Persuasion Games with Higher Order Uncertainty *

Frédéric KOESSLER[†]

June 2001

Abstract

In persuasion games, it is well known that a perfectly revealing equilibrium may fail to exist when the decision maker is uncertain about the interested party's payoff-relevant information. However, by explicitly integrating higher order uncertainty into the information structure, this paper shows that a perfectly revealing equilibrium *does* exist when disclosures are not restrained to intervals of the payoff-relevant state space. On the contrary, when payoff-irrelevant disclosures are impossible, a perfectly revealing equilibrium fails to exist as long as there is a strictly positive probability that the decision maker does not know whether the interested party is informed or not. In this case, a partially revealing equilibrium and associated inferences are characterized.

KEYWORDS: Strategic information revelation; Persuasion games; Higher order uncertainty; Provability.

JEL CLASSIFICATION: C72; D82.

1 Introduction

Persuasion games are sender-receiver games in which message sending is cheap, non-binding, and in which some information can be proved (or certified) by particular messages. This last requirement can also be substituted by the assumption that information can be verified or that there exist penalties against lying agents. In such games, one or more interested parties (such as salesmen, regulated firms, plaintiffs or defendants) try to influence a decision maker (such as a consumer, a regulator, a government, or an arbitrator) by strategically providing or concealing information relevant to the decision. The analysis of persuasion games is particularly fruitful to investigate welfare effects of lobbying, to understand how an adversary system provides information to a decision maker, when consumers should rely on a seller, etc.

Since the pioneering contributions of Grossman and Hart (1980), Grossman (1981), and Milgrom (1981), it is well known that efforts to manipulate a decision maker's choice by concealing or distorting information do not always succeed.

*I am grateful to François Laisney, Sandrine Spaeter, Gisèle Umbhauer, and Anthony Ziegelmeyer for helpful comments.

[†]Department of Economics, BETA–*Theme*, Université Louis Pasteur, 61 Avenue de la Forêt-Noire, F-67085 Strasbourg Cedex (France).
Tel.: (33) 3 90 24 20 90.
E-mail: koessler@cournot.u-strasbg.fr.
URL: http://cournot.u-strasbg.fr/koessler.

Indeed, by considering persuasion games to model a sales encounter in which the buyer is to purchase a product having an unknown quality, they show that in every sequential equilibrium the seller reveals all relevant information to the buyer. The intuition of this result is that the buyer will believe the worst information reported by the seller and the seller will only report favorable information. Therefore, voluntary communication results in full revelation as long as the buyer can make rational inferences which take into account the seller's incentives to withhold unfavorable information. The argument behind this fully revealing equilibrium is known as the *unravelling argument*.

Nonetheless, full revelation relies on the implicit assumption that the decision maker knows that the interested party is informed. Actually, the fragility of the unravelling argument has been recently emphasized by Shin (1994a,b) who considered persuasion games in which the decision maker is uncertain about interested parties' payoff-relevant information (second order uncertainty). By assuming that only fundamental events can be proved, he shows that a perfectly revealing equilibrium does not exist.[1] The reason for the failure of the unravelling argument in this setting is that the decision maker has no way to distinguish an interested party who possesses information and remains silent from an interested party who does not possess the information in the first place.

To show that the unravelling argument fails, Shin (1994a,b) assumed that announcements concern intervals of the payoff-relevant state space. He actually noticed that 'the only substantial assumption involved in my model is that announcements are of intervals of the set [of payoff-relevant states]. This is a restriction, but the corresponding gain in terms of the simplicity of the model would seem to justify its use. Future work may examine relaxing this assumption' (Shin, 1994b, p. 62). When truthful reports are justified by the fact that there are large enough penalties imposed on a party whose report is proved to be false at a subsequent date, there is indeed no a priori reason to forbid disclosures about the interested party's information. For example, opposing litigants may have the possibility to compel testimony under oath, subpoena witness, and discover documents that unable courts to know whether a party who claimed ignorance was really uninformed or not.

In this paper we scrutinize the failure of the unravelling argument by suggesting a way to *explicitly* integrate higher order knowledge in the information structure and in communication possibilities. First, we show that if we relax the assumption that announcements are of intervals of the payoff-relevant spate space, then the unravelling argument works again, i.e., a perfectly revealing equilibrium exists. This result neither depends on prior probabilities nor on the average precision of the interested party's information. In particular, the precision of the interested party's information has no effect on the decision maker's inferences and choices.

Second, we show that the failure of perfect revelation is, in effect, robust when disclosures are restrained to payoff-relevant information. More precisely, by expanding the state space in a way allowing third order uncertainty (i.e., uncertainty about the decision maker's knowledge about the interested party's information) we show that a perfectly revealing equilibrium does not exist as long as there is a strictly positive probability that the decision maker does not know whether the interested party is informed or not. In other words, as long as the interested party consid-

[1] This phenomenon has been already illustrated by Okuno-Fujiwara, Postlewaite, and Suzumura (1990) in a Cournot game with information revelation about production costs.

ers as possible, even with an arbitrarily small probability, that the decision maker does not know that he is informed, full revelation does not occur if payoff-irrelevant information cannot be disclosed.

In the last configuration a partially revealing equilibrium is characterized. In this equilibrium, the decision maker perfectly learns the state of Nature (the fundamental state) when he knows that the interested party is informed, but updates only partially his information when he does not know whether the interested party is informed. Then, when payoff-irrelevant revelations are impossible, our result strengthens Shin's result which was proved by assuming that the decision maker never knows whether the interested party is informed or not. Interestingly, the decision maker's inferences and the interested party's disclosure behavior do not depend on the probability that the decision maker receives information.

The plan of the paper is as follows. In Section 2, we show that perfect revelation occurs when all events can be proved. In Section 3, we present our results when only fundamental events can be proved and when the state space is expanded in a way allowing depth of knowledge two, i.e., third order uncertainty. Finally, we conclude in Section 4. Proofs can be found in Appendix.

2 Information Revelation with Full Provability

2.1 The Persuasion Game

We consider two players: an interested party (player 1) and a decision maker (player 2).[2] The decision maker is interested in predicting the future value of a project, denoted by the realization of a payoff-relevant variable $\theta \in \Theta = \{\theta_0, \theta_1\} \subseteq \mathbb{R}$. Hence, he seeks to evaluate the value of θ by choosing an action $V \in \mathbb{R}$ as close as possible to θ. The interested party seeks to maximize the value V estimated by the decision maker. To do this, the interested party tries to persuade the decision maker that the true state of Nature is high by revealing him some information.

The prior probabilities of the payoff-relevant state θ are

$$\Pr(\theta = \theta_0) = \beta$$
$$\Pr(\theta = \theta_1) = 1 - \beta,$$

where $\beta \in]0, 1[$.

There is an informative signal function $s : \Theta \to \{s^0, s^1\}$ such that

$$s(\theta) = \begin{cases} s^0 & \text{if } \theta = \theta_0 \\ s^1 & \text{if } \theta = \theta_1, \end{cases} \tag{1}$$

and a uninformative signal function $\overline{s} : \Theta \to \{\overline{s}\}$. An informative signal is received by the interested party with probability $\gamma \in]0, 1[$. Thus, the interested party knows the state of Nature (the fundamental state) only if he receives one of the informative signals. Let $S \equiv \{s^0, s^1, \overline{s}\}$ be the set of possible signals received by the interested party.

[2]Shin (1994a) considered two interested parties (a defendant and a plaintiff) whereas Shin (1994b) considered only one interested party (a firm) and two decision makers (two shareholders). Almost the same logic applies with one interested party and one decision maker.

Therefore, there are four possible states of the world: either the interested party receives the signal s^0 (and thus, $\theta = \theta_0$), or he receives the signal s^1 (and thus, $\theta = \theta_1$), or he receives thus signal $\overline{s}$ (and thus, either $\theta = \theta_0$ or $\theta = \theta_1$). Denote this state space by

$$\Omega = (\Theta \times S) \backslash \{(\theta_i, s^{1-i}) : i = 0, 1\} = \{(\theta_0, s^0), (\theta_1, s^1), (\theta_0, \overline{s}), (\theta_1, \overline{s})\}.$$

Hence, the prior probability distribution over this state space (before information is received) is given by the probability measure p such that $p(\theta_i, s^i) = \gamma \Pr(\theta = \theta_i)$ and $p(\theta_i, \overline{s}) = (1 - \gamma) \Pr(\theta = \theta_i)$, for $i = 0, 1$. The interested party's partition over this state space is then given by

$$H_1 = \{\{(\theta_0, s^0)\}, \{(\theta_1, s^1)\}, \{(\theta_0, \overline{s}), (\theta_1, \overline{s})\}\},$$

whereas the decision maker's partition is simply $H_2 = \{\Omega\}$. For every $j \in \{1, 2\}$, we denote by $h_j(\omega)$ player j's information set at $\omega \in \Omega$. Such a set includes all states player j conceives as possible at ω.

Clearly, the information structure described above involves second order uncertainty since there are uncertainties about the fundamentals and uncertainties about the interested party's information about the fundamentals. However, there are no further uncertainties since it is common knowledge that the decision maker does not know whether the interested party is informed about the fundamentals or not.[3] We will relax this assumption in the next section.

Player j's utility is characterized by a function $u_j : \Omega \times \mathbb{R} \to \mathbb{R}$ such that $u_j((\theta, s), V) = u_j((\theta, s'), V)$ for all $j \in \{1, 2\}$, $\theta \in \Theta$, $V \in \mathbb{R}$, and $s, s' \in S$. The interested party's utility function u_1 is strictly increasing in V. For example, his utility function can be $u_1((\theta, s), V) = V - \theta$. The decision makers' utility is given by $u_2((\theta, s), V) = -(V - \theta)^2$. More generally, we require that V is a one to one function of the decision maker's estimate of θ.

2.2 Strategies, Beliefs, and Equilibrium

After having observed his signal $s \in S$, the interested party chooses what to disclose to the decision maker. Let $\mathcal{Y}$ be the Boolean algebra generated by the interested party's partition H_1.[4] Since announcements are assumed truthful, he chooses to reveal a message (event) $x \in \mathcal{Y}$ such that the true state ω belongs to x (or, equivalently, $h_1(\omega) \subseteq x$). In other words, the interested party reveals that the real state of the world belongs to the union of his actual information set $h_1(\omega)$ with, eventually, other of his information sets over Ω. For example, when $x = h_1(\omega)$ he reveals exactly what he knows, and when $x = \Omega$ he reveals nothing.

By denoting $Y(\omega) = \{x \in \mathcal{Y} : \omega \in x\}$ the set of events containing ω which are self-evident to player 1,[5] the interested party's *communication* or *disclosure strategy* is given by a H_1 measurable function $c : \Omega \to \mathcal{Y}$ satisfying $c(\omega) \in Y(\omega)$ for each $\omega \in \Omega$.[6]

[3] Readers interested by the formal definition of the depth of knowledge in an information structure are referred to Morris, Postlewaite, and Shin (1995). See also Hart, Heifetz, and Samet (1996).

[4] The Boolean algebra generated by a partition is the set of unions of the elements of this partition, plus the empty set. More generally, a Boolean algebra is a set of sets closed under intersection, union, and complementation.

[5] A self-evident event to a player is an event that is known by this player whenever it occurs.

[6] Mixed communication strategies are defined in the Proof of Proposition 1.

After having observed the announcement x, the decision maker chooses $V \in \mathbb{R}$ which maximizes the accuracy of his prediction. By denoting μ his belief on Ω, the optimal strategy of the decision maker is a function $\sigma : \mathcal{Y} \to \mathbb{R}$ such that

$$\sigma(x) = E_\mu(\theta \mid x) \equiv \sum_{(\theta,s)\in\Omega} \theta\mu(\theta, s \mid x), \text{ for all } x \in \mathcal{Y}. \tag{2}$$

We say that a belief μ is a *skeptical belief* if the decision maker always assumes the worst from the interested party's disclosure. Formally, μ is skeptical if

$$\mu(\omega \mid x) = \begin{cases} p(\omega \mid \{(\theta_0, s^0)\}) & \text{if } (\theta_0, s^0) \in x \\ p(\omega \mid \{(\theta_1, s^1)\}) & \text{if } x = \{(\theta_1, s^1)\} \\ p(\omega \mid \{(\theta_0, \overline{s}), (\theta_1, \overline{s})\}) & \text{if } (\theta_0, s^0) \notin x \text{ and } x \neq \{(\theta_1, s^1)\}. \end{cases} \tag{3}$$

Note that this belief satisfies the provability constraint. That is, we have $\mu(\omega \mid x) = 0$ if $\omega \notin x$. Call a communication strategy c *perfectly revealing* if $E_\mu(\theta \mid c(\omega)) = E_p(\theta \mid h_1(\omega)) = \sum_{(\theta,s)\in\Omega} \theta p(\theta, s \mid h_1(\omega))$. In other words, a communication strategy is perfectly revealing if the decision maker acquires the interested party's relevant information after the communication stage.

In the following proposition we show that there exists a sequential equilibrium in which the communication strategy is perfectly revealing and the decision maker's belief is skeptical. This result contrasts with the partially revealing equilibrium obtained by Shin (1994a,b) who assumed that only fundamental information can be disclosed. It is worth mentioning that this result does *not* depend on the specific values of β and γ.

Proposition 1 *There is a sequential equilibrium (σ, c, μ) of the persuasion game with second order uncertainty in which $\sigma(x) = \sum_{(\theta,s)\in\Omega} \theta\mu(\theta, s \mid x)$ for all $x \in \mathcal{Y}$, the communication strategy c is perfectly revealing, and μ is a skeptical belief.*

Proof. See Appendix. □

3 Payoff-Relevant Provability

Payoff-irrelevant information may or may not be proved depending on the context in which the game is played. In this section, by assuming contrary to the previous section that only disclosures about payoff-relevant information are possible, we show that no perfectly revealing equilibrium exists as long as the decision maker knows, with a probability $\alpha < 1$, that the interested party is informed about the payoff-relevant state of Nature. A partially revealing equilibrium and associated updated beliefs are also examined.

3.1 Information Structure and Communication Possibilities

Now, the decision maker can receive a signal informing him that the interested party is informed about the payoff-relevant states. Formally, let $k : S \to \{k_1, k_{\overline{1}}\}$ be the informative signal of the decision maker such that

$$k(s) = \begin{cases} k_1 & \text{if } s \in \{s^0, s^1\} \\ k_{\overline{1}} & \text{if } s = \overline{s}, \end{cases} \tag{4}$$

and let $\overline{k} : S \to \{\overline{k}\}$ be his uninformative signal. When the signal is k_1 the decision maker knows that the interested party knows that the state of Nature is either $\theta = \theta_0$ or $\theta = \theta_1$. When the signal is $k_{\overline{1}}$ the decision maker knows that the interested party is uninformed. Finally, if the signal is $\overline{k}$, i.e., if the decision maker receives no signal, then he does not know whether the interested party is informed about the state of Nature θ or not. The informative signal is received by the decision maker with probability $\alpha \in]0, 1[$.

Let $K = \{k_1, k_{\overline{1}}, \overline{k}\}$ be the set of possible signals of the decision maker. Now, there are eight possible states of the world.[7] The state space is

$$\begin{aligned}\Omega &= (\Theta \times S \times K)\Big\backslash\Big(\{(\theta_i, s^{1-i}, k) : i = 0, 1,\ k \in K\} \\ &\quad \cup \{(\theta_i, s^i, k_{\overline{1}}) : i = 0, 1\} \cup \{(\theta_i, \overline{s}, k_1) : i = 0, 1\}\Big) \\ &= \big\{(\theta_0, s^0, k_1), (\theta_1, s^1, k_1), (\theta_0, s^0, \overline{k}), (\theta_1, s^1, \overline{k}), (\theta_0, \overline{s}, k_{\overline{1}}), (\theta_1, \overline{s}, k_{\overline{1}}), \\ &\quad (\theta_0, \overline{s}, \overline{k}), (\theta_1, \overline{s}, \overline{k})\big\}.\end{aligned}$$

The interested party's partition over Ω is now

$$H_1 = \big\{\big\{(\theta, s, k) \in \Omega : s = s'\big\} : s' \in S\big\}\,,$$

and the decision maker's partition is

$$H_2 = \big\{\big\{(\theta, s, k) \in \Omega : k = k'\big\} : k' \in K\big\}\,.$$

To simplify the notations, let $\{S^0, S^1, \overline{S}\} \equiv H_1$, $\{K_1, \overline{K}, K_{\overline{1}}\} \equiv H_2$, $\Theta_0 \equiv \{(\theta, s, k) \in \Omega : \theta = \theta_0\}$, and $\Theta_1 \equiv \{(\theta, s, k) \in \Omega : \theta = \theta_1\}$.[8] Hence, the event Θ_i is "the state of Nature is θ_i".

The utility functions and the interested party's communication strategies are defined as in the previous section. However, to allow only payoff-relevant information to be disclosed we will restrain the set of possible messages (events) to payoff-relevant messages. That is, we consider the set of messages $\mathcal{X} \subseteq \mathcal{Y}$ which only includes payoff-relevant events, i.e., $\mathcal{X} = \{S^0, S^1, \Omega\}$. In particular, the interested party cannot reveal that he does not know the fundamentals since $\overline{S} \notin \mathcal{X}$. In this case, a communication strategy is a H_1 measurable function $c : \Omega \to \mathcal{X}$ such that $c(\omega) \in X(\omega) = \{x \in \mathcal{X} : \omega \in x\}$.

3.2 Equilibrium

The decision maker's beliefs and strategies are slightly more difficult to characterize than in the previous section because they also depend on his own information. The decision maker's belief at ω is a conditional distribution $\mu(\cdot \mid x, h_2(\omega))$ when $x \in \mathcal{X}$ has been revealed by the interested party. His strategy is a H_2 measurable function $\sigma : \mathcal{X} \times \Omega \to \mathbb{R}$. Clearly, this strategy is optimal iff $\sigma(x, \omega) = E_\mu(\theta \mid x, h_2(\omega)) = \sum_{(\theta,s,k)\in\Omega} \theta\mu(\theta, s, k \mid x, h_2(\omega))$ for all $\omega \in \Omega$ and $x \in X(\omega)$.

[7] Either the interested party knows the payoff-relevant state and the decision maker knows that he knows, or the interested party knows the payoff-relevant state and the decision maker does not know that he knows, or the interested party does not know the payoff-relevant state and the decision maker knows that he does not know, or the interested party does not know the payoff-relevant state and the decision maker does not know that he does not know.

[8] It can be verified that this information structure has depth of knowledge two in the sense of Morris et al. (1995).

In the next proposition we show that no perfectly revealing sequential equilibrium exists for any $\alpha, \beta, \gamma \in]0,1[$.

Proposition 2 *There is* no *sequential equilibrium* (σ, c, μ) *of the persuasion game with third order uncertainty and payoff-relevant provability in which a perfectly revealing communication strategy* c *is used.*

Proof. See Appendix. □

In the following lines we characterize a partially revealing equilibrium. A communication strategy is called a *sanitization communication strategy* (Shin, 1994a,b) if the interested party reveals only favorable information, and conceals unfavorable one. In our model it can be formally characterized by

$$c(\omega) = \begin{cases} \Omega & \text{if } \omega \in S^0 \cup \overline{S} \\ S^1 & \text{if } \omega \in S^1. \end{cases}$$

In this setting we will say that a belief μ is *partially skeptical* if it satisfies Bayes' rule against a sanitization communication strategy, and if it assigns probability one to S^0 when the decision maker knows that the interested party is informed but reveals nothing. Formally, μ is a partially skeptical belief if it satisfies the following conditions:

$$\mu(\theta_0, s^0, k_1 \mid x, K_1) = \begin{cases} 1 & \text{if } x \in \{S^0, \Omega\} \\ 0 & \text{if } x = S^1, \end{cases} \tag{5}$$

$$\mu(\theta_0, s^0, \overline{k} \mid x, \overline{K}) = \begin{cases} 1 & \text{if } x = S^0 \\ 0 & \text{if } x = S^1 \end{cases} \tag{6}$$

$$\mu(\omega \mid \Omega, K_{\overline{1}}) = \frac{p(\omega)}{p(K_{\overline{1}})} \quad \text{if } \omega \in K_{\overline{1}}, \tag{7}$$

$$\mu(\omega \mid \Omega, \overline{K}) = \frac{p(\omega)}{p(\overline{K} \cap (S^0 \cup \overline{S}))} \quad \text{if } \omega \in \overline{K} \cap (S^0 \cup \overline{S}). \tag{8}$$

It is to be noticed that the only possible revelation of the interested party at $\omega \in \overline{S}$ is Ω since only payoff-relevant information can be revealed and the interested party has no payoff-relevant information at $\omega \in \overline{S}$. Hence, the only observable outside equilibrium move from a sanitization strategy occurs when the interested party reveals S^0. Otherwise, Bayes' rule directly applies to determine the decision maker's inferences.

Proposition 3 *There is a sequential equilibrium* (σ, c, μ) *of the persuasion game with third order uncertainty and payoff-relevant provability in which* $\sigma(x, \omega) = \sum_{(\theta,s,k)\in\Omega} \theta\mu(\theta, s, k \mid x, h_2(\omega))$ *for all* $\omega \in \Omega$ *and* $x \in X(\omega)$*, the communication strategy* c *is the sanitization strategy, and* μ *is a partially skeptical belief.*

Proof. See Appendix. □

3.3 Inferences Beliefs at the Partially Revealing Equilibrium

Along the equilibrium path of the partially revealing equilibrium, it is easy to see that, after the communication stage, the decision maker's partition over Ω becomes

$$H_2^* = \left\{ \{(\theta_0, s^0, k_1)\}, \{(\theta_1, s^1, k_1)\}, \{(\theta_1, s^1, \overline{k})\}, \overline{K} \cap (S^0 \cup \overline{S}), K_{\overline{1}} \right\}.$$

Hence, he knows the payoff-relevant state whenever he initially knew that the interested party was informed ($k = k_1$), and when $\theta = \theta_1$ and $s \neq \overline{s}$. If he knew that the interested party was uninformed, then he keeps prior probabilities about θ since $p(\Theta_0 \mid K_{\overline{1}}) = \beta$ and $p(\Theta_1 \mid K_{\overline{1}}) = 1 - \beta$. The interesting configuration corresponds to the case in which the decision maker does not know whether the interested party is informed and when nothing was revealed (i.e., when $\omega \in \overline{K} \cap (S^0 \cup \overline{S})$). In this case, the decision maker's beliefs about the payoff-relevant events are

$$\begin{aligned}
\mu(\Theta_0 \mid \Omega, \overline{K}) &= p(\Theta_0 \mid \overline{K} \cap (S^0 \cup \overline{S})) &&= \frac{p(\theta_0, s^0, \overline{k}) + p(\theta_0, \overline{s}, \overline{k})}{p(\theta_0, s^0, \overline{k}) + p(\theta_0, \overline{s}, \overline{k}) + p(\theta_1, \overline{s}, \overline{k})} \\
& &&= \frac{\beta}{\beta + (1-\beta)(1-\gamma)} \\
\mu(\Theta_1 \mid \Omega, \overline{K}) &= p(\Theta_1 \mid \overline{K} \cap (S^0 \cup \overline{S})) &&= \frac{(1-\beta)(1-\gamma)}{\beta + (1-\beta)(1-\gamma)},
\end{aligned}$$

which do not depend on α. We remark that if γ increases, i.e., if the average precision of the interested party's information increases and if he reveals nothing, then the decision maker puts more weight on his belief about Θ_0. As a consequence, the value V estimated by the decision maker is decreasing with the average precision of the interested party's information.

4 Conclusion

The main lesson of this paper is that full revelation occurs and is robust in traditional persuasion games as long as fundamental *and* non-fundamental events can be disclosed. On the contrary, if only payoff-relevant information can be disclosed, then full revelation will not occur whenever there is a strictly positive probability that the decision maker does not know that the interested party is informed. In this case, the interested party's communication strategy is to reveal favorable information, and the decision maker learns only partially the fundamentals. As a consequence, even without specifying the information structure and higher order uncertainties into details, predictions are stable: we only need to know the configuration of provability. In other words, predictions are *not* sensitive to the information structure and to uncertainty concerning what the informed party actually knows, but they are very sensitive to the structure of provability. An advantage is that this latter detail is often available in economic or legal situations we consider, contrary to details concerning the precise information that players possess when they are called to make a choice.

Appendix. Proofs

Proof of Proposition 1. The proof of Proposition 1 is made in two steps. First, we show that the interested party has no incentive to deviate from the perfectly revealing communication strategy c satisfying $c(\omega) = h_1(\omega)$ for all $\omega \in \Omega$. Then, we show that if μ is skeptical, then (σ, c, μ) satisfies the consistency condition of Kreps and Wilson (1982). The fact that σ is an optimal strategy for the decision maker is obvious.

When the interested party receives the signal s^0, he must necessarily send a message $x \ni (\theta_0, s^0)$ because $(\theta_0, s^0) \in x$ for all $x \in Y(\theta_0, s^0)$ by the definition of the function Y (he can only reveal information he actually possesses). Therefore, since $\mu(\omega \mid x) = \mu(\omega \mid \{(\theta_0, s^0)\})$ for all $x \ni (\theta_0, s^0)$ from the skeptical belief (3), the decision maker's belief (and hence his decision) does not change when the interested party deviates. Afterwards, note that

$$\sum_{(\theta,s)\in\Omega} \theta\mu(\theta, s \mid \{(\theta_1, s^1)\}) = \theta_1, \qquad \sum_{(\theta,s)\in\Omega} \theta\mu(\theta, s \mid \{(\theta_0, s^0)\}) = \theta_0,$$

$$\sum_{(\theta,s)\in\Omega} \theta\mu(\theta, s \mid \{(\theta_0, \overline{s}), (\theta_1, \overline{s})\}) = \theta_0 \frac{p(\theta_0, \overline{s})}{p(\{(\theta_0, \overline{s}), (\theta_1, \overline{s})\})} + \theta_1 \frac{p(\theta_1, \overline{s})}{p(\{(\theta_0, \overline{s}), (\theta_1, \overline{s})\})},$$

$$\Longrightarrow \quad \sigma(\{(\theta_1, s^1)\}) > \sigma(\{(\theta_0, \overline{s}), (\theta_1, \overline{s})\}) > \sigma(\{(\theta_0, s^0)\}),$$

because $\beta \in]0, 1[$. Moreover, from (3) we have

$$\sigma(\{(\theta_0, s^0)\}) = \sigma(\{(\theta_0, s^0), (\theta_1, s^1)\}) = \sigma(\{(\theta_0, s^0), (\theta_0, \overline{s}), (\theta_1, \overline{s})\}) = \sigma(\Omega),$$

and

$$\sigma(\{(\theta_0, \overline{s}), (\theta_1, \overline{s})\}) = \sigma(\{(\theta_1, s^1), (\theta_0, \overline{s}), (\theta_1, \overline{s})\}).$$

Therefore, when $s = s^1$ or $s = \overline{s}$, the interested party cannot increase the decision maker's evaluation by revealing $x \in \mathcal{Y}$, with $(\theta, s) \in x$.

To prove that (σ, c, μ) is consistent, we define a sequence of perturbed persuasion games in which the interested party reveals every possible message (or provable event) with positive probability. We then show that the decision maker's belief obtained via Bayes' rule tends to the skeptical belief μ given by (3). To allow perturbed games we have to define mixed communication strategies. A mixed communication strategy is a H_1 measurable function $\pi : \Omega \to \Delta(\mathcal{Y})$ such that $\mathrm{supp}(\pi(\omega)) \subseteq Y(\omega)$ for all $\omega \in \Omega$, where $\mathrm{supp}(\pi(\omega))$ is the support of the probability distribution $\pi(\omega)$. Let Π be the set of mixed communication strategies and let Π^0 be the set of all strictly positive mixed communication strategies, i.e., $\Pi^0 \equiv \{\pi \in \Pi : \pi(x \mid \omega) > 0 \text{ for all } \omega \in \Omega \text{ and } x \in Y(\omega)\}$. If $\pi^t \in \Pi^0$, then we can associate a belief μ^t with π^t via Bayes' rule. More precisely, if $\pi^t \in \Pi^0$, then for all $\omega \in \Omega$ and $x \in Y(\omega)$ we have

$$\mu^t(\omega \mid x) = \begin{cases} 0 & \text{if } \omega \notin x \\ \dfrac{\pi^t(x \mid \omega) p(\omega)}{\sum_{\omega' \in \Omega} p(\omega') \pi^t(x \mid \omega')} & \text{otherwise.} \end{cases} \tag{9}$$

Let $\pi \in \Pi$ be the mixed communication strategy associated with the perfectly revealing communication strategy c, i.e., $\pi(c(\omega) \mid \omega) = 1$ for all $\omega \in \Omega$. Let μ be the skeptical belief defined by (3). We have to show that there exists a sequence $\{(\mu^t, \pi^t)\}_t$ such that $\pi^t \in \Pi^0$ for all t, μ^t is associated with π^t via Bayes' rule, and

$(\mu, \pi) = \lim_{t\to\infty}(\mu^t, \pi^t)$. Let $\{\varepsilon^t\}_t$ be a sequence such that $\lim_{t\to\infty} \varepsilon^t = 0$. We consider the following "trembling" communication strategy:

$$\pi^t(x \mid \{(\theta_0, s^0)\}) = \begin{cases} \varepsilon^t & \text{if } x \in Y(\theta_0, s^0),\ x \neq \{(\theta_0, s^0)\} \\ 1 - 3\varepsilon^t & \text{if } x = \{(\theta_0, s^0)\}, \end{cases}$$

$$\pi^t(x \mid \{(\theta_0, \overline{s})\}) = \pi^t(x \mid \{(\theta_1, \overline{s})\}) =$$

$$\begin{cases} (\varepsilon^t)^2 & \text{if } x \in Y(\theta_0, \overline{s}),\ x \neq \{(\theta_0, \overline{s}), (\theta_1, \overline{s})\} \\ 1 - 3(\varepsilon^t)^2 & \text{if } x = \{(\theta_0, \overline{s}), (\theta_1, \overline{s})\}, \end{cases}$$

and

$$\pi^t(x \mid \{(\theta_1, s^1)\}) = \begin{cases} (\varepsilon^t)^3 & \text{if } x \in Y(\theta_1, s^1),\ x \neq \{(\theta_1, s^1)\} \\ 1 - 3(\varepsilon^t)^3 & \text{if } x = \{(\theta_1, s^1)\}. \end{cases}$$

Form Bayes' rule (9) it is not difficult to verify that $\lim_{t\to\infty}(\mu^t, \pi^t) = (\mu, \pi)$, i.e., (σ, c, μ) is consistent.

□

Proof of Proposition 2. Assume on the contrary that there is a perfectly revealing sequential equilibrium. This implies that $c(\omega) = h_1(\omega)$ for all $\omega \in S^0 \cup S^1$ and $c(\omega) = \Omega$ for all $\omega \in \overline{S}$ (remember that $X(\omega) = \{\Omega\}$ if $\omega \in \overline{S}$ because of the restriction to payoff-relevant revelations). We show that for any consistent belief μ, the interested party deviates at $\omega \in S^0$ and reveals Ω instead of S^0.

When the interested party reveals S^0 at $\omega \in S^0$ his expected utility is

$$\begin{aligned} & p(\theta_0, s^0, k_1 \mid S^0)\sigma(S^0, (\theta_0, s^0, k_1)) + p(\theta_0, s^0, \overline{k} \mid S^0)\sigma(S^0, (\theta_0, s^0, \overline{k})) \\ = \ & p(\theta_0, s^0, k_1 \mid S^0)\theta_0 + p(\theta_0, s^0, \overline{k} \mid S^0)\theta_0 = \theta_0. \end{aligned}$$

When he reveals Ω at $\omega \in S^0$ his expected utility is

$$\begin{aligned} & p(\theta_0, s^0, k_1 \mid S^0)\sigma(\Omega, (\theta_0, s^0, k_1)) + p(\theta_0, s^0, \overline{k} \mid S^0)\sigma(\Omega, (\theta_0, s^0, \overline{k})) \\ = \ & p(\theta_0, s^0, k_1 \mid S^0)(\mu(\theta_0, s^0, k_1 \mid \Omega)\theta_0 + \mu(\theta_1, s^1, k_1 \mid \Omega)\theta_1) \\ & + p(\theta_0, s^0, \overline{k} \mid S^0)\left(\frac{p(\theta_0, \overline{s}, \overline{k})}{p(\overline{K} \cap \overline{S})}\theta_0 + \frac{p(\theta_1, \overline{s}, \overline{k})}{p(\overline{K} \cap \overline{S})}\theta_1\right), \text{ by Bayes' rule,} \\ \geq \ & p(\theta_0, s^0, k_1 \mid S^0)\theta_0 + p(\theta_0, s^0, \overline{k} \mid S^0)\left(\frac{p(\theta_0, \overline{s}, \overline{k})}{p(\overline{K} \cap \overline{S})}\theta_0 + \frac{p(\theta_1, \overline{s}, \overline{k})}{p(\overline{K} \cap \overline{S})}\theta_1\right) \\ > \ & \theta_0, \text{ because } \frac{p(\theta_0, \overline{s}, \overline{k})}{p(\overline{K} \cap \overline{S})} = \beta \neq 0 \text{ and } p(\theta_0, s^0, \overline{k} \mid S^0) = 1 - \alpha \neq 0. \end{aligned}$$

Hence, the interested party always deviate from full revelation at $\omega \in S^0$. □

Proof of Proposition 3. We have to check for sequential rationality of the interested party and we must prove that (σ, c, μ) is consistent. Let c be the sanitization communication strategy. First note that the interested party cannot deviate at $\omega \in \overline{S}$ since $X(\omega) = \{\Omega\}$ for all $\omega \in \overline{S}$. If he deviates and reveals Ω at $\omega \in S^1$, then his expected utility strictly decreases since

$$\begin{aligned} & p(\theta_1, s^1, k_1 \mid S^1)\sigma(\Omega, (\theta_1, s^1, k_1)) + p(\theta_1, s^1, \overline{k} \mid S^1)\sigma(\Omega, (\theta_1, s^1, \overline{k})) \\ = \ & p(\theta_1, s^1, k_1 \mid S^1)\theta_0 + p(\theta_1, s^1, \overline{k} \mid S^1)\big(p(\Theta_0 \mid \overline{K} \cap (\overline{S} \cup S^0))\theta_0 \\ & + p(\Theta_1 \mid \overline{K} \cap (\overline{S} \cup S^0))\theta_1\big) \\ < \ & \theta_1 = p(\theta_1, s^1, k_1 \mid S^1)\sigma((\theta_1, s^1, k_1), S^1) + p(\theta_1, s^1, \overline{k} \mid S^1)\sigma((\theta_1, s^1, \overline{k}), S^1). \end{aligned}$$

Similarly, the interested party does not deviate at $\omega \in S^0$ because

$$\begin{aligned}
& p(\theta_0, s^0, k_1 \mid S^1)\sigma(\Omega, (\theta_0, s^0, k_1)) + p(\theta_0, s^0, \overline{k} \mid S^1)\sigma(\Omega, (\theta_0, s^0, \overline{k})) \\
= \; & p(\theta_0, s^0, k_1 \mid S^1)\theta_0 + p(\theta_0, s^0, \overline{k} \mid S^1)\big(p(\Theta_0 \mid \overline{K} \cap (\overline{S} \cup S^0))\theta_0 \\
& + p(\Theta_1 \mid \overline{K} \cap (\overline{S} \cup S^0))\theta_1\big) \\
> \; & \theta_0 = p(\theta_0, s^0, k_1 \mid S^0)\sigma(S^0, (\theta_0, s^0, k_1)) + p(\theta_0, s^0, \overline{k} \mid S^0)\sigma(S^0, (\theta_0, s^0, \overline{k})).
\end{aligned}$$

To show that the combination of partially skeptical beliefs and the sanitization communication strategy is consistent it suffices to consider perturbed games, as in the Proof of Proposition 1, with the following "trembling" communication strategy:

$$\text{If } \omega \in S^0, \text{ then } \pi^t(x \mid \omega) = \begin{cases} 1 - \varepsilon^t & \text{if } x = \Omega \\ \varepsilon^t & \text{if } x = S^0, \end{cases}$$

$$\text{If } \omega \in S^1, \text{ then } \pi^t(x \mid \omega) = \begin{cases} (\varepsilon^t)^2 & \text{if } x = \Omega \\ 1 - (\varepsilon^t)^2 & \text{if } x = S^1, \end{cases}$$

and, of course, $\pi^t(\Omega \mid \omega) = 1$ if $\omega \in \overline{S}$. Again, we can verify that $\lim_{t\to\infty}(\mu^t, \pi^t) = (\mu, \pi)$, where μ is a partially skeptical belief and π is the mixed communication strategy associated with the sanitization strategy. □

References

Grossman, S. J. (1981): "The Informational Role of Warranties and Private Disclosure about Product Quality," *Journal of Law and Economics*, 24, 461–483.

Grossman, S. J. and O. D. Hart (1980): "Disclosure Laws and Takeover Bids," *Journal of Finance*, 35.

Hart, S., A. Heifetz, and D. Samet (1996): "'Knowing Whether', 'Knowing That', and the Cardinality of State Spaces," *Journal of Economic Theory*, 70, 249–256.

Kreps, D. M. and R. Wilson (1982): "Sequential Equilibria," *Econometrica*, 50, 863–894.

Milgrom, P. (1981): "Good News and Bad News: Representation Theorems and Applications," *Bell Journal of Economics*, 12, 380–391.

Morris, S., A. Postlewaite, and H. S. Shin (1995): "Depth of Knowledge and the Effect of Higher Order Uncertainty," *Economic Theory*, 6, 453–467.

Okuno-Fujiwara, A., M. Postlewaite, and K. Suzumura (1990): "Strategic Information Revelation," *Review of Economic Studies*, 57, 25–47.

Shin, H. S. (1994a): "The Burden of Proof in a Game of Persuasion," *Journal of Economic Theory*, 64, 253–264.

——— (1994b): "News Management and the Value of Firms," *Rand Journal of Economics*, 25, 58–71.