural in applications where the costs are not exclusively monetary. As Spence (1973, p. 359) already noticed: "Signaling costs are to be interpreted broadly to include psychic and other costs ...". Negative signals are also less problematic when one takes the seller's consumption good value of signaling into account: "The signal cost function does, in principle, capture education as a consumption good, an effect that simply reduces the cost of education" (cf. Spence, 1973, p. 364).

her prior for any signal z observed. Her buying decision is then fully determined by her prior belief and independent of the signal received. Given this, the seller lacks any incentive to (stochastically) increase the signal and therefore only pooling on $m = 0$ can occur. Second, adding noise also severely restricts the set of separating equilibria. We illustrate this by focusing on the case considered in the experiment where F equals the standard normal distribution $N(0, 1)$ (which will be denoted by Φ). In Appendix A we show that all the results discussed here generalize to any distribution function F that satisfies the assumptions made in the previous subsection. Proofs of propositions are relegated to this appendix as well.

We first consider pure strategy equilibria before we deal with mixed strategy equilibria. Proposition 1 below characterizes the set of separating equilibria that may (but not necessarily do) exist besides pooling on $m = 0$.

Proposition 1. *Let $F = \Phi$ and assume that players are restricted to use pure strategies. (i) A pooling equilibrium in which both seller types choose $m = 0$ always exists. In this equilibrium the buyer never [always] buys when $p < [>]\beta^* \equiv \frac{y}{x+y}$. Pooling on some $m > 0$ cannot occur. (ii) Generically, i.e. for all $p \neq \frac{\beta^* \cdot g}{(1-\beta^*) \cdot b + \beta^* \cdot g}$, it holds that in any separating equilibrium the bad type of seller chooses $m = 0$ whereas the good type chooses some positive level of signal costs $m = m_g > 0$. The buyer buys if $z > z^*$ and refrains from buying otherwise. Signal costs m_g and cutoff signal z^* are given by:*

$$m_g = z^* + \sqrt{2\sigma^2 \cdot \left[\ln\left(\frac{g}{\sigma\sqrt{2\pi}}\right) \right]} \tag{1}$$

$$z^* = z_h^* \equiv \sqrt{2\sigma^2 \cdot \left[\ln\left(\frac{g}{\sigma\sqrt{2\pi}}\right) - \ln\left(\frac{p(1-\beta^*)}{(1-p)\beta^*}\right) \right]} \quad \text{if } p \leq \beta^* \tag{2}$$

$$z^* = z_h^* \quad \text{or} \quad z^* = z_l^* \equiv -z_h^* \quad \text{if } p > \beta^* \tag{3}$$

Note that Proposition 1 only characterizes the set of pure strategy equilibria. Conditions under which the separating equilibria indeed do exist will be discussed shortly.

The intuition behind the separating equilibria of Proposition 1 is as follows. Given that the noise distribution satisfies MLRP, the buyer necessarily uses a cutoff strategy; the seller's product is bought if and only if a signal larger than some cutoff level z^* is observed. For a given value of z^* , the equilibrium level of signal costs m_g the good type seller chooses then follows from equalizing the marginal benefits of raising m with the marginal costs (equal to one) of doing so. This yields expression (1).⁷ The exact value of cutoff z^* subsequently follows from the requirement that the buyer's posterior belief after observing z^* should make her indifferent between buying or not. Because posterior beliefs are determined by Bayes' rule everywhere, this requirement puts some strong characterizing restrictions on the player's equilibrium strategies. Expressions (2) and (3) follow from these. As the latter expression makes clear, for $p > \beta^*$ there are two solutions for z^* , so actually two separating equilibria may potentially exist side by side.

Our next proposition concerns the actual existence of a separating equilibrium. In this regard we are particularly interested in how the amount of noise – as reflected by parameter σ – affects existence.

Proposition 2. *Let $F = \Phi$ and assume that players are restricted to use pure strategies. A necessary condition for a separating equilibrium to exist is that $\sigma \leq \frac{g}{\sqrt{2\pi}} \cdot \min\left\{\frac{(1-p)\beta^*}{p(1-\beta^*)}, 1\right\} \equiv \bar{\sigma}$. Assuming $\sigma \leq \bar{\sigma}$, it holds that:*

- (i) $p \leq \beta^*$: a separating equilibrium does not exist if σ becomes sufficiently small;
- (ii) $\beta^* < p < \frac{\beta^* \cdot g}{(1-\beta^*) \cdot b + \beta^* \cdot g}$: a separating equilibrium always exists. For this equilibrium it holds that $\lim_{\sigma \downarrow 0} m_g = 0$;
- (iii) $p > \frac{\beta^* \cdot g}{(1-\beta^*) \cdot b + \beta^* \cdot g}$: a separating equilibrium does not exist.

Two main observations follow from Proposition 2. First, when there is a lot of noise separation cannot occur; for $\sigma > \bar{\sigma}$ separating simply becomes too difficult or too costly for the good type seller. Formally this can be understood from expressions (1) and (2). The terms within square brackets become negative for σ sufficiently large and no sensible solutions for m_g and z^* exist.

Second, a separating equilibrium also fails to exist when only a small amount of noise is introduced in the unfavorable prior case $p \leq \beta^*$. A priori the buyer then refrains from buying, but observing a small positive signal would already induce her to change her mind. (Formally, cutoff z_h^* as given by (2) is low when σ is small.) But given that the buyer is persuaded so easily, the bad type seller may want to deviate from $m = 0$. Similar to (1) above, his best candidate deviation level equals:

$$m_b = z^* + \sqrt{2\sigma^2 \cdot \left[\ln\left(\frac{b}{\sigma\sqrt{2\pi}}\right) \right]} \tag{4}$$

⁷ In fact, this expression incorporates the second order condition as well; from the SOC it follows that m_g necessarily exceeds z^* .

Table 2
Overview of separating equilibria.

	Amount of noise σ							
	≤ 20	25	40	75	100	120	140	≥ 145
z_h^*	–	53.14	75.66	114.26	132.12	141.02	145	–
m_g	–	101.28	142.21	206.43	228.80	231.62	216.63	–

Remarks: This table is based on the parameter values used in the experiment: $p = 0.5$, $g = 400$, $b = 90$, $x = 300$ and $y = 450$ (so $\beta^* = 0.6$). In the experiment we only consider four different values of σ , viz. 0, 10, 40 and 120. A dash implies that a separating equilibrium does not exist.

One requirement for a separating equilibrium to exist is thus that $m = m_b$ should yield the bad type seller weakly less than $m = 0$. This reduces to⁸:

$$b \cdot \left(\Phi\left(\frac{z^*}{\sigma}\right) - \Phi\left(\frac{z^* - m_b}{\sigma}\right) \right) \leq m_b \tag{5}$$

Because this condition depends on Φ , no closed form expression for the cutoff value on σ can be obtained. But it can be shown that it is necessarily violated for σ small enough.⁹

The prediction that just a small amount of noise destroys separation is a priori somewhat counter-intuitive. To better understand the underlying driving force, let the payoff parameters be such as in the experiment (where $p = 0.5 < 0.6 = \beta^*$, see Table 2) and consider first the least cost separating equilibrium of the no noise game. In this so-called Riley outcome the bad type seller chooses $m = 0$ and the good type $m = 90 (= b)$. The buyer buys only if a signal of $z = 90$ or higher is observed. Fig. 1(a) reflects both the densities of the signals generated by the two seller types (labeled f_B and f_G) and the buyer's posterior belief $\beta(z)$. The signal densities are degenerate at $z = 0$ and $z = 90$, respectively. This implies that $\beta(z)$ is determined by Bayes' rule only for these two values of z . To support the equilibrium, out-of-equilibrium beliefs must be such that $\beta(z) \leq \beta^*$ for $z < 90$. The figure depicts the equilibrium where $\beta(z) = 0$ in that case. Importantly, for any signal between 0 and 90 the buyer may hold skeptical beliefs that the signal quite likely came from the bad type seller, even when it is close to 90.

Now consider what happens if some noise is added, of size $\sigma = 10$ say. Naively one would then expect that the good type seller simply moves up his signal cost a bit, to 100 say, while the bad type stays put. Fig. 1(b) depicts this situation. In contrast to the no noise case, the two signal densities are now non-degenerate and posterior beliefs $\beta(z)$ for in between signals $0 < z < 100$ now follow from Bayes' rule. Given the highly concentrated signal densities, signals closest to 0 most likely come from the bad type while signals closest to 100 most likely come from the good type. The actual cutoff for which $\beta(z^*) = \beta^*$ is roughly in the middle at $z^* \approx 50.4$. But if the buyer uses this cutoff value, the bad type seller wants to jump away from $m = 0$ towards $m_b \approx 66.4$ (cf. expression (4)). This destroys the separation outcome. Another destabilizing factor is that for $z^* \approx 50.4$ the good type seller also wants to deviate from 100 to $m_g \approx 73.9$.

From the above it follows that crucial for separation to unravel when some noise is introduced, is that the buyer realizes that she should set a (much) lower cutoff level above which she decides to buy. If she does not do so and keeps the cutoff at $z^* = 90$ (or somewhat higher), the two seller types still have an incentive to separate.

The overall comparative statics in σ are illustrated in Table 2. For σ roughly below 24 separation cannot occur. The same holds for high amounts of noise (viz. σ roughly above 142) whereas for in between levels a separating equilibrium exists besides the pooling one.¹⁰ In these separating equilibria the level of signal costs chosen by the good type increases with noise (for σ not too high). This level m_g can actually become quite large relative to g and b , as the case with $\sigma = 120$ illustrates. The good type thus may be 'forced' to use very costly signals.

Mixed strategy equilibria. Up till now we have assumed that seller and buyer may only use pure strategies. Theorem 3 in Appendix A shows that if we allow them to use mixed strategies as well, two additional types of equilibria may potentially exist as well. In the first type of mixed equilibrium the bad type seller mixes between $m = 0$ and signal costs m_b as given in (4), whereas the good type chooses m_g from expression (1) for sure. The bad type's mixing probability $q_b \equiv \Pr(m = m_b)$ follows from the requirement that observing z^* should make the buyer indifferent (i.e. $\beta(z^*) = \beta^*$). For the bad type to be willing to mix, condition (5) now has to hold with equality. Given the appearance of Φ here, the resulting equilibrium

⁸ The mirror image requirement for the good type seller is that he should not have an incentive to deviate from m_g towards $m = 0$, i.e. $g \cdot (\Phi(\frac{z^*}{\sigma}) - \Phi(\frac{z^* - m_g}{\sigma})) \geq m_g$ is needed. This actually has a bite in case (i) of Proposition 2. There a separating equilibrium may not exist when σ becomes large within the relevant interval $(0, \bar{\sigma}]$.

⁹ Also in part (ii) where $\beta^* < p < \frac{\beta^* \cdot g}{(1 - \beta^*) \cdot b + \beta^* \cdot g}$ the separating equilibrium based on z_h^* vanishes for σ sufficiently small. In that case, however, an equilibrium based on $z_l^* < 0$ always exists (given that $\sigma \leq \bar{\sigma}$). In this equilibrium the buyer a priori intends to buy and even receiving a moderately negative signal does not lead her to behave differently. The bad type seller therefore does not have an incentive to deviate from $m = 0$. Part (iii) of Proposition 2 can be intuitively understood from considering how cutoff level z^* (either equal to z_l^* or z_h^*) varies with prior belief p . Because z^* is close to zero when p is high, the bad type seller obtains a strong incentive to deviate from $m = 0$, i.e. no-deviation condition (5) is violated. This upsets the separating equilibria.

¹⁰ Note that standard belief-based equilibrium refinements (like the intuitive criterion of Cho and Kreps, 1987) are ineffective in our setup, because for $\sigma > 0$ there are no out-of-equilibrium beliefs.

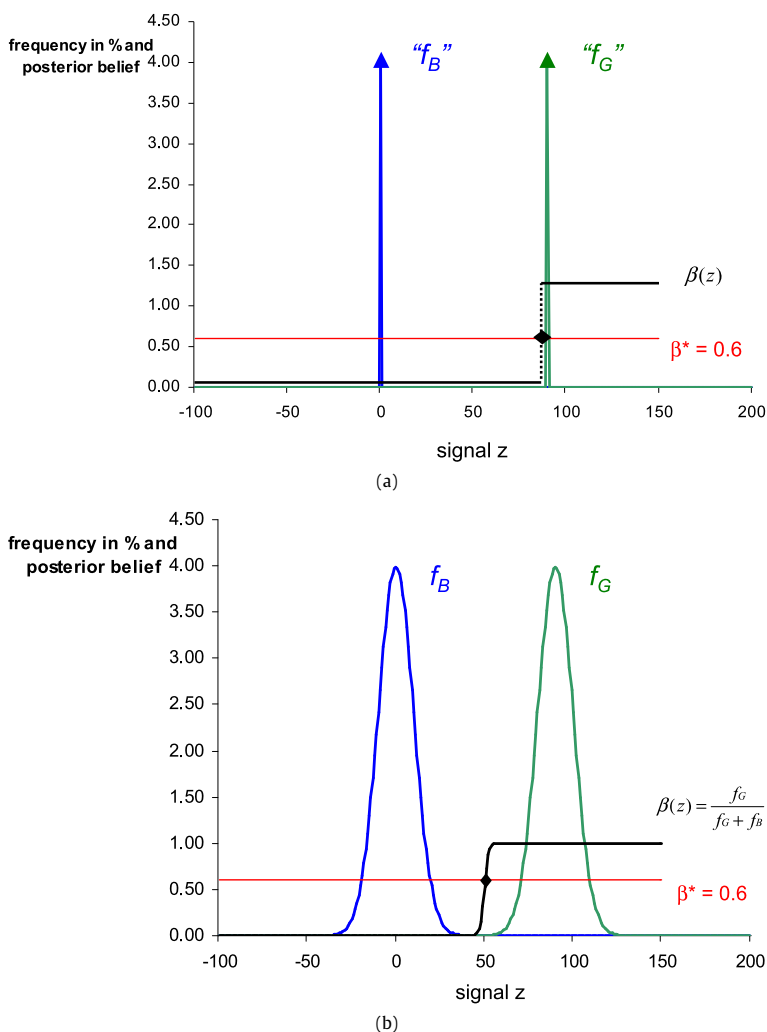


Fig. 1. Signal densities and posterior beliefs in (a) the no noise case and (b) the $\sigma = 10$ case.

Table 3

Overview of mixed strategy equilibria.

σ	Bad type mixes				Good type mixes		
	z_h^*	q_b	m_b	m_g	z_h^*	q_g	m_g
1	88.98	0.15	89.66	90.16	396.53	≈ 0	399.71
5	77.95	0.15	87.88	91.11	385.14	≈ 0	398.30
10	69.06	0.15	85.05	92.59	372.74	≈ 0	396.28
40	–	–	–	–	314.21	≈ 0	380.75
120	–	–	–	–	200.35	0.25	290.95

Remarks: This table is based on the parameter values used in the experiment: $g = 400$, $b = 90$, $x = 300$, $y = 450$ and $p = 0.5$. In the experiment we only consider the four different values of σ of 0, 10, 40 and 120. A dash implies that a mixed strategy equilibrium does not exist for the given value of σ .

values of z^* , m_b and m_g have to be calculated numerically. Table 3 provides an overview for some relevant parameter values. For the values of σ considered in the experiment, this mixed equilibrium appears to exist for $\sigma = 10$ only.¹¹ It exists for smaller amounts of noise as well. In fact, it holds that m_g and m_b tend to b as σ tends to 0 and that $\lim_{\sigma \downarrow 0} q_b > 0$. When the noise vanishes this equilibrium thus converges to a mixed equilibrium of the no noise game that is *insufficiently revealing*; with strictly positive probability the buyer takes the opposite decision of what she would do under complete information.¹²

¹¹ When σ becomes too large, the defining equation of q_b results in a negative value.

¹² In this particular case the buyer always buys if $z = b$ is observed, although with positive (but small) probability the seller is of the bad type.

In the other mixed strategy equilibrium the good type seller mixes between $m = 0$ and $m = m_g$, while the bad type chooses $m = 0$ for sure. For the good type to be indifferent it now must hold that $g \cdot (\Phi(\frac{z^*}{\sigma}) - \Phi(\frac{z^* - m_g}{\sigma})) = m_g$, so also here closed form expressions cannot be obtained. What can be shown theoretically, however, is that this equilibrium converges to a pooling (on $m = 0$) equilibrium when the noise becomes small, i.e. $\lim_{\sigma \downarrow 0} q_g = 0$. Moreover, signal costs m_g become large for low values of σ : $\lim_{\sigma \downarrow 0} m_g = \lim_{\sigma \downarrow 0} z^* = g$. The parametric examples in Table 3 illustrate this.

The first mixed strategy equilibrium discussed above challenges the prediction that for low values of σ separation will not occur. In the other mixed equilibrium the positive signal costs m_g decrease with σ . Because these predictions run counter to the main predictions based on pure strategy equilibria only, we explicitly check in our empirical analysis whether actual behavior is consistent with these mixed strategy equilibria.

2.3. Related theoretical literature

Matthews and Mirman (1983) consider a limit pricing model with a potential entrant and an incumbent monopolist who is privately informed about industry conditions. The actual price the incumbent charges depends on both his output choice and a random demand shock that occurs after the output decision. Owing to this noise, equilibria are obtained that differ from standard signaling game equilibria in three ways: (i) there is a great reduction in the number of equilibria, (ii) (separating) equilibrium strategies now directly depend on prior beliefs, and (iii) different amounts of information are revealed in different separating equilibria, leading to richer comparative statics.¹³

These three features apply in our setting as well. First, with noise (generically) only five different equilibria exist at most, as opposed to the continuum of equilibria in the no noise case. For a range of parameter values (e.g. p sufficiently high) the equilibrium is even unique. Second, as expressions (1) through (3) reveal, separating equilibrium strategies directly vary with prior belief p . For instance, cutoff value z_h^* decreases with p , implying that the buyer is more easily persuaded to buy if she is already more inclined to do so a priori. This seems a much more intuitive prediction than the irrelevance of p for the required level of separation that the no noise case predicts. Similar remarks apply to variations in g . Third, with noise even a separating equilibrium is insufficiently revealing. Different separating equilibria may therefore lead to different amounts of information being revealed. In our setting two different separating equilibria may actually exist side by side when the prior is favorable. In the one based on z_h^* the signal costs that the two seller types choose are more dispersed and the buyer obtains more information than in the one based on z_l^* (cf. Proposition 1).

Carlsson and Dasgupta (1997) focus on equilibrium selection in signaling games without noise, by studying the limiting set of 'noise-proof' equilibria that results from letting the noise vanish. Among other things, they show that every noise-proof equilibrium of the no noise game is necessarily insufficiently revealing. Our theoretical analysis replicates this finding for our setup. In particular, by letting σ go to zero Theorems 2 and 3 in Appendix A show that there exist only two noise-proof equilibria in the no noise game, viz. pooling on $m = 0$ and the mixed equilibrium described before in which only the bad type mixes between $m = 0$ and $m_b = b$. Both are insufficiently revealing, as with positive probability the buyer takes a decision that she will regret ex post.¹⁴

Despite these similarities, our theoretical findings differ in some other important respects from these two earlier studies. In both Matthews and Mirman (1983) and Carlsson and Dasgupta (1997) the equilibria of the noisy signaling games they study are always separating.¹⁵ In contrast, our setup allows for pooling equilibria as well. At the same time, only in our game the existence of a separating equilibrium is not guaranteed and depends on the amount of noise in the signal. Unlike these previous authors, therefore, we are able to study the question of how the existence of different types of equilibria (and thus different types of equilibrium strategies) varies with noise.

The main driving force why we obtain results that differ from both Matthews and Mirman (1983) and Carlsson and Dasgupta (1997) is that in these earlier papers the seller's (message) choice actually serves two purposes. Besides a pure signaling function geared towards influencing the behavior of the receiver, the seller's choice also allows him to optimally

¹³ The introduction of noise may have profound implications in complete information contexts as well. Bagwell (1995) studies a noisy leader game in which a follower observes the actual choice a leader made with some noise. The striking result he obtains is that the standard first-mover advantage then completely disappears. As pointed out by van Damme and Hurkens (1997), this conclusion depends on the restriction to pure strategies. When the noise is small, there always exists a mixed strategy equilibrium that approximates the standard Stackelberg outcome of the game without noise. Also the assumption of complete information is crucial. Maggi (1999) shows that the value of commitment reappears when the leader's choice is based on private information that is payoff-irrelevant for the follower. The latter assumption allows Maggi to abstract away from signaling considerations, which are the main focus of this paper.

¹⁴ Some other main findings of Carlsson and Dasgupta (1997) do not carry over to our setting. For example, by Proposition 1(a) it follows that pooling on $m = 0$ is noise-proof. But for $p < \beta^*$ this equilibrium does not survive the never-a-weak-best response (NWBR) refinement in the game without noise. Unlike in Carlsson and Dasgupta (1997), therefore, in our setup not every noise-proof equilibrium satisfies NWBR. Likewise, we do not find a unique noise-proof equilibrium whereas in Carlsson and Dasgupta it "often" is.

¹⁵ The same holds for the noisy signaling model used by Kanodia et al. (2005) to study a firm's optimal investment in the presence of capital market imperfections. Calveras (2003) embeds a noisy signaling game in a model of a bank that can manipulate the noisy information a regulator observes. His signaling subgame allows for both pooling and separating equilibria (see his Proposition 3). Apart from his model being much more specific than ours, Calveras does not study how the existence of the different types of equilibria varies with the amount of noise nor does he consider mixed strategy equilibria. The latter also applies to Hertzendorf's (1993) analysis of noisy advertising in the multi-dimensional signaling model of Milgrom and Roberts (1986). The purpose of his study is to show that in the presence of noise, no separating equilibrium exists in which prices and advertising are simultaneously used as informative signals.

adapt to changing circumstances. In Matthews and Mirman this adaptation purpose for example follows from the fact that the output choice (message) of the incumbent monopolist (seller) varies with industry conditions (type), even when the receiver (potential entrant) is fully informed on these industry conditions. Optimal output is higher the more favorable industry conditions are. The seller's output choice therefore does not perform a pure signaling function alone. The same applies for the setup in Carlsson and Dasgupta (1997).¹⁶ The important consequence of this two-folded purpose is that (in the noisy games) the seller's best response correspondence is *strictly* monotonic in his type (cf. Proposition 3.1 in Carlsson and Dasgupta, 1997). Loosely put, the two purposes together pull sellers towards separation.

In our setup the seller's message choice only serves a pure signaling function, such that (costly) signaling constitutes a pure social waste.¹⁷ If in our game the seller's type would be public information, the seller's ('message') choice would be independent of his type and equal $m = 0$ for both types.¹⁸ As a result, sellers' best responses are only *weakly* monotonic in types (cf. the proof of Lemma 1 in Appendix A). The pull towards separation is therefore much weaker. This implies in turn that our setup allows for pooling equilibria as well while at the same time the existence of a separating equilibrium is not guaranteed.

3. Experimental design and procedures

The computerized experiment was run at the University of Amsterdam where subjects were recruited from the student population. Subjects read the on-screen instructions at their own pace. At the end of the instructions, subjects had to answer some test questions correctly before they could proceed. They also received a summary of the instructions on paper.¹⁹ Subjects knew that the experiment consisted of two parts. Part 1 lasted for 40 periods. In part 2 subjects formulated a strategy that automatically determined their play in another 10 periods. Subjects received instructions for the second part only after the first part was finished. At the beginning of the experiment, subjects received a starting capital of 5000 points. Their period earnings (losses) were added to (subtracted from) this starting capital. At the end of the experiment, points were exchanged into euros at a rate of 1 euro for 250 points. In 1.5 to 2 hours, a total of 184 subjects earned on average 37.05 euros with a standard error of 10.89.

At the start of the first part subjects were assigned to the role of seller or the role of buyer. Throughout the whole experiment subjects kept the same role. Each period, sellers and buyers were randomly matched in pairs within a (fixed) matching group of 8 subjects. Subjects knew that they were never matched with the same subject twice in a row. In most sessions, we ran 2 matching groups simultaneously. At the start of a period, the seller was privately informed of the quality of his product. In each matching group, 2 products had high quality and 2 products had low quality. Thus, the prior probability of a high-quality product was 0.5, a fact that was communicated to all subjects. The quality of the product of a seller in a given period was independent of the quality of his product in another period. After observing the quality of his product, the seller chose a signal cost, an integer amount between 0 and 400. The computer added a noise term, an independent draw from a $N(0, \sigma^2)$ distribution to the signal cost, and communicated the resulting signal, but not the signal cost nor the noise term, to the buyer. We communicated the density of the normal distribution with the help of a figure and some explanatory remarks about symmetry and confidence intervals. The buyer decided whether or not to buy the product, after which the payoffs of the pair were determined. The payoff table was common information to the subjects.

At the end of a period, both players were informed of the quality, the signal cost and the signal. In addition, subjects could view a social history window at the bottom of the screen that showed the results of all pairs in their own matching group for the last 10 periods. For buyers, the screen was ordered on signal (from high to low), quality (from high to low), signal cost (from high to low) and buy-decision (from yes to no), respectively. Subjects could recognize their own previous results as these were printed against a different, light-gray background. Fig. 2(a) shows a snapshot of this window. Fig. 2(b) shows the social history window for sellers, which was ordered on signal cost (from high to low), signal (from high to low), buy-decision (from yes to no) and quality (from high to low), respectively. We provided this information because it allows subjects to learn faster. Our paper deals with the topic of equilibrium selection and the comparative statics of the equilibrium predictions with respect to the amount of noise. To address these issues play must converge to equilibrium in the first place. We think that, compared to the world outside of the laboratory, the superior information provided to subjects balances their lack of experience with the game. A similar social history (on black board) was first provided in a signaling

¹⁶ Let π denote the probability with which the buyer buys and $u(t, m, \pi)$ the sender's utility. Carlsson and Dasgupta (1997) assume that $\frac{\partial u(t, m, \pi)}{\partial m} = 0$ has a unique (interior) solution $m_\pi(t)$, with $m_\pi(t)$ strictly increasing in t (cf. Assumptions (U3) through (U5) on p. 443). Note that in our setup we have $u(t, m, \pi) = t \cdot \pi - m$, so $\frac{\partial u(t, m, \pi)}{\partial m} = -1$ and this assumption is not met.

¹⁷ By focusing on costly signaling our setup differs from "cheap talk" games in which messages are costless. See Blume et al. (2007) for an analysis of the impact of introducing noise on the amount of information transmission in the cheap talk model of Crawford and Sobel (1982). Landeras and Pérez de Villarreal (2005) introduce noise into a screening model in which the *uninformed* party moves first.

¹⁸ Our different results are thus not due to the fact that in our setup the marginal costs of raising m are independent of the seller's type, as one a priori might have expected. As explained in the main text, our setup is completely isomorphic to the case where seller's utility equals $u(t, m, \pi) = \pi - \frac{m}{t}$, with π the probability with which the buyer buys. Just as in Carlsson and Dasgupta (1997), in this specification marginal costs are type dependent.

¹⁹ The instructions are available at journal's web site as online supplementary material.

row number	signal	quality	signal cost	buyer buys?
1	142	HIGH	146	YES
2	142	LOW	118	YES
3	98	HIGH	101	NO
4	-2	LOW	5	NO

(a)

row number	signal cost	signal	buyer buys?	quality
1	142	146	YES	HIGH
2	142	118	YES	HIGH
3	98	101	NO	LOW
4	5	-2	NO	LOW

(b)

Fig. 2. Social history window (a) buyers (example) and (b) sellers (example).

Table 4

Experimental design.

Treatment	σ	# of matching groups	# Subjects per matching group
σ_0	0	4	8
σ_{10}	10	5	8
σ_{40}	40	6	8
σ_{120}	120	8	8

Remarks: Per period payoffs for the subjects are given in Table 1, with $b = 90$, $g = 400$, $x = 300$, $y = 450$ and $p = 0.5$.

experiment of Miller and Plott (1985), who introduced it in the later sessions to help subjects recognize the relationship between types and choices.²⁰

The variance of the error term in the signal (σ^2) was the only variable that varied between the 4 treatments. We refer to the treatments as σ_0 , σ_{10} , σ_{40} and σ_{120} . Table 4 summarizes the details of the experimental design. Each subject participated in one treatment only. We correctly anticipated that behavior would become more volatile for higher noise levels in the signals. Therefore, we decided to collect a larger number of observations for the treatments with the higher noise levels.

In the second part of the experiment, we asked subjects to formulate a strategy for periods 41–50. Buyers were asked to provide a cutoff level for the signal received, at and above (below) which they would (not) buy the product. They could also indicate that they would never or always buy the product, independent of the signal. Sellers were asked to choose a signal cost for high-quality products as well as for low-quality products. We explicitly mentioned that it was up to the seller to decide whether he wanted to choose the same signal costs for high and low quality or different amounts. We emphasized that otherwise the game was exactly the same as the one they played in the first 40 periods. When all subjects had chosen their strategies, the computer automatically played out the final 10 periods.

4. Experimental results

We present the experimental results in two parts. In Section 4.1, we deal with the results at the aggregate level. We show how actual signal costs vary with the noise level and we investigate whether the data come closer to the pooling equilibrium or the separating equilibrium in the treatments where the equilibria coexist. Here, we focus on the data of the second half of part 1 (periods 21–40) when subjects had become familiar with the game. It turns out that the results for the strategy method confirm the results of part 1 to a large extent.²¹ We have chosen to limit the report of the results on

²⁰ Other papers have used role reversion to accomplish this. After senders have become receivers, it becomes easier for them to interpret the meaning of a signal (e.g., Brandts and Holt, 1992).

²¹ Brandts and Charness (forthcoming) provide a survey of methodological work on the strategy method and conclude that the strategy method does not affect results in a qualitative sense. Our results are in agreement with their main conclusion.

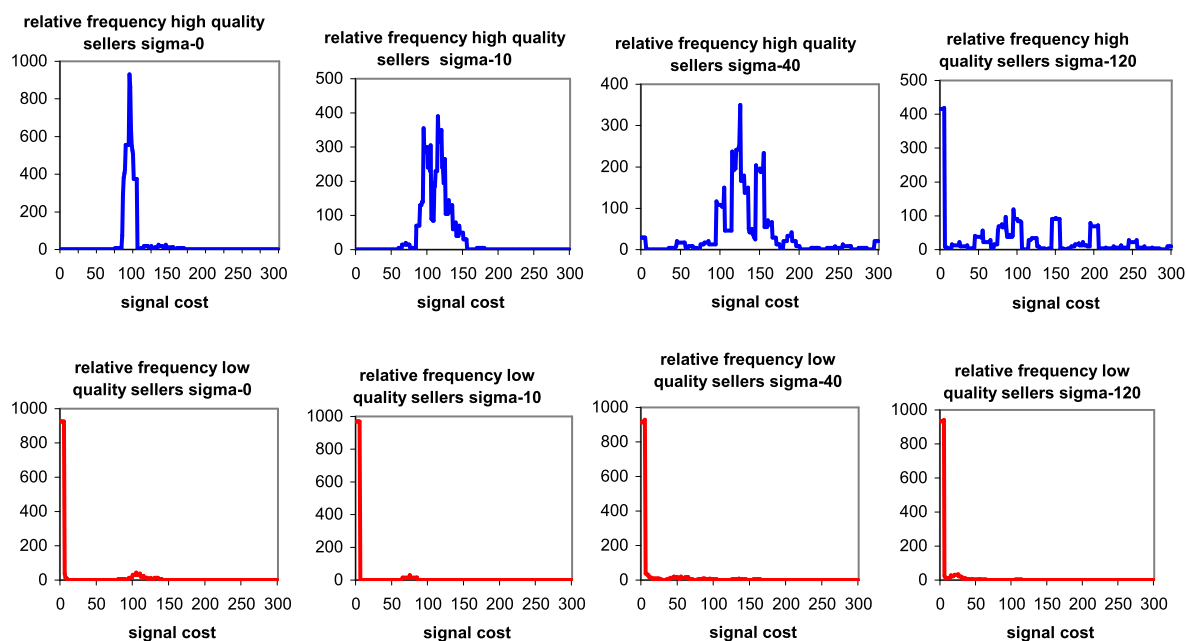


Fig. 3. Running histograms for periods 21–40; for each signal cost the relative frequency of cases in the interval [signal cost – 5, signal cost + 5] is displayed.

the strategy method to the extent that they provide additional insight. In Section 4.2, we zoom in on the behavior of our subjects and provide a coherent explanation of the main results. Here we use the data of the whole experiment.

4.1. Aggregate overview main results

First we take a look at how sellers behaved. Fig. 3 presents histograms of the signal costs chosen by high-quality sellers and low-quality sellers. In treatment σ_0 , high-quality sellers most often chose a signal cost of 100, followed by 91. A large majority of 93.2% of submitted signal costs lied between 90 and 100. A total of 88.8% of the signal-costs submitted by low-quality sellers in σ_0 equaled exactly 0. In this treatment, sellers' behavior provides clean evidence for separation.

Also in treatment σ_{10} low-quality sellers overwhelmingly chose 0, while high-quality sellers chose to separate. The latter tended to send higher signal costs, also at a higher variance than in σ_0 : 86.5% of the submitted signal costs were in between 90 and 130. This pattern of higher and more volatile signal costs extends to treatment σ_{40} , where high-quality sellers separated with 79.6% of the signal costs lying between 100 and 160, while low-quality sellers stuck to 0. The picture looks very different for treatment σ_{120} , though. With 40.9% of the high-quality sellers choosing a signal cost of exactly 0, the focus of sellers' attention seemed to be on pooling. Still a considerable fraction of 41.6% of the signal costs was at 90 or above. High signal costs were very spread out. In contrast, low-quality sellers by and large chose a signal cost of 0, like in the other treatments.

Very high signal costs were observed occasionally. In treatments σ_{40} and σ_{120} , high-quality sellers chose signal costs of at least 230 in 4.2% and 5.6% of the cases, respectively. In the other treatments such high signal costs were never observed.

A key prediction of the theoretical analysis is that conditional on the existence of a separating equilibrium, the signal costs should increase with noise as long as subjects coordinate on separating. Table 5 splits the submitted signal costs into pooling (below 90) and separating (at and above 90), for low-quality sellers as well as high-quality sellers.²² Conditional on a separating signal cost being sent, the signal cost indeed increased monotonically from on average 97.9 in σ_0 to 158.3 in σ_{120} . A Kruskal–Wallis test performed at the group level rejects the hypothesis that there are no differences in signal costs between the treatments ($p = 0.02$).²³

The table foreshadows the main finding regarding subjects' inclination to separate. The relative frequency of pooling signals increases with noise. In treatments σ_0 , σ_{10} , σ_{40} and σ_{120} , high-quality senders submitted pooling signal costs in 0.6%, 2.0%, 7.5% and 58.4% of the cases, respectively. In all treatments the overwhelming majority of low-quality sellers submitted signal costs of 0. Thus, sellers aimed for separation in treatments σ_0 , σ_{10} and σ_{40} , while the results are mixed for treatment σ_{120} , where the pooling equilibrium attracted sellers more than the separating equilibrium did.

²² This classification follows from the observation that in a separating equilibrium (if it exists), the good type seller always chooses a signal cost that exceeds 90 (see the case $p < \beta^*$ in Table 2). Positive signal costs below 90 thus cannot be interpreted as an attempt to separate and are therefore labeled pooling.

²³ Mann–Whitney tests reveal that all pair-wise comparisons of signal costs between treatments are significant or weakly significant, except for the comparison between σ_{40} and σ_{120} .

Table 5
Pooling and separating signal costs (SC) in periods 21–40.

	Pooling: SC < 90			Separating: SC ≥ 90			Predicted separating eq.
	n	Actual	Emp. best response	n	Actual	Emp. best response	
Low type							
σ_0	149	0.7 (7.0)	2.1 (13.0)	11	112.3 (11.7)	0.0(0.0)	0.0
σ_{10}	200	2.3 (12.9)	0.0 (0.0)	–	–	–	0.0
σ_{40}	235	2.8 (11.7)	0.0 (0.0)	5	124.0 (28.8)	0.0(0.0)	0.0
σ_{120}	319	1.6 (7.0)	0.0 (0.0)	1	110.0 (0.0)	0.0(0.0)	0.0
High type							
σ_0	1	80.0 (0.0)	80.0 (0.0)	159	97.9 (11.4)	93.8 (4.4)	90.0
σ_{10}	4	70.0 (4.1)	112.5 (15.0)	196	113.6 (15.9)	117.9 (11.6)	48.7 [*]
σ_{40}	18	38.3 (33.7)	149.6 (16.2)	222	140.5 (41.5)	156.3 (17.5)	142.2
σ_{120}	187	17.7 (29.9)	190.8 (73.8)	133	158.3 (55.5)	187.4 (62.5)	231.6

Remarks: n gives the number of observations. Standard deviations (based on individual observations) are in parentheses.

^{*} When $\sigma = 10$ only pooling on zero is a Nash equilibrium; a choice for 48.7 by the high type is a best response given that the low type chooses a signal cost of zero (but not vice versa).

Table 5 also reveals two differences between equilibrium predictions and submitted signal costs. The first difference is that the submitted signal cost increased on average monotonically with noise, while theoretically a separating equilibrium does not exist in treatment σ_{10} . The second difference is that high-quality sellers' signal costs did not increase as rapidly with noise (from σ_{40} to σ_{120}) as equilibrium predicts.

It only makes sense for sellers to play equilibrium when buyers play equilibrium. Therefore, a more relevant question is to what extent sellers chose best responses to the actual behavior of the buyers. We describe the latter by means of an empirical cutoff value \hat{z}^* . In particular, we determined in each matching-group and each period which cutoff level \hat{z}^* for the buyers minimized the sum of the buyers' errors against that cutoff level in the 10 most recently completed periods of that particular group.²⁴ We subsequently set the seller's empirical best response equal to the signal cost that maximized expected payoff given this cutoff level \hat{z}^* of the buyers.²⁵ Often, there was a range of cutoff levels that fitted the data equally well. In those cases, we determined the best response to the maximum cutoff level in the optimal range and the best response to the average cutoff level in the optimal range. It turns out that the best response on the basis of the maximum cutoff level was closer to the actual signal cost than the one based on the average. Therefore, we report statistics based on the maximum.

Table 5 includes a column that reports the sellers' best responses. In agreement with the actual data, the best responses of high-quality sellers increased monotonically with noise, with lesser increments than the equilibrium-predictions. Note that the high-quality sellers' best response in treatment σ_{10} equaled 117.9, quite close to the actual data, despite the fact that a separating equilibrium does not exist here. We will come back to this important finding in Section 4.2.

To assess whether subjects coordinated on a pooling or a separating equilibrium, buyers and sellers' behavior have to be scrutinized simultaneously. First, we deal with the possibility that subjects played in accordance with the logic of a mixed equilibrium. As explained in Section 2, there are two types of mixed equilibria. The one where the good type mixes has two features that are incompatible with the data. The first one is that the comparative statics prediction is violated. According to this equilibrium, the positive signal cost chosen by the high-quality seller should decrease with (increasing) noise, while it actually increased. Second, in σ_{120} , the good type should mix between 0 with probability 0.75 and 290.95 with probability 0.25. High-quality sellers submitted signal costs higher than 250 in only 2.2% of the cases, however. In all other treatments, the mixed equilibrium is observationally indistinguishable from pure pooling equilibria and we will deal with those later.

The next type of mixed equilibrium is the one where the low-quality seller mixes. This equilibrium exists in σ_{10} , but not in σ_{40} nor in σ_{120} (see Table 3). In σ_{10} , in 6 out of 200 cases low-quality sellers chose a signal cost larger than 0, three times 70 and three times 80. According to the mixed equilibrium, bad type sellers should choose a signal cost 85 with probability 0.15 and 0 otherwise. So the positive signal costs chosen are below the theoretically expected level. Moreover, the proportion of positive signal costs (3%) falls considerably short of the theoretically expected level (15%). In fact, a binomial test rejects the hypothesis that the proportion of positive observations is in line with the theoretical prediction ($p = 0.00$, test performed at the choice level). Choosing positive signal costs was not a great idea for bad type sellers: 5 out of these 6 positive signal costs led to signals above the equilibrium cutoff level of buyers, but only 3 actually led to a sale. Thus, sending these positive signal costs led to an average loss of 30. These 6 observations are probably best interpreted as unsuccessful attempts of low-quality sellers to fool the buyers, or simply mistakes, instead of mixed-equilibrium play. We conclude that mixed equilibria do not organize the data well. In the remainder, we will therefore focus on the pure strategy equilibria.

For each matching group, we computed the number of outcomes consistent with the pooling equilibrium and the number of outcomes consistent with the separating equilibrium. An outcome is consistent with the pooling equilibrium if and only

²⁴ Remember, sellers had access to a social history screen of 10 periods deep of their own matching-group, which also listed the behavior of the buyers.
²⁵ Depending on the level of \hat{z}^* , the best response of the good [bad] type seller equals either signal costs of m_g [m_b] as given by expression (1) [(4)], or zero signal costs.

if there was no sale. An outcome is consistent with the separating equilibrium if and only if the buyer's decision whether or not to buy was in accordance with the separating prediction that depended on the quality of the seller and the actual noise term in the signal.²⁶ (Thus, an outcome may be consistent with both types of equilibria and also with neither type.) Table 6 lists for each matching group the extent to which actual play agreed with either of the equilibria in periods 21–40. In treatments σ_0 , σ_{10} and σ_{40} , the outcomes of all groups agreed much better with the separating equilibrium than with the pooling equilibrium. In treatment σ_{120} , the results were less clear-cut; either equilibrium attracted half of the groups. The results based on the strategy method in periods 41–50 were the same as the ones reported for periods 21–40, except that two of the groups (#4 and #5 in treatment σ_{120}) that were playing in accordance with separation in periods 21–40 switched to pooling in periods 41–50. In fact, in periods 21–40, for these two groups the separating equilibrium only predicted marginally more outcomes than the pooling equilibrium. Our interpretation is that these groups had not yet converged to equilibrium in periods 21–40. After unsuccessfully trying to establish separating play, subjects in these groups switched to pooling in the final 10 periods. Thus, in σ_{120} , at the end of the experiment 6 of the 8 matching-groups agreed with pooling and the 2 others with separating.

Table 6 also shows how and to what extent actual results deviate from best responses. The procedure to calculate best responses for sellers was already explained above. For buyers we used the following procedure. In each matching-group and each period, we determined the average signal cost chosen by the good type sellers (\widehat{m}_g) and the average signal cost chosen by the bad type sellers (\widehat{m}_b) in the last 10 periods. Then we computed the best response cutoff level z^{BR} given these average signal costs.²⁷ This cutoff level and the received signal together determined the buyer's best response decision to buy or not. In most groups of the treatments with noisy signals high-quality sellers chose lower signal costs than the best response prediction. Note also that buyers bought less than the best response model predicted in treatments σ_0 , σ_{10} and σ_{40} , but more than the best response model in treatment σ_{120} . This limits the scope for an explanatory role of risk-aversion. Risk averse buyers should use higher cutoff levels than risk neutral buyers and therefore buy less often. The data in σ_0 , σ_{10} and σ_{40} deviate in the direction expected by risk aversion. In σ_{120} , the data deviate in the opposite direction, however.

We now take a closer look at buyers' behavior. For each individual buyer, we estimated her personal cutoff signal \widehat{z}^* below which she did not buy. The cutoff level was set such that the number of errors against the cutoff level was minimized. Table 7 presents the data separately for the groups that were classified as pooling and the ones that were classified as separating in Table 6. In σ_{120} , subjects in pooling groups employed much larger cutoff levels than subjects in separating groups. For the separating groups, subjects used higher cutoff levels than predicted by best response and equilibrium in treatments σ_0 , σ_{10} and σ_{40} , but lower cutoff levels in treatment σ_{120} . A Kruskal–Wallis test performed at the group level does not reject the hypothesis that there are no differences in estimated cutoff levels between the treatments ($p = 0.51$). Overall, subjects' cutoff levels were not sufficiently responsive to the noise in the signal.

Taken together, our main findings can be summarized as follows. In the treatments with no, low and intermediate noise the high-quality sellers clearly aim for separation, choosing signal costs that increase monotonically with noise. For high noise seller behavior is mixed with the majority of (high-quality) sellers focusing on pooling. The empirical best responses of high-quality sellers also increase with noise, albeit less steep than (separating equilibrium) theory predicts. Outcomes resulting from the joint behavior of seller and buyer are in agreement with separation for σ_0 , σ_{10} and σ_{40} . Yet for the case of σ_{120} the majority of the matching groups coordinate on a pooling equilibrium. For the separating groups, buyers' actual cutoff levels do not vary with noise, contrasting best response (and equilibrium) behavior which requires higher cutoffs for higher levels of noise.

4.2. Explaining the results

In the previous section we compared actual behavior to equilibrium behavior and the heuristic of (empirical) best response. Our main findings indicate that actual choices do not fully track both equilibrium predictions and empirical best responses. This particularly holds true for buyers, whose actual cutoff levels appear largely insensitive to noise. Because playing best response can be cognitively demanding, it is not unlikely that subjects make use of other plausible (and simpler) heuristics like reinforcement or imitation as well. In this section, we employ a maximum likelihood procedure to estimate an attraction learning model that allows sellers and buyers to use any of these three heuristics. We are not the first to develop learning models along these lines. To the contrary, our approach blends with the work of Brandts and Holt (1992, 1993), Cooper et al. (1997a, 1997b), McKelvey and Palfrey (1998), Camerer and Ho (1998, 1999) and Wilcox (2006). Our approach differs from previous work primarily because of the noise in the signaling technology.

We first describe the model for the sellers. In period t , seller i with product quality q (low or high) updates attraction $A_{i,c,q}[t]$ of choosing signal cost c according to:

$$A_{i,c,q}[t] = \varphi A_{i,c,q}[t - 1] + \delta_{BL} E\pi_{BL,i,c,q}[t] + \delta_{IM} E\pi_{IM,i,c,q}[t] + (1 - \delta_{BL} - \delta_{IM}) E\pi_{RE,i,c,q}[t] \tag{6}$$

²⁶ For $\sigma = 10$ where a separating equilibrium does not exist, we defined an outcome consistent with separation if either (i) the seller is of the bad type and there is no sale, or (ii) the seller is of the good type and trade takes place. Here the definition of a separating outcome thus corresponds with the $\sigma = 0$ case.

²⁷ This best response is given by $z^{BR} = \frac{1}{2} \left\{ (\widehat{m}_g + \widehat{m}_b) + \frac{2\sigma^2 \ln \left[\frac{\beta^* (1-\beta)}{(1-\beta^*)\beta} \right]}{\widehat{m}_g - \widehat{m}_b} \right\}$.

Table 6

Average actual outcomes per group and comparison with best response (periods 21–40).

			1		2		3		4		5		6		7		8		
			Act.	BR	Act.	BR	Act.	BR	Act.	BR	Act.	BR	Act.	BR	Act.	BR	Act.	BR	
σ_0	SC	Low	2.4	0.0	2.5	8.0	0.0	0.0	28.4	0.0									
		High	91.7	90.0	92.8	89.9	100.0	100.0	106.8	95.0									
	Buyer buys? Agrees with eq.	Pool	50.0	51.3	50.0	51.3	50.0	50.0	50.0	50.0	55.0	62.5							
		Separ.	100		97.5		100.0				45.0								
		P/S	S		S		S		S										
σ_{10}	SC	Low	6.0	0.0	5.3	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0					
		High	97.8	103.9	100.3	105.3	118.6	124.3	122.0	130.8	125.1	124.6							
	Buyer buys? Agrees with eq.	Pool	47.5	53.8	48.8	52.5	47.5	50.0	42.5	50.0	46.3	50.0							
		Separ.	52.5		51.3		52.5		57.5		53.8								
		P/S	95.0		93.8		97.5		92.5		96.3								
σ_{40}	SC	Low	2.0	0.0	0.5	0.0	2.8	0.0	3.0	0.0	1.4	0.0	22.3	0.0					
		High	90.0	136.1	112.2	156.0	119.3	148.2	131.3	155.2	153.3	162.7	191.0	176.4					
	Buyer buys? Agrees with eq.	Pool	43.8	50.0	42.5	52.5	46.3	55.0	48.8	55.0	50.0	51.3	53.8	53.8					
		Separ.	56.3		57.5		53.8		51.3		50.0		46.3						
		P/S	83.8		90.0		91.3		93.8		92.5		88.8						
σ_{120}	SC	Low	0.1	0.0	0.0	0.0	0.5	0.0	0.0	0.0	9.0	0.0	1.1	0.0	5.0	0.0	0.0	0.0	
		High	4.3	121.1	28.8	205.5	32.8	231.7	62.3	220.3	86.9	187.5	92.6	210.3	103.8	178.7	197.6	160.2	
	Buyer buys? Agrees with eq.	Pool	35.0	0.0	21.3	7.5	23.8	1.3	45.0	23.8	32.5	22.5	22.5	25.0	41.3	30.0	40.0	37.5	
		Separ.	65.0		78.8		76.3		55.0		67.5		77.5		58.8		60.0		
		P/S	42.5		68.8		58.8		72.5		80.0		75.0		68.8		87.5		
			P		P		P		S(P)		S(P)		P		S		S		

Remarks: We have 23 independent group observations, with 4, 5, 6 and 8 groups for treatments σ_0 , σ_{10} , σ_{40} and σ_{120} , respectively. The table reports the average signal cost SC within a group (by quality level), buyers' buy decision, and the percentage of outcomes that is in line with the pooling and the separating equilibrium outcomes. Groups are ordered on the basis of average signal cost when quality is high. 'Act.' means actual, 'BR' refers to best response. Groups 4 and 5 in σ_{120} converged to pooling with the strategy method: in periods 41–50, 87.5% (75.0%) and 77.5% (75.0%) of the outcomes agreed with pooling (separating) in groups 4 and 5, respectively.

Table 7
Average estimated cutoff levels and tests for equality (periods 21–40).

	Pooling groups			Separating groups		
	Estimated actual (Stan. dev.)	Predictions best response [equilibrium]	Actual vs. best response [equilibrium]	Estimated actual (Stan. dev.)	Predictions best response [equilibrium]	Actual vs. best response [equilibrium]
σ_0	–	–	–	94.0 (4.1)	53.0 [90.0]	0.07 [0.11]
σ_{10}	–	–	–	96.2 (15.1)	57.8 [25.6]	0.04 [0.04]
σ_{40}	–	–	–	98.5 (22.4)	75.3 [75.7]	0.03 [0.03]
σ_{120}	211.5 (166.2)	330.9 [800]	0.72 [0.07]	96.9 (84.8)	122.2 [141.0]	0.47 [0.07]

Remarks: For each buyer the actual cutoff level was estimated on the basis of the choices in periods 21–40. The cutoff level was set such that the number of errors against the cutoff level was minimized. Standard deviations (based on observations per person) are in parentheses; equilibrium predictions/comparisons appear in square brackets.

* When $\sigma = 10$ only pooling on zero is a Nash equilibrium. Within treatment comparisons are based on Wilcoxon tests performed at the matching group level.

where ϕ represents the weight on the attraction in the previous period, δ_{BL} represents the weight on belief learning (i.e., best response), δ_{IM} the weight on imitation learning and $(1 - \delta_{BL} - \delta_{IM})$ the weight on reinforcement learning. For each level of signal cost c , the seller's expected payoff from belief learning $E\pi_{BL,i,c,q}[t]$ is determined according to the best response procedure described in the previous section. That is, the error minimizing cutoff level for the buyer is estimated from the history screen. Then, given this cutoff level, the expected payoff for the seller for each signal cost and quality is calculated.

The seller's expected payoff from imitation $E\pi_{IM,i,c,q}[t]$ is determined with the following procedure. For each possible combination of signal cost and quality, each entry in the history screen is given a weight depending on how close the signal cost of that entry is to the particular signal cost c in question. For each entry a weight is defined by $1/(1 + |c_{entry} - c|)$, where c_{entry} represents the signal cost observed in the specific entry. Then the weight of each entry is normalized so that the weights add up to 1. The weights thus denote the 'resemblance' of the signal cost of the entry in the history screen to the signal cost considered. For each entry the seller's payoff given the buyer's actual decision is multiplied by the weight. Then the weighted payoff $E\pi_{IM,i,c,q}[t]$ is constructed by summing the resulting numbers over all the entries. Note that this payoff is largest for signal costs c that imitate the most successful signal costs observed in the past. These include relatively low signal costs that were combined with lucky signals (i.e. a high draw for the random noise term).

Finally, the expected payoff from reinforcement $E\pi_{RE,i,c,q}[t]$ depends on whether the signal cost c was actually chosen by the subject in the previous period. If it was not chosen in a particular combination of signal cost and quality, the expected payoff equals 0. If it was chosen, the expected payoff equals the seller's payoff given the buyer's actual decision and the quality (low or high).

Given a seller's profile of attractions in a certain period, he chooses the signal cost according to a logistic response function. In particular, the probability $P_{i,c,q}[t]$ that seller i with product quality q chooses signal cost c in period t equals:

$$P_{i,c,q}[t] = \frac{e^{\lambda_i A_{i,c,q}[t]}}{\sum_{c=0}^{350} e^{\lambda_i A_{i,c,q}[t]}} \tag{7}$$

where λ_i represents the seller's precision level.

For the buyers we employed an estimation procedure with exactly the same structure. That is, in period t buyer i who observed the signal s updates attraction $A_{i,a,s}[t]$ of choosing action a (with $a = 0$ if she does not buy and $a = 1$ if she buys) according to:

$$A_{i,a,s}[t] = \phi A_{i,a,s}[t - 1] + \delta_{BL} E\pi_{BL,i,a,s}[t] + \delta_{IM} E\pi_{IM,i,a,s}[t] + (1 - \delta_{BL} - \delta_{IM}) E\pi_{RE,i,a,s}[t] \tag{8}$$

Here ϕ , δ_{BL} , δ_{IM} and $(1 - \delta_{BL} - \delta_{IM})$ are defined like above. For each signal and action, the buyer's expected payoff from belief learning $E\pi_{BL,i,a,s}[t]$ is calculated according to the buyer's best response procedure previously described in Section 4.1. In the history screen, the average signal cost of the high-quality seller and the average signal cost of the low-quality seller is determined. Given these average signal costs of the two seller types, the expected payoff for the buyer for each possible combination of action and signal is calculated.

The following procedure explains how the buyer's expected payoff from imitation $E\pi_{IM,i,a,s}[t]$ is determined. For each possible combination of signal and action, each entry in the history screen receives a weight that correlates with the proximity of the signal of that entry to the particular signal in the combination. Specifically, for each entry the weight is given by $1/(1 + |s_{entry} - s|)$, where s_{entry} represents the signal observed in the specific entry. The weights are normalized so that they add up to 1. For each entry the weight is multiplied by the buyer's payoff given the buyer's action and the seller's actual quality, and the weighted payoff $E\pi_{IM,i,a,s}[t]$ is constructed by summing the resulting numbers over all the entries. We refer to this heuristic as imitation, because it is determined according to the same structure as for the seller. Unlike in the seller case, however, the heuristic could alternatively be interpreted as a "behavioral" belief learning model. The reason is that buyers observe signals and update the corresponding attractions, but do not choose signals. So where sellers can simply imitate the most successful signal costs of other sellers, a similar straightforward interpretation is lacking for the buyers. For them the procedure contains a belief learning element, because by allowing the buyer to update the expected

Table 8
Maximum likelihood estimation results learning model (periods 11–40).

	Buyer model			Seller model		
	$\sigma_\lambda = 0$	$\sigma_\lambda \geq 0$	$\sigma_\lambda \geq 0$	$\sigma_\lambda = 0$	$\sigma_\lambda \geq 0$	$\sigma_\lambda \geq 0$
δ_{BL}	0.23 (0.04)	0.48 (0.09)	–	0.26 (0.01)	0.28 (0.01)	–
δ_{IM}	0.27 (0.06)	0.35 (0.07)	–	0.18 (0.01)	0.18 (0.01)	–
$\delta_{BL}(\sigma_0)$	–	–	0.10 (0.15)	–	–	0.17 (0.01)
$\delta_{IM}(\sigma_0)$	–	–	0.79 (0.24)	–	–	0.41 (0.04)
$\delta_{BL}(\sigma_{10})$	–	–	0.00 (0.09)	–	–	0.46 (0.03)
$\delta_{IM}(\sigma_{10})$	–	–	0.94 (0.27)	–	–	0.24 (0.03)
$\delta_{BL}(\sigma_{40})$	–	–	0.36 (0.12)	–	–	0.29 (0.03)
$\delta_{IM}(\sigma_{40})$	–	–	0.64 (0.20)	–	–	0.45 (0.03)
$\delta_{BL}(\sigma_{120})$	–	–	0.93 (0.14)	–	–	0.53 (0.01)
$\delta_{IM}(\sigma_{120})$	–	–	0.07 (0.05)	–	–	0.10 (0.00)
ϕ	0.55 (0.08)	0.67 (0.05)	0.22 (0.20)	0.86 (0.01)	0.89 (0.00)	0.89 (0.01)
μ_λ	11.52 (2.00)	7.51 (1.43)	13.96 (3.47)	22.81 (0.60)	24.52 (2.10)	29.43 (1.50)
σ_λ	–	0.77 (0.11)	0.70 (0.09)	–	0.78 (0.06)	0.88 (0.07)
$-\log L$	810.2	753.6	696.4	8757.7	7984.1	7537.8

Remark: Standard deviations are in parentheses.

profitability of the buy decision corresponding to the received signal on the basis of the observed payoffs of buying after different signals, it implicitly assumes that buyers make inferences about seller behavior.

The buyer’s expected payoff from reinforcement $E\pi_{RE,i,a,s}[t]$ in a particular combination of signal and action equals 0 if the signal in the combination was not observed by the buyer in the previous period. If it was observed, the expected payoff equals the buyer’s payoff given the seller’s actual quality and the action considered (either buy or not buy).

Given that the buyer updates her attractions in this way, she chooses whether or not to buy according to a logistic response:

$$P_{i,a,s}[t] = \frac{e^{\lambda_i A_{i,a,s}[t]}}{\sum_{a=0}^1 e^{\lambda_i A_{i,a,s}[t]}} \tag{9}$$

where λ_i represents the buyer’s precision level.

Following Wilcox (2006), we correct for heterogeneity between subjects by drawing the precision levels λ_i from a log-normal distribution $(\mu_\lambda, \sigma_\lambda)$. We estimated the model from period 11 onwards, the first time that the history screen was completely filled with observations from the 10 most recent periods. In period 11, we set all the attractions of period 10 equal to 0.²⁸

Table 8 presents the results. First we discuss the results for the buyer model. The table presents three versions of the model, one where all precision levels are forced to be the same ($\sigma_\lambda = 0$, first column) and two with heterogeneity between subjects (second and third column). Unsurprisingly, the likelihood of the data improves significantly for the buyer model when we allow for heterogeneity (comparison first and second column, likelihood ratio test, $p < 0.001$). Remarkably, the estimates of the buyer model put substantially less weight on reinforcement learning when heterogeneity is introduced: the weight decreases from $(1 - \delta_{BL} - \delta_{IM}) = 0.49$ when $\sigma_\lambda = 0$ to 0.17 when $\sigma_\lambda \geq 0$ (second column). Thus, like Wilcox (2006), we find that ignoring heterogeneity leads to a gross overstatement of reinforcement learning in the buyer model. When the estimates are corrected for differences between subjects, the largest weight is assigned to belief learning (0.48) and imitation or behavioral belief learning (0.35). When we in addition allow the learning parameters to vary across treatments (third column), the likelihood again increases significantly and the weight assigned to reinforcement learning is further reduced. In none of the treatments the weight to reinforcement exceeds 0.17. Notice that the proportions assigned to belief learning and imitation (“behavioral” belief learning) vary substantially across treatments. As we argued above, for the buyer these two models are very much aligned and capture more or less the same behavior, so it is not surprising that these proportions vary.

The picture for the sellers looks very different. In the comparison of columns four and five, the estimates of the seller model appear more robust with respect to the heterogeneity of subjects. With or without heterogeneity, reinforcement learning receives the highest weight. When we allow the learning parameters to vary across treatments (final column), the weight on reinforcement diminishes somewhat but still remains sizable, viz. in the range of 0.26 to 0.42. The estimates thus suggest a genuine role for reinforcement learning. This makes sense if one takes the complexity of the problem into account (see below). Like in the buyer model, the likelihood of the data improves significantly for the seller model when heterogeneity is introduced (comparison fourth and fifth columns) and when difference across treatments are allowed (comparison fifth and sixth columns, in both cases likelihood ratio test yields $p < 0.001$).

²⁸ In the maximum likelihood procedure we had to integrate over the lognormal distribution. We did this numerically, approximating the integral by a discrete trapezoid. We will send the details of this procedure upon request. To keep the estimation problem manageable, we do not allow for heterogeneity in ϕ . Wilcox (2006) mentions that the bias from ignoring heterogeneity in ϕ is a ‘relatively minor problem’.

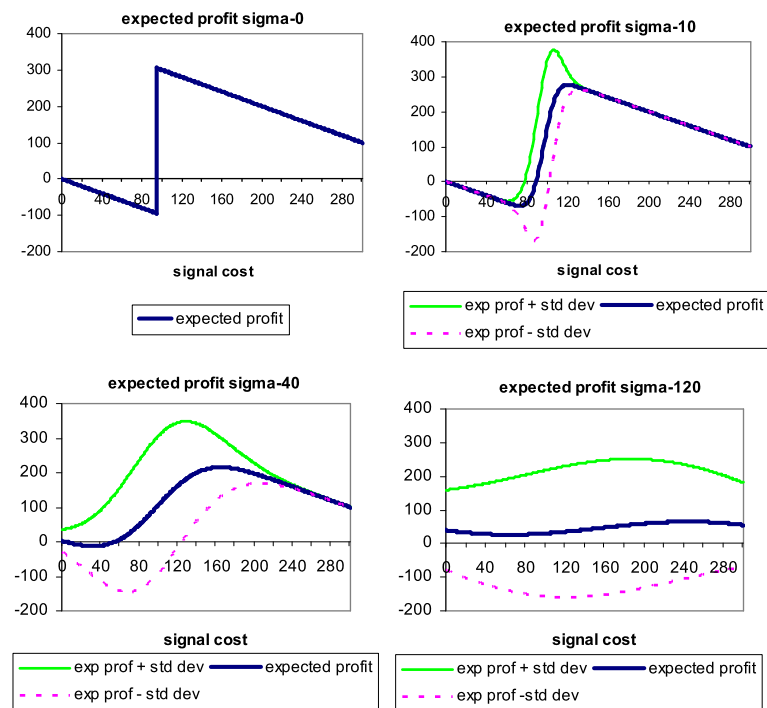


Fig. 4. Expected profit high-quality seller as function of signal cost given actual behavior of buyers in periods 21–40.

The sellers' decision problem is complicated by the fact that they choose from a very large action space. In fact, in our data we observe that sellers often repeat successful choices (for a given quality) and only experiment with a limited set of signal costs. This feature of the data credits reinforcement learning. The combination of a high weight on reinforcement and a high precision level correctly leads to high predicted probabilities of previously chosen signal costs. In contrast, buyers face a much simpler decision problem, because they only have to judge the profitability of two actions given the signal that they receive, with one of the actions (not buy) always leading to a payoff of zero. From this perspective, it makes perfect sense that buyers employ more sophisticated heuristics like ("behavioral") belief learning to a larger extent.

The estimated learning model helps explaining why subjects' actual choices deviate (to some extent) from best response behavior, as was observed in Section 4.1. First consider the behavior of sellers. Remember that in the treatments with noisy signals, the high-quality sellers' actual signal costs were smaller than their best responses (and also insufficiently responsive to the noise). A plausible explanation for this is that reinforcement receives a lot of weight in the seller model. In the early periods, sellers started with signal costs that were on average below the best response levels. Possibly high-quality sellers originally started with a simple rule of choosing a signal cost equal to the payoff that a low-quality seller would receive in case of a sale enhanced by a 'safety margin'. They subsequently learned to a considerable extent by reinforcement and therefore experimented too little, thus not finding out that their expected profits would be maximized at higher levels.

Another noteworthy feature of sellers' behavior is that the volatility of the signal costs chosen by high-quality sellers increased with the noise in the signal. As the histograms presented in Fig. 3 revealed, especially in treatment σ_{120} high-quality sellers chose signals all over the place when they attempted to separate. Fig. 4 provides a coherent reason why this may have occurred. This figure presents the expected profit of a high-quality seller and its variance conditional on the signal cost submitted, given the actual average cutoff level of the buyers in a treatment. In treatment σ_0 , the payoff function is steep and, indeed, high-quality sellers' signal costs in this treatment were clustered in a very small interval. In treatment σ_{120} , however, the payoff function has become very flat. Thus even though sellers partly employ a belief learning heuristic – see the estimates of δ_{BL} in the right-hand panel of Table 8 – their expected payoffs from belief learning provide little guidance in locating the exact best response. This also provides an additional explanation (on top of reinforcement playing a genuine role) why especially for σ_{120} actual signal costs fall short of the empirical best response.

Turning to the buyer, the lack of responsiveness of buyers' cutoff levels to noise reported in Section 4.1 can be explained in a similar way. It is quite natural for buyers to start with a 'myopic' cutoff level of 90, the amount that the low-quality seller earns when his good is bought. The question of interest is why in the treatments with noise buyers did not learn to change the cutoff sufficiently into the direction of the true best response, especially so given that belief learning receives the largest weight in the buyer's learning model (cf. Table 8). Fig. 5 provides an answer to this question. For each treatment, this figure shows the buyer's expected payoff and its volatility conditional on each possible cutoff level z , given the actual average signal costs chosen by low-quality and high-quality sellers. It appears that the expected payoff functions are very flat around 90, and in all treatments the profit at a cutoff level of 90 is close to the profit of the optimal cutoff level. Irrespective

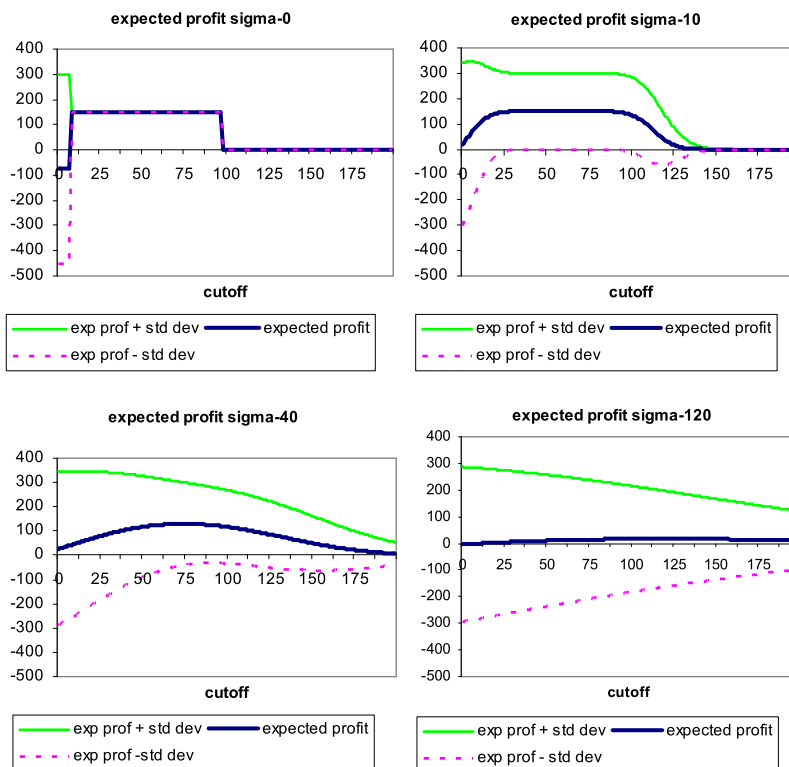


Fig. 5. Expected profit buyer as function of cutoff given actual behavior of sellers in periods 21–40.

of the amount of noise, the expected payoffs from belief learning thus provide hardly any pressure to change the original cutoff level. Our experimental finding that the effect of noise on cutoff levels is (much) smoother than predicted by theory lines up with experimental results in different games. For instance, Brandts and Figueras (2003) find that reputation building increases with the fraction of honest bankers, but not as steeply as theory predicts. In a similar find Georganas and Nagel (2011) find that previous ownership in a company varies less with takeover bids than predicted.

The learning model together with the flatness of the buyer's expected profit functions also explains our important experimental finding that subjects separate even in treatment σ_{10} where the noise is so small that a pure strategy separating equilibrium does not exist. As reported in Table 7, the actual average cutoff level buyers employ in this treatment is well above the best response level: 96.2 versus 57.8. In periods 21–40 buyers actually earned on average 133.9 (at a standard error of 157.6). If they would consistently have used the much lower optimal best response cutoff level of 57.8,²⁹ their profit would have been 141.8 (standard error 166.0). In the large majority of 95% of the cases, the optimal cutoff leads to exactly the same choice and profit as the buyer actually made. In only 4% of the cases the optimal cutoff would have led to a higher profit, while in 1% of the cases it would have led to a lower profit. The difference between the actual profit and the optimal profit generated by a hypothetical cutoff strategy of 57.8 is not significant according to a Wilcoxon rank test using all observations as data points ($p = 0.15$). Thus, as already indicated above, there was no noticeable pressure on buyers to lower their cutoff. Given the high cutoff levels actually used by the buyers, it is no surprise that high-quality sellers continued to send messages with high signal costs. In fact, due to reinforcement learning being a significant behavioral driver for sellers as well, their signal costs were even a bit lower than the actual best response (112.8 versus 117.9). In sum, the non-equilibrium separation outcome materializes in σ_{10} because buyers employed higher cutoff levels than they should in equilibrium, but were hardly punished for doing so, while sellers were close to best responding.

A potential concern is that the observed separation in σ_{10} is an artefact of the limited length of our sessions. To address this concern, we used the estimates for treatment σ_{10} reported in Table 8 to simulate how play unfolds after period 10. We calculated the averages of the actual signal costs employed by high and low type sellers and their respective variances, and the average rejection rates by buyers in the first 10 periods.³⁰ We used these data to create the history screen for the first 10 periods. The learning model of both buyer and seller and the estimated parameters then determine how behavior

²⁹ If buyers would have used this cutoff level of $z^{BR} = 57.8$, the theoretical best response for the bad (good) type seller is to choose signal costs equal to zero (81.34). So even in that case the bad type would not have an incentive to deviate from 0, i.e. to upset the separating outcome.

³⁰ The sellers' averages are reported in Table 9. The average rejection rates of buyers in the first 10 periods equal 94.67% for low-quality sellers and 12.89% for high-quality sellers.

Table 9Simulations learning model treatment σ_{10} .

Simulation starts	Quality seller	Input: periods 1–10	Result: periods 501–1000
Like observed in experiments	Signal cost low	8.5 (29.4)	13.8 (26.4)
	Signal cost high	113.4 (27.6)	110.7 (2.6)
Hypothetical from pooling	Signal cost low	0.0 (0.0)	11.6 (23.1)
	Signal cost high	0.0 (0.0)	113.6 (3.4)

Remarks: The simulations employ the parameter estimates of treatment σ_{10} and project how play unfolds in 1000 periods. The upper part reports results using the actual average signal cost chosen by high and low seller types (and their respective variances). The lower part reports results from a hypothetical start from the pooling equilibrium. The simulation results are the averages of 50 runs with matching groups consisting of 16 sellers and 16 buyers each.




Fig. 6. Signal cost in simulations for typical matching groups. The upper panel presents results starting from a history screen like observed in the experiments. The lower panel presents results starting from a (hypothetical) pooling outcome. High-quality signal costs are displayed by black dots, low-quality signal costs by grey dots.

develops after period 10. The upper part of Table 9 presents the results of this exercise and the upper panel of Fig. 6 shows the results of a typical matching group. The results clearly indicate that the separation outcome observed in treatment σ_{10} is a stable phenomenon. Even after 1000 periods high and low-quality sellers continue to separate in our simulations.

The lower part of Table 9 provides insight in how the simulation results depend on initial conditions. These simulations are based on an artificial history screen where high- and low-quality sellers start choosing zero signal costs corresponding to the pooling equilibrium (and the buyer's initial acceptance rate is set equal to zero). Remarkably, with the estimated parameters of the learning model simulated play rather quickly drifts toward the separating play observed in our experiment. The reason behind this is that sellers' and buyers' initial behavior is noisy. Thus, some signal costs get accepted by mistake. When this happens in the case that a high-quality seller accidentally employs a high signal cost, both the seller's and buyer's behavior is reinforced. A similar incident with a low-quality seller is not reinforced. Thus, gradually high-quality

Table 10

Seller profits conditional on signal and type (periods 21–40)

	Low type		High type	
	Pooling SC < 90	Separating SC ≥ 90	Pooling SC < 90	Separating SC ≥ 90
σ_0	−0.7 (7.0)	−46.8 (46.3)	−80.0 (0.0)	294.5 (60.2)
σ_{10}	−0.9 (9.6)	–	30.0 (200.0)	257.8 (100.6)
σ_{40}	3.3 (24.6)	−52.0 (37.0)	117.2 (206.3)	221.7 (119.1)
σ_{120}	22.6 (40.4)	−110 (0.0)	80.7 (170.2)	73.3 (189.8)

Remarks: The cells list average profits. Standard deviations (based on individual observations) in parentheses.

sellers learn to employ high signal costs, and buyers learn to accept high signals. The lower panel of Fig. 6 displays the result of a typical matching group that starts from the pooling outcome.

Average (actual) profits also shed light on another main finding from our experiment, viz. which equilibrium organizes the data best in treatments σ_0 , σ_{40} and σ_{120} . In the introduction we hypothesized that sellers are less willing to pursue the separating equilibrium when the noise increases. With noise, there is always a chance that a separating signal cost of a high-quality seller is pushed below the cutoff of the buyer, in which case the seller incurs a loss. This becomes more likely the higher the level of noise is. For instance, in σ_{40} , the probability that the (equilibrium) separating signal cost of the high-quality seller is not accepted equals 0.05. In σ_{120} , the probability that a high-quality seller incurs a loss increases to 0.23. At the same time, the equilibrium markup in case of a sale decreases from 258 in σ_{40} to 169 in σ_{120} . Thus, for the seller the prospects of the separating equilibrium deteriorate when the noise in the signal increases.

Average (actual) profits shed light on another main finding from our experiment, viz. which equilibrium organizes the data best in treatments σ_0 , σ_{40} and σ_{120} . In the introduction we hypothesized that sellers are less willing to pursue the separating equilibrium when the noise increases. With noise, there is always a chance that a separating signal cost of a high-quality seller is pushed below the cutoff of the buyer, in which case the seller incurs a loss. This becomes more likely the higher the level of noise is. For instance, in σ_{40} , the probability that the (equilibrium) separating signal cost of the high-quality seller is not accepted equals 0.05. In σ_{120} , the probability that a high-quality seller incurs a loss increases to 0.23. At the same time, the equilibrium markup in case of a sale decreases from 258 in σ_{40} to 169 in σ_{120} . Thus, for the seller the prospects of the separating equilibrium deteriorate when the noise in the signal increases.

The actual profit data are in accordance with this explanation. For the treatments where a pooling and separating equilibrium coexist as well as for treatment σ_{10} where this is not the case, Table 9 lists seller profits, separated for high types and low types, and for pooling signals (< 90) and separating signals (≥ 90). For high-quality types, the attractiveness of the separating equilibrium is highest in σ_0 . With noise the loss-gain tradeoff worsens, which favors pooling. In fact, for σ_{120} high-quality sellers made on average slightly more when they selected pooling signal costs. This result agrees with the earlier finding that in the end 6 of the 8 matching-groups in σ_{120} converged to pooling.

From Brandts and Holt (1992, 1993), Cooper et al. (1997a, 1997b) and our paper a unified picture emerges about how people play signaling games. In all these papers, subjects start playing at their myopic optimum after which play unfolds in accordance with adaptive learning dynamics. Brandts and Holt show that this process can lead to the unintuitive equilibrium being played. In the context of a limit pricing game, Cooper et al. find that this process converges to the pooling equilibrium if it exists. In our paper, this process reveals how Spence's wasteful signaling result is saved for behavioral reasons even when (a small amount of) noise in the receiver's perception of the signal is introduced.

5. Conclusion

In this paper we introduced noise in Spence's pure signaling game. Besides being more realistic, allowing for noise in the signal is appealing because it substantially cuts down the number of equilibria. With an unfavorable prior belief, the (pure strategy) separating equilibrium even completely disappears for low levels of noise. It reappears for intermediate noise levels, where the (high type's) signal costs increase with the noise in the signal up to a ceiling. In contrast, a pooling (on no signaling) equilibrium always exists.

In our experiment, noise systematically affected the signal costs of high-quality sellers. High-quality sellers tended to aim for separation. The signal costs of the sellers who opted for separation rose with the noise in the signaling technology. For high noise levels, the separating equilibrium lost ground to the pooling equilibrium, though. With high noise, high-quality sellers faced the risk that a signal cost aimed at separation would fail to accomplish its goal because it received a bad draw for the noise term. In addition, with high noise higher signal costs were required to convince the buyer, which decreased the markup in case of a sale. Thus, the separating equilibrium became much less attractive, which is reflected in the frequency that it was chosen. We did, however, observe a couple of very high signal costs when the noise in the signal cost was high. Noise thus adds to the explanation why in some real life cases wasteful signaling appears so excessive and prominent.

We observed a smoother pattern in the effect of noise on the signal cost than predicted by theory. Conditional on choosing a separating level of signal costs, the signal costs of high-quality sellers increased monotonically with noise. Although intuitive this is a remarkable empirical finding, because it means that subjects separated even in the case where no separating equilibrium existed. It turns out that with little noise subjects initially played as if there were no noise. In particular,

buyers used higher cutoffs than prescribed by equilibrium. The strategic nature of the game was such that there was negligible pressure on buyers to change their initial behavior. Buyers almost made the same amount as they would have earned with their best response and their best response hardly guided them to lower cutoffs. As a result, sellers did not have an incentive to cease separating. Thus, separation did not unravel. A simple attraction learning model incorporating belief learning, imitation and reinforcement, accounts for this intuitive anomaly.

We believe that the mechanisms identified in this paper are also relevant in the field. In some real life situations, senders employ rather inexpensive signals. For instance, a GRE score is an almost perfect predictor of how well a student performs in a PhD program. Notice that this is an example with only a very limited amount of noise in the signaling technology, so that separating may not be supported in equilibrium. Still, students seem to separate successfully through their GRE scores, just like our subjects in treatment σ_{10} . In other cases, like advertising during the Super Bowl, or fighting for a fierce image among the Yanomamö, senders use much more expensive signals with a possibly lower success rate. These examples differ in many dimensions. One crucial dimension may be the signaling technology though. The task that directors of graduate schools face when they interpret GRE scores seems easier compared to the task of average sports fans who interpret the cost behind a Super Bowl commercial, or the task of Yanomamö women who distill a man's fierceness from his scars.

Our results illustrate that in particular senders suffer from a malfunctioning signaling technology. They thus have the largest incentive to improve the signaling technology. Endogenizing the noise in the receiver's perception of the signal may be an interesting avenue for future research.

Appendix A

In this appendix we formally derive the theoretical predictions discussed in Section 2. Recall that $z = m + \sigma \cdot \varepsilon$, with z the signal observed by the buyer, m the signal costs chosen by the seller and ε a random variable with distribution F . With regard to F we make the following three assumptions:

- (F.1) F is continuously differentiable, i.e. density f is continuous.
- (F.2) The density f is strictly positive on the entire real line.
- (F.3) The conditional density of z given m (denoted $g(z | m)$) satisfies the strict monotone likelihood ratio property (MLRP):

$$\frac{g(z|m)}{g(z|m')} = \frac{\frac{1}{\sigma} \cdot f(\frac{z-m}{\sigma})}{\frac{1}{\sigma} \cdot f(\frac{z-m'}{\sigma})}$$

is strictly increasing in z for $m > m'$.

These three assumptions facilitate the equilibrium analysis. First, in the setup of both Matthews and Mirman (1983) and Carlsson and Dasgupta (1997), MLRP implies that the receiver necessarily uses a cutoff strategy in equilibrium. Building on their theoretical analysis, Lemma 1 below reveals that essentially the same result applies in our model where signaling constitutes a pure social waste. Second, the three assumptions together also imply that f is 'nicely' shaped, see our Lemma 2.

Lemma 1. Let $\mu_b(m)$ [$\mu_g(m)$] denote the probability with which the bad [good] type seller chooses signal costs m in equilibrium. Consider non-pooling equilibria only, i.e. $\mu_b(m) \neq \mu_g(m)$ for some $m \geq 0$. Assumptions (F.1) through (F.3) then imply that the buyer's best response rule is of the following form (with $z^* \in (-\infty, +\infty)$):

$$\pi(z) = \begin{cases} 0 & \text{if } z < z^* \\ 1 & \text{if } z > z^*, \end{cases} \tag{A.1}$$

with $\pi(z)$ the probability that the buyer 'Buys' after observing signal z .

Proof. To prove the lemma, we first show that the seller's equilibrium strategy is weakly monotonic in his type. Using this we subsequently show that the buyer's equilibrium payoff from buying is monotonically increasing in the observed signal, positive for large signals and negative for small signals.³¹

Let $p(m | \pi(z)) = \int \pi(z) \cdot \frac{1}{\sigma} \cdot f(\frac{z-m}{\sigma}) dz$ denote the probability with which the buyer buys, given that she uses strategy $\pi(z)$ and the seller chooses signal cost m . For the type $t \in \{b, g\}$ seller, expected payoffs then equal $t \cdot p(m | \pi(z)) - m$. Let M_t be the set of maximizers of this expected payoff function. This set is non-empty because f , and thus $p(m | \pi(z))$, is continuous and the relevant range $[0, t]$ of signal costs m is compact. Suppose there exists a $m'_b \in M_b$ with $m'_b > 0$. It then holds that $g \cdot [p(m'_b | \pi(z)) - p(m | \pi(z))] > b \cdot [p(m'_b | \pi(z)) - p(m | \pi(z))] \geq m'_b - m$ for all $m \in [0, m'_b)$. The second inequality directly follows from $m'_b \in M_b$ whereas the first follows from $g > b$. Thus, the good type strictly prefers $m'_b > 0$ over any lower level of signal costs. Because this holds for any $m'_b > 0$ in M_b , there exists a cutoff level $m^c \geq 0$ such that the bad (good) type necessarily chooses $m \leq m^c$ ($m \geq m^c$) in equilibrium. The seller's equilibrium strategy is thus weakly monotonic.

The buyer's expected payoffs of buying when she observes signal z and the seller plays (μ_b, μ_g) equal $V(z | (\mu_b, \mu_g)) = -y + [x + y] \cdot \beta(z)$. Here $\beta(z)$ denotes the buyer's posterior belief that the seller is of the good type after observing signal

³¹ The proof is in the spirit of Lemmas 1 and 2 in Matthews and Mirman (1983). Because they do not consider mixed strategies, however, our proof more closely follows the one of Proposition 3.1 in Carlsson and Dasgupta (1997).

z and given that the seller uses strategy (μ_b, μ_g) . By Assumption (F.2) this belief is determined by Bayes' rule everywhere:

$$\beta(z) \equiv \Pr(t = \text{good} \mid z, (\mu_b, \mu_g)) = \frac{p \cdot \int f\left(\frac{z-m}{\sigma}\right) \cdot \mu_g(m) dm}{p \cdot \int f\left(\frac{z-m}{\sigma}\right) \cdot \mu_g(m) dm + (1-p) \cdot \int f\left(\frac{z-m}{\sigma}\right) \cdot \mu_b(m) dm} \tag{A.2}$$

Given that f is continuous it follows that $\beta(z)$ is continuous in z . Moreover, if $\mu_b(m) \neq \mu_g(m)$ for some $m \geq 0$, MLRP together with the weak monotonicity of the seller's strategy imply that $\beta(z)$ is strictly increasing in z (cf. Milgrom, 1981, Proposition 2).³² This in turn implies that the buyer's expected payoffs $V(z \mid (\mu_b, \mu_g))$ are continuous and strictly increasing in z . Suppose $V(z \mid (\mu_b, \mu_g)) > [<]0$ for all z . Then the buyer always [never] buys irrespective of the value of z and both seller types would strictly prefer $m = 0$. This contradicts $\mu_b \neq \mu_g$. Hence there is a unique solution z^* to $V(z \mid (\mu_b, \mu_g)) = 0$. \square

Lemma 2. Assumptions (F.1) through (F.3) imply that $f(u)$ is uni-modal and strictly increasing [decreasing] in u for $u < [>]M$, with M denoting the mode.

Proof. Let $c_1 < c_3$. We first show that $f(c) > \min\{f(c_1), f(c_3)\}$ for all $c \in (c_1, c_3)$. Suppose not. Then by the continuity of f there exists a $c_2 \in (c_1, c_3)$ for which $f(c_2) \leq f(c)$ for all $c \in [c_1, c_3]$ (i.e. c_2 is an interior global minimum of f on the compact set $[c_1, c_3]$). Assumption (F.1) then also implies that there exists a $\Delta c \leq \min\{c_3 - c_2, c_2 - c_1\}$ (with $\Delta c > 0$) such that $f(c_2 - \Delta c) \geq f(c_2)$ and $f(c_2 + \Delta c) \geq f(c_2)$. Pick z_h and m' such that $z_h - m' = \sigma(c_2 + \Delta c)$, and take $z_l = \sigma c_2 + m'$ and $m = m' + \sigma \Delta c$. Then $\frac{f\left(\frac{z_h-m}{\sigma}\right)}{f\left(\frac{z_l-m'}{\sigma}\right)} = \frac{f(c_2)}{f(c_2+\Delta c)} \leq \frac{f(c_2-\Delta c)}{f(c_2)} = \frac{f\left(\frac{z_l-m}{\sigma}\right)}{f\left(\frac{z_h-m'}{\sigma}\right)}$. This contradicts that $g(z \mid m)$ satisfies MLRP.

Given that f is continuous and strictly positive on \mathbb{R} it follows that f cannot be monotonically increasing; otherwise $\int_{c_1}^{\infty} f(u) du > \int_{c_1}^{\infty} f(c_1) du = \infty$, contradicting that f is a density. Together with $f(c) > \min\{f(c_1), f(c_3)\}$ for all $c \in (c_1, c_3)$ it follows that f is uni-modal. \square

Assuming that players are restricted to use pure strategies only, Theorem 1 below characterizes the set of possible equilibria. Proposition 1 in the main text directly follows from this theorem.

Theorem 1. Assume that players are restricted to use pure strategies.

- (i) A pooling equilibrium in which both seller types choose $m = 0$ always exists. In this equilibrium the buyer never [always] buys when $p < [>]\beta^* \equiv \frac{y}{x+y}$. Pooling on some $m > 0$ cannot occur.
- (ii) Generically, i.e. for all $p \neq \frac{\beta^* \cdot g}{(1-\beta^*) \cdot b + \beta^* \cdot g}$, it holds that in any separating equilibrium the bad type seller chooses $m = 0$ whereas the good type chooses some positive level of signal costs $m = m_g > 0$. The buyer buys if $z > z^*$ and refrains from buying otherwise. Necessary and sufficient conditions for $((0, m_g); z^*)$ to constitute equilibrium strategies are:

$$\frac{g}{\sigma} \cdot f\left(\frac{z^* - m_g}{\sigma}\right) = 1 \quad \text{with } m_g > z^* - \sigma M \tag{A.3}$$

$$\frac{1}{\sigma} \cdot f\left(\frac{z^*}{\sigma}\right) = \frac{p \cdot (1 - \beta^*)}{\beta^* \cdot (1 - p)} \cdot \frac{1}{g} \tag{A.4}$$

$$b \cdot \left(F\left(\frac{z^*}{\sigma}\right) - F\left(\frac{z^* - m_b}{\sigma}\right) \right) \leq m_b \quad \text{for } m_b > z^* - \sigma M \text{ that solves } \frac{b}{\sigma} \cdot f\left(\frac{z^* - m_b}{\sigma}\right) = 1, \tag{A.5}$$

$$g \cdot \left(F\left(\frac{z^*}{\sigma}\right) - F\left(\frac{z^* - m_g}{\sigma}\right) \right) \geq m_g \tag{A.6}$$

³² To see this directly, note that the sign of $\frac{\partial \beta(z)}{\partial z}$ equals the sign of $\frac{\partial l(z)}{\partial z}$, with $l(z)$ equal to:

$$l(z) \equiv \frac{p \cdot \int f\left(\frac{z-m}{\sigma}\right) \cdot \mu_g(m) dm}{(1-p) \cdot \int f\left(\frac{z-m}{\sigma}\right) \cdot \mu_b(m) dm} = \frac{p \cdot \int \left[\frac{f\left(\frac{z-m}{\sigma}\right)}{f\left(\frac{z-m^c}{\sigma}\right)} \right] \cdot \mu_g(m) dm}{(1-p) \cdot \int \left[\frac{f\left(\frac{z-m}{\sigma}\right)}{f\left(\frac{z-m^c}{\sigma}\right)} \right] \cdot \mu_b(m) dm}$$

The second equality simply follows from dividing both numerator and denominator by $f\left(\frac{z-m^c}{\sigma}\right)$, with m^c defined in the main text. By MLRP, the numerator is strictly increasing in z whenever there exists a $m > m^c$ for which $\mu_g(m) > 0$. Similarly so, the denominator is strictly decreasing in z if there exists a $m < m^c$ for which $\mu_b(m) > 0$. Hence $l(z)$ is strictly increasing in z whenever $\mu_b \neq \mu_g$.

Proof. (i) For any pooling strategy it follows from (A.2) that the buyer’s posterior belief equals her prior belief for any signal z observed; $\beta(z) = p \forall z$. Therefore, if $p \cdot x - (1 - p) \cdot y > 0$, i.e. $p > \beta^* \equiv \frac{y}{x+y}$, the buyer buys for sure and neither type of seller wants to spend positive signaling costs. Similarly, in case $p < \beta^*$ the buyer never buys and the seller’s unique best response is $m = 0$. Only in the knife-edge case $p = \beta^*$ the buyer is indifferent and her probability of buying $\pi(z)$ may vary with z . For the type t seller (with $t \in \{b, g\}$) to be willing to choose some $m > 0$ it must then hold that $t \cdot \frac{\partial p(m|\pi(z))}{\partial m} = 1$, where $p(m|\pi(z)) = \int \pi(z) \cdot \frac{1}{\sigma} \cdot f\left(\frac{z-m}{\sigma}\right) dz$ denotes the probability of trade, given buyer’s strategy $\pi(z)$ and the seller choosing signal costs m .³³ With $b < g$ this first order condition cannot hold simultaneously for both types, so pooling on some $m > 0$ cannot occur.

(ii) From Lemma 1 it follows that in any non-pooling equilibrium the buyer necessarily uses a cutoff strategy like in (A.1). The expected payoff of choosing m for the type t seller then equals $t \cdot (1 - F(\frac{z^*-m}{\sigma})) - m$. Hence the following necessary first order condition for an interior maximum:

$$\frac{t}{\sigma} \cdot f\left(\frac{z^* - m}{\sigma}\right) = 1 \tag{A.7}$$

The l.h.s. (r.h.s.) equals the marginal benefits (costs) of raising m . With $b < g$, for given m condition (A.7) cannot hold simultaneously for both types. The two seller types thus never put positive probability on the same $m > 0$ in equilibrium.

From Lemma 2 it follows that (A.7) allows at most two solutions. If so, only the largest one satisfies the second order condition, because at the optimum f should be increasing. (If f is decreasing, an increase in m at the margin increases the marginal benefits of raising m further.) Therefore, necessarily $\frac{z^*-m_t}{\sigma} < M$. In equilibrium the type t seller thus chooses between the two levels $m = 0$ and $m_t > z^* - \sigma M$ satisfying (A.7) only.

We next show that the bad type seller necessarily chooses $m = 0$. Suppose to the contrary that he chooses $m_b > 0$. This requires that m_b yields the bad type weakly more than $m = 0$ does. Given $g > b$ the good type then already strictly earns more by choosing m_b rather than $m = 0$, so certainly this is the case for signal costs equal to m_g . Hence for all levels of signal costs chosen in equilibrium first order condition (A.7) holds. From Bayes’ rule in (A.2) we then obtain that after observing cutoff signal z^* , the buyer’s posterior belief equals:

$$\beta(z^*) = \frac{p \cdot \frac{1}{g}}{p \cdot \frac{1}{g} + (1 - p) \cdot \frac{1}{b}} = \frac{p \cdot b}{p \cdot b + (1 - p) \cdot g}$$

Because $\beta(z)$ is continuous in z , it necessarily must be such that the buyer is indifferent between her two actions after observing cutoff signal z^* (cf. Carlsson and Dasgupta, 1997):

$$\beta(z^*) = \beta^* \equiv \frac{y}{x + y} \tag{A.8}$$

Since generically $\frac{p \cdot b}{p \cdot b + (1 - p) \cdot g} \neq \frac{y}{x + y}$, an equilibrium in which neither type chooses $m = 0$ cannot exist. Therefore, the bad type seller necessarily chooses $m = 0$.

Only two possible types of pure strategy equilibria remain: (1) both seller types choose $m = 0$ (cf. case (i)), and (2) the bad type chooses $m = 0$ whereas the good type chooses $m_g > 0$ satisfying (A.3). Consider the latter case. From (A.2) and (A.3) we obtain that:

$$\beta(z^*) = \frac{p}{p + (1 - p) \cdot g \cdot \frac{1}{\sigma} \cdot f\left(\frac{z^*}{\sigma}\right)}$$

Together with requirement (A.8) then equality (A.4) follows. Note that generically, this latter equality has either two, or zero solutions (cf. Lemma 2).

Given cutoff z^* implicitly defined in (A.4), the bad type seller should not have an incentive to deviate from $m = 0$. The best candidate deviation level of signal costs m_b necessarily satisfies first order condition (A.7) and $m_b > z^* - \sigma M$ as well. This level m_b should give the bad type seller weakly less than choosing $m = 0$, i.e.:

$$b \cdot \left(1 - F\left(\frac{z^*}{\sigma}\right)\right) \geq b \cdot \left(1 - F\left(\frac{z^* - m_b}{\sigma}\right)\right) - m_b$$

Rewriting this yields requirement (A.5). (From Lemma 2 it follows that only when $\frac{b}{\sigma} \cdot f\left(\frac{z^*}{\sigma}\right) < 1$ no $m_b > z^* - \sigma M$ exists that solves the first order condition (A.7); in that case condition (A.5) is automatically satisfied.) Similarly so, the good type seller should (weakly) prefer choosing $m_g > 0$ over no signal costs at all. This is what condition (A.6) requires. □

The actual existence of a separating equilibrium depends on whether expressions (A.3) through (A.6) in Theorem 1 allow a feasible solution. Theorem 2 below, which generalizes Proposition 2 in the main text to general distribution functions F ,

³³ By Assumption (F.1) $\frac{\partial p(m|\pi(z))}{\partial m}$ does exist. In fact, $\frac{\partial p(m|\pi(z))}{\partial m} = \int \frac{1}{\sigma} \cdot f\left(\frac{z-m}{\sigma}\right) d\pi(z)$, see Carlsson and Dasgupta (1997, fn. 8).

in particular considers how this varies with the value of σ . For the situation in which the level of noise becomes small, the following lemma will appear helpful in proving this theorem. In words it says that the tails of density $f(u)$ become ‘thin’ if we move sufficiently far away from the mode M (see part (a)). As a consequence, we should stay sufficiently close to the mode M if we want $f(u)$ to equal a particular given value v (cf. part (b)).

Lemma 3. Assumptions (F.1) through (F.3) imply that:

- (a) $\forall k > 0 \exists \bar{U}(k) > 0$ such that $f(|u|) \leq \frac{k}{|u|}$ for all $|u| \geq \bar{U}(k)$.
- (b) Let f_+^{-1} (f_-^{-1}) denote the inverse of $f(u)$ on the interval $u \geq M$ ($u \leq M$), with M the mode of $f(u)$.³⁴ It holds that: $\forall k > 0 \exists \underline{V}(k) > 0$ such that $-\frac{k}{v} \leq f_-^{-1}(v) \leq M \leq f_+^{-1}(v) \leq \frac{k}{v}$ for all v satisfying $0 < v \leq \underline{V}(k)$.

Proof. (a) First consider the case $u > 0$. Let $D \equiv \{u \mid f(u) = \frac{k}{u}\}$ denote the set of intersection points of $f(u)$ and $h(u; k) \equiv \frac{k}{u}$. First suppose that this set is bounded (which includes the case that D is empty), i.e. $\exists U(k) > 0$ such that $u \leq U(k)$ for all $u \in D$. Then by the continuity of f and $h(u; k)$ on $(0, \infty)$, for all $u > U(k)$ either $f(u) > \frac{k}{u}$ or $f(u) < \frac{k}{u}$. Now the former would imply:

$$\int_{U(k)}^{\infty} f(u) du > \int_{U(k)}^{\infty} \frac{k}{u} du = k \cdot \left(\lim_{u \rightarrow \infty} \ln u - \ln U(k) \right) = \infty$$

This contradicts $\int_{U(k)}^{\infty} f(u) du \leq 1$ given that f is a density. Hence it must hold that $f(u) < \frac{k}{u}$ for all $u > U(k)$.

Next assume that the set of intersection points D is unbounded. Let $u_1, u_3 \in D$ with $u_1 < u_3$ and $f(u) < \frac{k}{u}$ for all u satisfying $u_1 < u < u_3$. Such u_1 and u_3 do exist when D is unbounded, because otherwise $f(u) \geq \frac{k}{u}$ for all $u \geq u_1$ and $\int_{u_1}^{\infty} f(u) du = \infty$, a contradiction. Define $u_2 \equiv \frac{u_1 + u_3}{2}$. MLRP then requires $\frac{f(u_1)}{f(u_2)} < \frac{f(u_2)}{f(u_3)}$.³⁵ Given $u_1, u_3 \in D$ this becomes $\frac{k^2}{u_1 \cdot u_3} < [f(u_2)]^2$. From $f(u_2) < \frac{k}{u_2}$ this can only be satisfied whenever $\frac{k^2}{u_1 \cdot u_3} < \frac{k^2}{u_2^2}$, i.e. $(\frac{u_1 + u_3}{2})^2 < u_1 \cdot u_3$. Rewriting this we get $(u_1 - u_3)^2 < 0$, a contradiction. Hence D cannot be unbounded.

The case $u < 0$ is simply the mirror image of $u > 0$ and thus immediately follows from the above.

(b) This part follows from part (a). To see this, consider the case $u \geq M$. Here f – and thus f_+^{-1} – is decreasing. Let $\bar{U}(k) > 0$ be the cutoff value as given in part (a), i.e. $f(u) \leq \frac{k}{u}$ for all $u \geq \bar{U}(k)$. Consider values $v \leq f(\bar{U}(k)) \equiv \underline{V}(k)$. From f_+^{-1} decreasing it follows that for all these values $f_+^{-1}(v) \geq f_+^{-1}(f(\bar{U}(k))) = \bar{U}(k)$. Now suppose there exists a v' with $0 < v' \leq \underline{V}(k)$ for which $f_+^{-1}(v') > \frac{k}{v'}$. Given that function $h(u; k) \equiv \frac{k}{u}$ is strictly decreasing, it holds that $h(f_+^{-1}(v'); k) < h(\frac{k}{v'}; k) = v' = f(f_+^{-1}(v'))$. So, at point $f_+^{-1}(v')$ function $f(u)$ lies above function $h(u; k)$. Together with $f_+^{-1}(v') \geq \bar{U}(k)$ this contradicts part (a). Hence necessarily $f_+^{-1}(v) \leq \frac{k}{v}$ for all v satisfying $0 < v \leq \underline{V}(k)$. (Note that $M \leq f_+^{-1}(v)$ follows by definition.) Again, the case $u \leq M$ is the mirror image of $u \geq M$. \square

Theorem 2. Assume that players are restricted to use pure strategies. A necessary condition for a separating equilibrium to exist is that $\sigma \leq g \cdot f(M) \cdot \min\{\frac{(1-p)\beta^*}{p(1-\beta^*)}, 1\} \equiv \bar{\sigma}$. Assuming $\sigma \leq \bar{\sigma}$, it holds that:

- (i) $p \leq \beta^*$: a separating equilibrium does not exist if σ becomes sufficiently small.
- (ii) $\beta^* < p < \frac{\beta^* \cdot g}{(1-\beta^*) \cdot b + \beta^* \cdot g}$: a separating equilibrium always exists. For this equilibrium it holds that $\lim_{\sigma \downarrow 0} m_g = 0$.
- (iii) $p > \frac{\beta^* \cdot g}{(1-\beta^*) \cdot b + \beta^* \cdot g}$: a separating equilibrium does not exist.

Proof. When $\sigma > \bar{\sigma}$ either (A.3) or (A.4) in Theorem 1 does not have a solution, so a separating equilibrium cannot exist. Therefore, $\sigma \leq \bar{\sigma}$ is a necessary condition.

Before proving (i) through (iii) separately, we first show that $\lim_{\sigma \downarrow 0} z^* = \lim_{\sigma \downarrow 0} m_g = \lim_{\sigma \downarrow 0} m_b = 0$ (with m_b the solution to (A.7) for $t = b$). Consider the defining equation (A.4) of z^* and let $c = \frac{p \cdot (1-\beta^*)}{\beta^* \cdot (1-p)} \cdot \frac{1}{g}$. From Lemma 2 it follows that (generally) this equation has either two or no solutions. For σ low enough, (A.4) admits two solutions. Denote these solutions z_l^* and z_h^* respectively, with $z_l^* < z_h^*$. Note that necessarily $\frac{z_l^*}{\sigma} < M$ and $\frac{z_h^*}{\sigma} > M$. First consider the latter solution z_h^* . With f_+^{-1} denoting the inverse of f on the interval above M , we obtain $z_h^* = \sigma \cdot f_+^{-1}(\sigma c)$ from (A.4). From Lemma 3(b) it then follows that for all σ satisfying $0 < \sigma \leq \frac{V(k)}{c} \equiv \bar{\sigma}(k; c)$, necessarily $f_+^{-1}(\sigma c) \leq \frac{k}{\sigma c}$. Hence $z_h^* = \sigma \cdot f_+^{-1}(\sigma c) \leq \sigma \cdot \frac{k}{\sigma c} = \frac{k}{c}$ for all $\sigma \leq \bar{\sigma}(k; c)$. Because this holds for any arbitrary $k > 0$, and $z_h^* > \sigma M$, we obtain $\lim_{\sigma \downarrow 0} z_h^* = 0$.

³⁴ Given that f is monotonically increasing below M and monotonically decreasing above M (cf. Lemma 2), these inverses do exist.

³⁵ To see this, pick z_h and m' such that $z_h - m' = \sigma u_3$. Then let $z_l = \sigma u_2 + m'$ and $m = m' + \sigma(u_2 - u_1)$. This gives $\frac{z_h - m'}{\sigma} = u_3$, $\frac{z_h - m}{\sigma} = u_2 = \frac{z_l - m'}{\sigma}$ and $\frac{z_l - m}{\sigma} = u_1$. The requirement then follows from MLRP.

Next consider a solution z_l^* to (A.4) for which $\frac{z_l^*}{\sigma} < M$. In this case $z_l^* = \sigma \cdot f^{-1}(\sigma c)$ from (A.4), with f^{-1} the inverse of f on the interval below M . From Lemma 3(b) it then follows that for all $\sigma \leq \frac{V(k)}{c} \equiv \bar{\sigma}(k; c)$, necessarily $f^{-1}(\sigma c) \geq -\frac{k}{\sigma c}$. Hence $z_l^* = \sigma \cdot f^{-1}(\sigma c) \geq -\frac{k}{c}$ for all $\sigma \leq \bar{\sigma}(k; c)$. Because this holds for any arbitrary $k > 0$, and $z_l^* < \sigma M$, we obtain $\lim_{\sigma \downarrow 0} z_l^* = 0$. Hence, overall $\lim_{\sigma \downarrow 0} z^* = 0$ for any solution z^* to (A.4).

By inserting $z_l^* = z^* - m_g$ and $c = \frac{1}{g}$ in the reasoning for z_l^* above we immediately obtain $\lim_{\sigma \downarrow 0} (z^* - m_g) = 0$ from Eq. (A.3). Together with $\lim_{\sigma \downarrow 0} z^* = 0$ this implies $\lim_{\sigma \downarrow 0} m_g = 0$. Similarly so for $\lim_{\sigma \downarrow 0} m_b = 0$.

(i) Inequality $p \leq \beta^*$ is equivalent to $\frac{p \cdot (1 - \beta^*)}{\beta^* \cdot (1 - p)} \cdot \frac{1}{g} \leq \frac{1}{g}$. From conditions (A.3) and (A.4) in Theorem 1 we obtain that $f\left(\frac{z_h^* - m_g}{\sigma}\right) \geq f\left(\frac{z_h^*}{\sigma}\right)$ necessarily. Together with Lemma 2 this implies that necessarily $\frac{z_h^*}{\sigma} > M$. Hence in this case only a separating equilibrium based on z_h^* may exist. We show that this separating equilibrium disappears for σ sufficiently small, because the bad type seller obtains an incentive to deviate from $m = 0$ to a positive level of signal costs equal to m_b . To see this, from equality (A.4) we have that $\lim_{\sigma \downarrow 0} f\left(\frac{z_h^*}{\sigma}\right) = 0$. With Lemma 2 and $\frac{z_h^*}{\sigma} > M$ this in turn implies that $\lim_{\sigma \downarrow 0} \frac{z_h^*}{\sigma} = \infty$. Taking the limit in the l.h.s. of condition (A.5) we then obtain that:

$$\lim_{\sigma \downarrow 0} b \cdot \left(F\left(\frac{z_h^*}{\sigma}\right) - F\left(\frac{z_h^* - m_b}{\sigma}\right) \right) = b > 0$$

Here $\lim_{\sigma \downarrow 0} \frac{z_h^* - m_b}{\sigma} = -\infty$ (and thus $\lim_{\sigma \downarrow 0} F\left(\frac{z_h^* - m_b}{\sigma}\right) = 0$) follows from the fact that m_b satisfies both first order condition (A.7) and $m_b > z_h^* - \sigma M$ (such that $\frac{z_h^* - m_b}{\sigma} < M$). Because $\lim_{\sigma \downarrow 0} m_b = 0$ as derived above, requirement (A.5) cannot be satisfied for σ sufficiently small.

(ii) When $\beta^* < p < \frac{\beta^* \cdot g}{(1 - \beta^*) \cdot b + \beta^* \cdot g}$ separating equilibria based on z_l^* and z_h^* may exist side by side. The one based on z_h^* vanishes for low σ , see the proof of part (i). We show that the one based on z_l^* continues to exist under the stated conditions. Note that when $\sigma \leq \bar{\sigma}$, condition (A.4) allows a solution $z_l^* \leq \sigma M$. Given this solution z_l^* and $\frac{p(1 - \beta^*)}{(1 - p)\beta^*} > 1$ from $\beta^* < p$, we can always find a corresponding solution m_g to (A.3). Because $p < \frac{\beta^* \cdot g}{(1 - \beta^*) \cdot b + \beta^* \cdot g}$, it follows from (A.4) that $\frac{b}{\sigma} \cdot f\left(\frac{z_l^*}{\sigma}\right) = \frac{b}{g} \cdot \frac{p(1 - \beta^*)}{(1 - p)\beta^*} < 1$. This implies that condition (A.5) is automatically satisfied, because no $m_b > z^* - \sigma M$ exists that satisfies the first order condition (A.7) (see the proof of Theorem 1). Moreover, from (A.3) and Lemma 2 we have:

$$g \cdot \left(F\left(\frac{z_l^*}{\sigma}\right) - F\left(\frac{z_l^* - m_g}{\sigma}\right) \right) = g \cdot \int_{z_l^* - m_g}^{z_l^*} \frac{1}{\sigma} \cdot f\left(\frac{u}{\sigma}\right) du > g \cdot \int_{z_l^* - m_g}^{z_l^*} \frac{1}{g} du = m_g$$

Hence condition (A.6) is satisfied as well.

(iii) In case $p > \frac{\beta^* \cdot g}{(1 - \beta^*) \cdot b + \beta^* \cdot g}$, (A.4) implies that $\frac{b}{\sigma} \cdot f\left(\frac{z_l^*}{\sigma}\right) = \frac{b}{g} \cdot \frac{p(1 - \beta^*)}{(1 - p)\beta^*} > 1$. At $m = 0$ the marginal benefits for the bad type of raising the signal costs thus exceed the marginal costs of doing so and he wants to deviate from choosing $m = 0$. That is, condition (A.5) cannot be satisfied. \square

The non-existence of a separating equilibrium in cases (i) and (iii) is based on the same intuition. Eq. (A.4) in Theorem 1 provides a precise characterization of the buyer's equilibrium cutoff value z^* on the basis of her posterior beliefs. But the feasible value(s) of z^* may be incompatible with seller's best response behavior given the buyer's cutoff strategy. In particular, no-deviation condition (A.5) for the bad type seller may not be satisfied for the value(s) of z^* that solve (A.4). He thus obtains an incentive to deviate from $m = 0$.

We finally turn to mixed strategy equilibria. Theorem 3 below characterizes the set of mixed strategy equilibria that potentially may exist.

Theorem 3. *When players are allowed to use mixed strategies, only two additional types of (mixed strategy) equilibria may potentially exist:*

- (i) *The bad type chooses $m = 0$ with probability $1 - q_b$ and some $m_b > 0$ with probability q_b , whereas the good type chooses some $m_g \geq m_b$ for sure. A necessary condition for existence is: $p < \frac{\beta^* \cdot g}{(1 - \beta^*) \cdot b + \beta^* \cdot g}$. Conditional on existence it holds that: $\lim_{\sigma \downarrow 0} z^* = \lim_{\sigma \downarrow 0} m_b = \lim_{\sigma \downarrow 0} m_g = b$ and $\lim_{\sigma \downarrow 0} q_b = \frac{(1 - \beta^*) \cdot p}{(1 - p) \cdot \beta^*} \cdot \frac{b}{g} \equiv q_0$.*
- (ii) *The bad type chooses $m = 0$ for sure while the good type chooses $m = 0$ with probability $1 - q_g$ and some $m_g > 0$ with probability q_g . A necessary condition for existence is: $p < \beta^*$. Conditional on existence it holds that: $\lim_{\sigma \downarrow 0} z^* = \lim_{\sigma \downarrow 0} m_g = g$ and $\lim_{\sigma \downarrow 0} q_g = 0$.*

Proof. From the proof of Theorem 1 we obtain the following three observations: (1) in equilibrium the type t seller chooses between the two levels of signal costs $m = 0$ and $m_t > z^* - \sigma M$ satisfying (A.7) only, (2) the bad type seller necessarily

puts positive probability on $m = 0$, and (3) if the bad type puts positive probability on $m_b > 0$ as well, the good type strictly prefers level $m_g \geq m_b$ over $m = 0$ and thus chooses m_g for sure. Together these three observations imply that only mixed strategy equilibria of types (i) and (ii) may potentially exist. The remainder of the proof characterizes these mixed equilibria in more detail and considers the limit equilibria of letting σ become infinitely small.

First note that in a mixed equilibrium it necessarily holds that $z^* > \sigma M$. Suppose $z^* \leq \sigma M$. Then $f(\frac{z^*-m}{\sigma}) < f(\frac{z^*}{\sigma})$ by Lemma 2 and neither type wants to mix between 0 and m (at m marginal benefits equal marginal costs, so at all inframarginal levels below m marginal benefits exceed marginal costs).

(i) Equilibrium values of z^*, m_b, m_g and q_b are characterized by the following four equations:

$$\frac{b}{\sigma} \cdot f\left(\frac{z^* - m_b}{\sigma}\right) = 1 \quad \text{with } m_b > z^* - \sigma M \tag{A.9}$$

$$\frac{g}{\sigma} \cdot f\left(\frac{z^* - m_g}{\sigma}\right) = 1 \quad \text{with } m_g > z^* - \sigma M \tag{A.10}$$

$$\frac{b}{\sigma} \cdot f\left(\frac{z^*}{\sigma}\right) = \frac{(1 - \beta^*) \cdot p \cdot b - \beta^* \cdot (1 - p) \cdot g \cdot q_b}{\beta^* \cdot (1 - q_b) \cdot (1 - p) \cdot g} \quad \text{with } z^* > \sigma M \tag{A.11}$$

$$b \cdot \left(F\left(\frac{z^*}{\sigma}\right) - F\left(\frac{z^* - m_b}{\sigma}\right) \right) = m_b \tag{A.12}$$

Given the last condition, no-divergence requirement (A.6) for the good type is automatically satisfied. Existence depends on whether (A.9) through (A.12) admit a feasible solution. Because these are four non-linear equations (in four unknowns), it is in general hard to determine whether a solution exists. We therefore look at the equilibrium properties when σ becomes small, assuming the mixed strategy equilibrium to exist. The latter requires necessarily $\frac{(1-\beta^*) \cdot p \cdot b - \beta^* \cdot (1-p) \cdot g \cdot q_b}{\beta^* \cdot (1-q_b) \cdot (1-p) \cdot g} < 1$, for otherwise $f(\frac{z^*-m_b}{\sigma}) < f(\frac{z^*}{\sigma})$ and the bad type does not want to mix. This reduces to $p < \frac{\beta^* \cdot g}{(1-\beta^*) \cdot b + \beta^* \cdot g}$. Because $f > 0$, from (A.11) it follows that $q_b < \frac{(1-\beta^*) \cdot p}{(1-p) \cdot \beta^*} \cdot \frac{b}{g} \equiv q_0$ necessarily. (Note that for $p < \frac{\beta^* \cdot g}{(1-\beta^*) \cdot b + \beta^* \cdot g}$ it holds that $q_0 < 1$.)

From (A.9) and (A.10) it follows that $\lim_{\sigma \downarrow 0} (z^* - m_b) = \lim_{\sigma \downarrow 0} (z^* - m_g) = 0$, see the proof of Theorem 2. Moreover, equality (A.12) implies $\lim_{\sigma \downarrow 0} m_b = b$. Together, $\lim_{\sigma \downarrow 0} z^* = \lim_{\sigma \downarrow 0} m_b = \lim_{\sigma \downarrow 0} m_g = b$. Now suppose $\lim_{\sigma \downarrow 0} q_b \neq q_0$. Then there exists some $r < q_0$ such that $q_b \leq r$ for some subsequence $\sigma_r \downarrow 0$. Let z_r solve $\frac{b}{\sigma} \cdot f(\frac{z_r}{\sigma}) = \frac{(1-\beta^*) \cdot p \cdot b - \beta^* \cdot (1-p) \cdot g \cdot r}{\beta^* \cdot (1-r) \cdot (1-p) \cdot g} \equiv k(r)$. From the proof of Theorem 2 it follows that $\lim_{\sigma \downarrow 0} z_r = 0$. With $k(r)$ decreasing in r (given $p < \frac{\beta^* \cdot g}{(1-\beta^*) \cdot b + \beta^* \cdot g}$) and $f(\frac{z}{\sigma})$ decreasing in z for $z > \sigma M$ (cf. Lemma 2), it follows that $z^* \leq z_r$ for all $\sigma_r \downarrow 0$, and thus $z^* \rightarrow 0$ along this subsequence. This contradicts $\lim_{\sigma \downarrow 0} z^* = b > 0$. Hence necessarily $\lim_{\sigma \downarrow 0} q_b = q_0$.

(ii) In this case equilibrium values of z^*, m_g and q_g are characterized by the following three equations:

$$\frac{g}{\sigma} \cdot f\left(\frac{z^* - m_g}{\sigma}\right) = 1 \quad \text{with } m_g > z^* - \sigma M \tag{A.13}$$

$$\frac{g}{\sigma} \cdot f\left(\frac{z^*}{\sigma}\right) = \frac{(1 - \beta^*) \cdot p \cdot q_g}{(\beta^* - p) + (1 - \beta^*) \cdot p \cdot q_g} \quad \text{with } z^* > \sigma M \tag{A.14}$$

$$g \cdot \left(F\left(\frac{z^*}{\sigma}\right) - F\left(\frac{z^* - m_g}{\sigma}\right) \right) = m_g \tag{A.15}$$

Given (A.15), requirement (A.5) for the bad type seller is satisfied. We again look at the equilibrium properties for low values of σ , assuming the mixed strategy equilibrium to exist. This requires $\frac{(1-\beta^*) \cdot p \cdot q_g}{(\beta^* - p) + (1-\beta^*) \cdot p \cdot q_g} < 1$, for otherwise $f(\frac{z^*-m_g}{\sigma}) < f(\frac{z^*}{\sigma})$ and the good type does not want to mix. Therefore $p < \beta^*$ is needed.

From (A.13) it follows that $\lim_{\sigma \downarrow 0} \frac{z^* - m_g}{\sigma} = -\infty$ and $\lim_{\sigma \downarrow 0} (z^* - m_g) = 0$, see the proof of Theorem 2. Because the r.h.s. of (A.14) is bounded from above by $\frac{(1-\beta^*) \cdot p}{(1-p) \cdot \beta^*}$, the l.h.s. is bounded as well. This implies $\lim_{\sigma \downarrow 0} f(\frac{z^*}{\sigma}) = 0$ and thus $\lim_{\sigma \downarrow 0} \frac{z^*}{\sigma} = \infty$. Together with equality (A.15) we obtain $\lim_{\sigma \downarrow 0} m_g = g$, and thus $\lim_{\sigma \downarrow 0} z^* = g$ as well. Now suppose $\lim_{\sigma \downarrow 0} q_g \neq 0$. Then there exists some $r > 0$ such that $q_g \geq r$ for some subsequence $\sigma_r \downarrow 0$. Let z_r solve $\frac{g}{\sigma} \cdot f(\frac{z_r}{\sigma}) = \frac{(1-\beta^*) \cdot p \cdot r}{(\beta^* - p) + (1-\beta^*) \cdot p \cdot r} \equiv l(r)$. From the proof of Theorem 2 it follows that $\lim_{\sigma \downarrow 0} z_r = 0$. With $l(r)$ increasing in r and $f(\frac{z}{\sigma})$ decreasing in z for $z > \sigma M$ (cf. Lemma 2), it follows that $z^* \leq z_r$ for all $\sigma_r \downarrow 0$, and thus $z^* \rightarrow 0$ along this subsequence. This contradicts $\lim_{\sigma \downarrow 0} z^* = g > 0$. Hence necessarily $\lim_{\sigma \downarrow 0} q_g = 0$. \square

Theorem 3 reveals that whenever an equilibrium exists in which the good seller type mixes, this equilibrium converges to the pooling on $m = 0$ equilibrium when the noise becomes small. The other equilibrium in which the bad type mixes converges to a mixed equilibrium that is insufficiently revealing; the bad type chooses $m = b$ with probability q_0 and $m = 0$ otherwise, while the good type chooses $m = b$ for sure. Upon receiving message b the buyer then decides to buy. Interestingly, Theorems 2 and 3 together reveal that for a very favorable prior, only pooling on $m = 0$ can occur. Some amount of noise thus precludes informative signaling altogether in this case.

Appendix B. Supplementary material

The online version of this article contains additional supplementary material.
Please visit doi:10.1016/j.geb.2011.04.006.

References

- Aoyagi, Masaki, Fréchet, Guillaume R., 2009. Collusion as public monitoring becomes noisy: experimental evidence. *J. Econ. Theory* 144, 1135–1165.
- Bagwell, Kyle, 1995. Commitment and observability in games. *Games Econ. Behav.* 8, 271–280.
- Blume, Andreas, Board, Oliver, Kawamura, Kohei, 2007. Noisy talk. *Theoretical Econ.* 2, 395–440.
- Brandts, Jordi, Charness, Gary, forthcoming. The strategy versus the direct-response method: a first survey of experimental comparisons. *Exper. Econ.* doi:10.1007/s10683-011-9272-x.
- Brandts, Jordi, Figueras, Neus, 2003. An exploration of reputation formation in experimental games. *J. Econ. Behav. Organ.* 50, 89–115.
- Brandts, Jordi, Holt, Charles A., 1992. An experimental test of equilibrium dominance in signaling games. *Amer. Econ. Rev.* 82, 1350–1365.
- Brandts, Jordi, Holt, Charles A., 1993. Adjustment patterns and equilibrium selection in experimental signaling games. *Int. J. Game Theory* 22, 279–302.
- Cai, Hongbin, Wang, Joseph Tao-Yi, 2006. Overcommunication in strategic information transmission games. *Games Econ. Behav.* 56, 7–36.
- Calveras, Aleix, 2003. Closure policy when bank inspection can be manipulated. *Europ. Finance Rev.* 7, 385–408.
- Camerer, Colin F., Ho, Teck, 1998. EWA learning in games: heterogeneity, time-variation, and probability form. *J. Math. Psychology* 42, 305–326.
- Camerer, Colin F., Ho, Teck, 1999. Experience-weighted attraction learning in normal-form games. *Econometrica* 67, 827–874.
- Carlsson, Hans, Dasgupta, Sudipto, 1997. Noise-proof equilibria in two-action signaling games. *J. Econ. Theory* 77, 432–460.
- Chagnon, Napoleon A., 1988. Life histories, blood revenge, and warfare in a tribal population. *Science* 239, 985–992.
- Chagnon, Napoleon A., 1992. *Yanomamö – The Last Days of Eden*. Harcourt Brace Jovanovich College Publishers.
- Cho, In-Koo, Kreps, David M., 1987. Signaling games and stable equilibria. *Quart. J. Econ.* 102, 179–221.
- Cooper, David, Garvin, Susan, Kagel, John, 1997a. Signalling and adaptive learning in an entry limit pricing game. *RAND J. Econ.* 28, 662–683.
- Cooper, David, Garvin, Susan, Kagel, John, 1997b. Adaptive learning vs. equilibrium refinements in an entry limit pricing game. *Econ. J.* 107, 553–575.
- Crawford, Vincent P., Sobel, Joel, 1982. Strategic information transmission. *Econometrica* 50, 1431–1451.
- Feltovich, Nick, Harbaugh, Richmond, To, Ted, 2002. Too cool for school? Signalling and countersignalling. *RAND J. Econ.* 33, 630–649.
- Georganas, Sotiris, Nagel, Rosemarie, 2011. Auctions with thresholds: an experimental study of company takeovers. *J. Ind. Organ.* 29, 34–45.
- Güth, Werner, Müller, Wieland, Spiegel, Yossi, 2006. Noisy leadership: an experimental approach. *Games Econ. Behav.* 57, 37–62.
- Hertzenndorf, Mark N., 1993. I'm not a high-quality firm – but I play one on TV. *RAND J. Econ.* 24, 236–247.
- Jeitschko, T.D., Normann, H.T., 2009. Signalling in deterministic and stochastic settings: a theoretical and experimental investigation. Working paper.
- Kanodia, Chadra, Singh, Rajdeep, Spero, Andrew E., 2005. Imprecision in accounting measurement: can it be value enhancing? *J. Acc. Res.* 43, 487–519.
- Kihlstrom, Richard E., Riordan, Michael H., 1984. Advertising as a signal. *J. Polit. Economy* 92, 427–450.
- Landeras, Pedro, Pérez de Villarreal, J.M., 2005. A noisy screening model of education. *Labour* 19, 35–54.
- Maggi, Giovanni, 1999. The value of commitment with imperfect observability and private information. *RAND J. Econ.* 30, 555–574.
- Matthews, Steven A., Mirman, Leonard J., 1983. Equilibrium limit pricing: the effects of private information and stochastic demand. *Econometrica* 53, 981–996.
- McKelvey, Richard D., Palfrey, Thomas R., 1998. Quantal response equilibria for extensive form games. *Exper. Econ.* 1, 9–41.
- Milgrom, Paul R., 1981. Good news and bad news: representation theorems and applications. *Bell J. Econ.* 12, 380–391.
- Milgrom, Paul R., Roberts, D. John, 1986. Price and advertising signals of product quality. *J. Polit. Economy* 94, 796–821.
- Miller, Ross M., Plott, Charles R., 1985. Product quality signaling in experimental markets. *Econometrica* 53, 837–872.
- Nelson, Phillip, 1970. Information and consumer behavior. *J. Polit. Economy* 78, 311–329.
- Nelson, Phillip, 1974. Advertising as information. *J. Polit. Economy* 82, 729–754.
- Spence, Michael, 1973. Job market signaling. *Quart. J. Econ.* 87, 355–374.
- Spence, Michael, 1974. *Market Signaling-Informational Transfer in Hiring and Related Processes*. Harvard University Press, Cambridge.
- van Damme, Eric, Hurkens, Sjaak, 1997. Games with imperfectly observable commitment. *Games Econ. Behav.* 21, 282–308.
- Veblen, Thorstein, 1899. *The Theory of the Leisure Class*. MacMillan, New York.
- Wang, Joseph Tao-Yi, Spezio, Michael, Camerer, Colin F., 2010. Pinocchio's pupil: using eyetracking and pupil dilation to understand truth telling and deception in sender-receiver games. *Amer. Econ. Rev.* 100, 984–1007.
- Wilcox, Nathaniel T., 2006. Theories of learning in games and heterogeneity bias. *Econometrica* 74, 1271–1292.