



Cowles Foundation

for Research in Economics
at Yale University

Cowles Foundation Discussion Paper No. 1743

SELLING INFORMATION

Johannes Hörner and Andrzej Skrzypacz

December 2009
Revised August 2010

An author index to the working papers in the
Cowles Foundation Discussion Paper Series is located at:
<http://cowles.econ.yale.edu/P/au/index.htm>

This paper can be downloaded without charge from the
Social Science Research Network Electronic Paper Collection:
<http://ssrn.com/abstract=1526606>

Selling Information*

Johannes Hörner[†] and Andrzej Skrzypacz[‡]

June 24, 2010

Abstract

We study a dynamic buyer-seller problem in which the good is information and there are no property rights. The potential buyer is reluctant to pay for information whose value to him is uncertain, but the seller cannot credibly convey its value to the buyer without disclosing the information itself. Information comes as divisible hard evidence. We show how and why the seller can appropriate a substantial fraction of the value through gradual revelation, and how the entire value can be extracted with the help of a mediator.

Keywords: value of information, dynamic game.

JEL codes: C72, D82, D83

1 Introduction

In the absence of intellectual property rights, it is difficult for the possessor of private information that is relevant to others' decisions to appropriate its value. First, the potential buyer must

*We thank Daron Acemoglu, Drew Fudenberg, David Kreps, Max Kwiek, R.Vijay Krishna, Romans Pancs, Arthur Robson and seminar participants at the Barcelona JOCS, the Collegio Carlo Alberto, Turin, Essex University, the European University Institute, Harvard-MIT, Illinois at Urbana-Champaign, Oxford, Simon Fraser University, Stanford, University of British Columbia, University of Western Ontario, UC San Diego, Yale, X-HEC Paris and the SED 2009 meeting for useful comments and suggestions.

[†]Yale University, 30 Hillhouse Ave., New Haven, CT 06520, USA. johannes.horner@yale.edu.

[‡]Stanford University, Graduate School of Business, 518 Memorial Way, Stanford, CA 94305, USA. andy@gsb.stanford.edu.

know that the seller really does possess relevant information. As importantly, how much this information is worth typically depends on its nature. Suppose, for instance, that the information is some bad news regarding the opportunity to make an investment that the buyer would abstain from making in the absence of any news. In this case, the nature of the information is such that it is not worth anything to the buyer, because having it will not affect what he decides to do. Now suppose that the information is some good news about this investment. In this case, the information is of value to the buyer. However, a problem arises. In order to sell the information, the seller must somehow convey to the buyer that it is of value. Yet how can the seller convey the fact that the information is of value to the buyer without giving away at least part of it, and hence part of its value? At the very least, the seller must convey that the information is good news. However, the buyer may not care about the content of this information beyond the mere fact that it is good news, i.e. that investing is the optimal decision. In this case, if the seller discloses that she has such good news, assuming she can do so, the buyer will no longer be willing to pay anything for it.

This conundrum is known as Arrow's *information paradox* (Arrow, 1959). Information, unlike physical goods, cannot be taken away. Yet rewarding the possession of knowledge appropriately is essential to the fostering of innovative activity and to the provision of incentives for the acquisition of information and expertise.

This paper solves for the best sales strategy, and the maximal sales revenue, of the seller of information when information is (i) *verifiable*, and (ii) *divisible*. To say that the information is verifiable means that whatever information is released is credible and cannot be faked. Producing a working prototype at an industrial fair, market research regarding the profitability of a given investment, or exploration data regarding the size of an oil field are all examples of pieces of information that may be difficult to forge. To say that the information is divisible means that it can be conveyed in parts that constitute evidence that the information has value. The working prototype, the favorable market research, and positive exploration data are all encouraging signs that the seller's information has value, but they are not necessarily the end of the story. That is, they increase the subjective probability assigned by the buyer to the event that the information has value to him, thereby containing good news for him, although they need not lead him to rule

out just yet the possibility that the news is actually bad. More formally, information is modeled as a collection of bits, and the belief of the buyer that news is good increases in the number and importance of the bits that are disclosed by the seller.

As mentioned above, we are interested in the case in which there are no intellectual property rights. That is, information is non-contractible. The buyer cannot be coerced into making payments, and the seller cannot be coerced into disclosing information (that she may not possess anyway). Further, our seller cares neither about the decision that the buyer takes, nor about the information *per se*. All she cares about is maximizing the total payment she can get for it. The buyer must trade off the payments that he must make with the potential value of this information in his investment choice.

Our objective is to identify which strategy maximizes the profit of a seller who possesses good news. This is quite natural if we think of this news as an invention, and it also maximizes the incentive to acquire information in the first place, as we shall see in Section 5. However, most of our analysis takes this as a premiss and explores how information should be revealed over time. In our model, the seller is an Agent who can either know that the state is good, in which case the buyer, a Firm, should optimally invest in a venture that he is considering, or she knows that the state is bad, in which case the Firm should not, which is also what he would do without any further information. Producing a piece of evidence in favor of investment might not rule out the other state of the world entirely. That is, it is possible, although unlikely, that some partial evidence in favor of the good state can be provided by the Agent, who knows that the state is actually bad. In each round of our finite-horizon game, the Firm and Agent simultaneously make payments to each other first, and then the Agent discloses some information if she wishes to. After the last round, the Firm chooses to invest or not. We ask what equilibrium maximizes the payoff of the Agent who knows that the state is good. We obtain two results.

First, we show that *revealing information over time is valuable*. The possessor of good news gains from releasing it slowly over time. In our leading example, we show that the release should involve an initial burst of information given away for free, followed by gradual information sales. While the Firm is actually indifferent between this procedure and paying the expected value of the information in one shot, the Agent who has good news gains from such timing. These gains,

then, must come at the expense of the Agent who possesses bad news. In this example, we obtain an explicit expression for the profit of the Agent who possesses good news, and the strategy that she should follow. In general, we show that her profit can be expressed as an integral of the difference between the Firm's average value and marginal value for information, over all beliefs between the Firm's prior belief and one. This profit is larger than what she could hope for in a one-shot game, but it still falls short of the full value of this information to the buyer.

Second, we show that *with the help of a mediator, the seller can extract the full value of the information*. Here, a mediator, or intermediary, should be interpreted as being a disinterested and trusted third party; it may be a person, but it might just as well be a computer program. The authority of such a mediator is minimal: it is able to perform randomizations at the stage of information disclosure without it being the case that the seller is indifferent over all elements in its support (that is, it makes mixed actions observable). In addition, because it is trusted, the buyer might change his belief in response to its messages, without the mediator releasing any actual hard information (for example, if the mediator states that the seller has some hard information, the buyer will believe it).

Surprisingly, the optimal scheme also involves payments by the seller to the buyer. Loosely speaking, a seller who possesses good news and the buyer disagree on the likely evolution of the buyer's future belief (that news is good). Whereas the possessor of good news anticipates that she will disclose information that will make the buyer more optimistic over time about the relevance of the information to an investment that he is considering, the buyer expects to become neither more optimistic, nor pessimistic, on average. Due to the fact that they have different beliefs, there is room for profitable trade between the buyer and the seller of good news. This trade corresponds to a bet on the change in the buyer's belief and involves payments that can go both ways. Again, there is a loser from such trades, namely, a seller who possesses bad news, and this limits the scope for trade: the bets are implemented as voluntary payments contingent upon the signals revealed by the intermediary and have to be small enough that the seller does not renege on paying out. Due to the fact that these trades cannot be informative *per se* about the type of the seller, the seller of bad news must prefer going along with such schemes rather than being found out.

Nevertheless, as mentioned, the seller of good news is able, quite generally, to achieve a profit equal to the actual value of this information to the buyer, despite the fact that, if the buyer knew exactly how much this information were worth to him, he would know the information itself, and would not be willing to pay a dime for it!

This paper is related to several strands of the literature. In terms of one of its most important economic applications, the sale of information about an invention by its inventor, it is close to Anton and Yao (1994 and 2002). In Anton and Yao (1994), an inventor possesses private information regarding the value of an invention, and as in our paper, attempts to sell it at the highest price in the absence of any intellectual property rights. The authors show how competition among firms can provide the inventor with the necessary leverage to negotiate favorable terms with a firm, even after revealing this information to the firm privately. Due to the fact that the agent can still disclose this information to the firm's competitors, she is able to appropriate a sizable share of the invention's market value. This explanation is complementary to ours, inasmuch as it identifies competition as an alternative channel that mitigates the risk of expropriation. Anton and Yao (2002) introduces the further possibility that the inventor publicly discloses once, at the beginning of the game (before any payments occur), some partial information. The two buyers then offer contracts, among which the seller chooses. This is a signaling game, in which the spread in the contractual payments provides differential incentives to disclose information to the firms, depending on whether or not the buyer's innovation is eventually successful. The success of the innovation is a verifiable event, the probability of which increases in proportion to the inventor's ability and skill. In contrast, in our model, there are no contingent payments, and therefore, no differential incentives can be provided.

The gradualism that appears in equilibrium is related to findings of the literature on contribution games. See, for instance, Admati and Perry (1991), Marx and Matthews (2000) and Compte and Jehiel (2004). See also Gul (2001) and Che and Sákovics (2004) for the dynamic resolution of the hold-up problem. Indeed, there are similarities between contributing to a public good and disclosing information. In both cases, concessions are irreversible. However, unlike in the public goods case, there is a strong asymmetry between the players in our case and the payoff structure is quite different. As a result, some of the findings are different as well. For instance, in

the literature on contribution games, there is no counterpart to what happens in the first round in our leading example, in which a big chunk of information is given away for free. Gradualism appears in our case only after the first round.

The formal maximization problem, and in particular the structural constraints on the revealing of information, are reminiscent of the literature on long cheap talk. See, in particular, Forges (1990) and Aumann and Hart (2003), and, more generally, Aumann and Maschler (1995). As is the case here, the problem is how to “split” a martingale optimally over time. That is, the Firm’s belief is a martingale, and the optimal strategy specifies its distribution over time. The similarity is especially clear when there is a mediator. There are important differences, however. In particular, unlike in the literature on long cheap-talk, payoff-relevant actions are taken before information disclosure is finished, because the Firm pays the Agent as information is revealed over time. In fact, with a mediator, the Agent also makes payments to the Firm during the communication phase. As in Forges and Koessler (2008), messages are type-dependent, in that the Agent is constrained in the messages she can send by the information she actually possesses. Cheap talk, that is, the possibility of sending payoff-irrelevant messages from sets that are type-independent, is of no help in our model. Rosenberg, Solan and Vieille (2009) consider the problem of information exchange between two informed parties in a repeated game without transfers, and establish a folk theorem. In all these papers, the focus is on identifying the best equilibrium from the Agent’s perspective, in the *ex ante* sense, i.e. before her type is known. In our case, this is trivial (see Lemma 1.1) and does not deliver differential payoffs to the Agent’s types. Therefore, such an equilibrium does not provide the Agent with incentives to engage in research activities in the first place. To do so requires identifying the best equilibrium from the point of view of a particular type of the Agent.

Finally, there is a vast literature directly related to the value of information. See, among others, Admati and Pfleiderer (1988 and 1990). Esó and Szentes (2007) take a mechanism design approach to this problem, while Gentzkow and Kamenica (2009) apply ideas similar to Aumann and Maschler (1995) to study the optimal disclosure of information in the standard (two-period) signaling model. There is, finally, a fascinating literature in computer science concerned with the possibility of (probabilistically) proving one’s knowledge of some fact without disclosing

this fact. This literature on zero-knowledge proofs, which starts with Goldwasser, Micali and Rackoff's (1985) paper is too large to survey here.

2 Splitting Information: A Simple Example

Suppose that the Agent may possess information that is valuable to another party, the Firm. More precisely, she either knows nothing, that is, she just knows (that the true state of the world lies in) Ω , or some relatively coarse information (say, $A \subseteq \Omega$), or some rather precise information (say, $B \subsetneq A$). In particular, if the Agent knows B , she could choose to disclose A if she wishes. The nature of this information is irrelevant. What matters here is that: (i) the possible types of the Agent, i.e., the information she may have, admit a natural ordering, from a least informed type to a most informed type; (ii) the Firm's payoff, as a function of the information he obtains, is increasing in this ordering. That is, denoting by v_Ω, v_A and v_B the payoff that the Firm can secure if he acquires the corresponding information, we assume that $v_\Omega \leq v_A < v_B$, and normalize v_Ω to zero, which is also the payoff that the Firm obtains if the Agent refuses to disclose any information. Finally, it is also assumed here that the Agent does not care about the information that she discloses *per se*, but only about the payments she can obtain from the Firm in exchange for it. The probabilities with which the Agent is of a given type are common knowledge. Let p_A and p_B denote these probabilities (so that the Agent is uninformed with probability $1 - p_A - p_B$).

Suppose first that the Agent has the opportunity to sell this information to the Firm in one attempt. How much is the Firm willing to pay for the Agent's potential information, assuming that the Agent will release all she knows if and only if the Firm makes this payment? Obtaining the information is worth in expectation

$$p_A v_A + p_B v_B$$

to the Firm, while not receiving any information is worth zero to the Firm. Therefore, this value is also what the Firm is willing to pay for full disclosure.

Let us now assume instead that information is disclosed in two stages (without discounting), with payments being made before each disclosure. More precisely, assume the following: (i) the

Firm makes the appropriate payment, so the Agent discloses that she knows A , if she is able to so disclose it; then (ii) the Firm makes the second payment, so the Agent discloses B . How much is the Firm willing to pay at each stage? Let us solve the game backwards: if A is disclosed, the Firm assigns probability $\mathbb{P}[\text{type } B \mid \text{type } A \text{ or type } B]$, or $\mathbb{P}[B \mid A \text{ or } B]$, for short, to the event that the Agent also knows B . Such knowledge would be worth v_B , but given that he knows A already, the Firm can secure v_A anyhow. Therefore, the Firm is willing to pay up to

$$\mathbb{P}[B \mid A \text{ or } B](v_B - v_A)$$

up front for the possibility of additional information. The best for the Agent is then to ask for that much. However, this leaves the Firm with a net continuation payoff of v_A at the beginning of the second round. Therefore, at the beginning of the first round, the Firm is willing to pay at most

$$(p_A + p_B)v_A$$

for the possibility of receiving this first piece of information. All in all, the Firm expects to pay

$$(p_A + p_B)[v_A + \mathbb{P}[B \mid A \text{ or } B](v_B - v_A)] = p_A v_A + p_B v_B$$

over the course of the two rounds. Therefore, the Firm is indifferent over the two scenarios.

Now consider the fully informed Agent, who knows B . In the first scenario, she gets a single payment. However, in the second scenario, she knows that there will be a second round after A is disclosed. In light of this knowledge, she expects to obtain the sum of the Firm's two payments:

$$(p_A + p_B)v_A + \mathbb{P}[B \mid A \text{ or } B](v_B - v_A),$$

which is clearly larger than the payment she receives in one shot, because the second term is not multiplied by the probability $p_A + p_B$. Therefore, an informed Agent gains from the information being sold slowly over time (whether she knows A or B). Given that the Firm is indifferent between both scenarios, and because in both scenarios all information gets disclosed, it must be that the uninformed Agent who knows nothing prefers the scenario in which information gets

disclosed in one shot.

This example illustrates why splitting information might be a good idea from the point of view of more informed agents. Rewarding such agents might be desirable from an *ex ante* perspective, if, for instance, agents have to be given incentives to acquire information.

Note, however, that the Agent who knows B does not extract the full value v_B of her information, even if she discloses this information progressively. Is this an artefact of the imperfect divisibility of information in this specific example, or a robust and inevitable feature of the sale of information without commitment?

To answer this question, and understand how information should be optimally disclosed when information can be divided more finely, we now turn to a more elaborate set-up, which will constitute our leading example throughout the paper. We will refine the informational structure, but simplify the payoff structure. We will then discuss how the example's ingredients affect the results (see Subsection 4.4).

3 The Model

There are two risk-neutral players: an Agent and a Firm. There are two states of the world, $\omega \in \Omega := \{0, 1\}$. The Agent is privately informed of the state of the world at the beginning of the game, but the Firm is not. The Firm's initial belief is that the state is 1 is p_0 , which is common knowledge. The fact that the Agent is perfectly informed is a normalization. See below for how to adjust the analysis in case that she is not.

The game lasts K rounds, but our focus will be on what happens as K grows arbitrarily large. After the K rounds have elapsed, the Firm must act in one of two ways, I or N . Either the Firm chooses to “Invest” in some venture (I) or to “Not Invest” (N). Not investing yields a safe (i.e., state-independent) payoff normalized to 0. Investing yields a payoff 1 when the state of the world is $\omega = 1$ and $-\gamma$ if the state of the world is $\omega = 0$ (for some $\gamma > 0$). That is, “Investing” is risky: it can pay more than the safe action, but only in one state.¹ The parameter γ measures the cost of taking this action, if it is inappropriate.

¹Given that we normalize to 1 the payoff of the investment in state 1, comparative statics with respect to γ have more than one interpretation.

Because the Agent knows the state, we call her a type-1 Agent, if $\omega = 1$, and a type-0 Agent otherwise. Indeed, from the point of view of payoffs, these are the only relevant events. However, as we shall see, the Agent might have a greater or lesser amount of information about the state of the world; hence, from a Bayesian viewpoint, the designation of these events as “types” is a convenient abuse of terminology.

Note that, if no information has been revealed, the Firm’s optimal action is to invest if and only if

$$p \geq p^* := \frac{\gamma}{1 + \gamma},$$

and obtain thereby a payoff of

$$w(p) := (p - (1 - p)\gamma)^+,$$

where $x^+ := \max\{0, x\}$. While our analysis will cover both the case in which the prior belief p_0 is below or above p^* , we will implicitly assume the more interesting case in which p_0 is smaller than p^* , unless stated otherwise. We will refer to the payoff $w(p)$ as the Firm’s *outside option*, and we will generalize the analysis to outside options with rather arbitrary specifications in Subsection 4.4.

In each of the K rounds before the action is taken, the Firm and Agent can make a monetary transfer, and the Agent can reveal some information if she so wishes. More precisely, the strategy has two parts. In round $k = 1, \dots, K$, as a function of the history of transfers and information disclosures up to that point, the Agent and the Firm can simultaneously make a non-negative transfer t_k^A and t_k^F , respectively, to the other party.² Second, once these transfers are made and observed, the Agent may disclose some verifiable information.

The information available to the Agent is type-dependent (thus, it is not “cheap talk”). The Agent might be one of many “informational” types. Possible evidence (that the state is 1) that the Agent might disclose includes some pieces of information that the type-0 Agent is very likely to possess, some that she is very unlikely to possess, and everything in between. That is, it spans the entire range of probabilities with which the type-0 Agent may possess this piece of

²The Reader might wonder what is gained by allowing the Agent to pay the Firm. After all, it is the Agent who possesses the valuable good in this model, which is the information. Indeed, as we shall see, such payments are not essential in the equilibria of the baseline model. However, such payments will play a critical role once we allow for an intermediary.

information. The type-1 Agent, on the other hand, possesses all these pieces of information, and, in addition, definitive evidence that she is a type-1 Agent.

This implies that, given any belief $p \in (0, 1)$ that the Firm might assign to state 1, following some arbitrary history, and for any $p' \geq p$, there exists some piece of information whose disclosure would lead the Firm to update his belief to p' . That is, for any pair of beliefs $\{p, p'\}$, there exists a piece of information $\iota(p, p') \in \mathcal{I}$ that the type-1 Agent possesses for sure, but that the type-0 Agent possesses only with probability

$$q = \frac{1 - p'}{p'} \frac{p}{1 - p}.$$

By construction, if the Agent is expected to disclose this piece of information whenever she can, it follows from Bayes' rule that the posterior probability assigned to $\omega = 1$ is equal to

$$\frac{p}{p + (1 - p)q} = p'.$$

If the Agent fails to disclose this piece of information, while she is expected to do so if she can, the Firm correctly updates his belief to probability zero. Note that, considering the case $p' = 1$, there is some piece of information that the type-0 Agent does not possess, that is, a “proof” that the state is 1.³

The specific nature of this information and of the set \mathcal{I} need not concern us here (some interpretations and additional formal definitions are provided below). What matters here is that this information is perfectly *divisible*, that is, there exists such a piece of information for each $p' \geq p$, and *uniquely defined* in equilibrium, so that, if there are two pieces of information that the type-1 Agent possesses for sure and that the type-0 Agent possesses with probability q for each piece, the equilibrium specifies which piece the Agent is meant to reveal. (Otherwise, recognizing that the type-0 Agent had an incentive to reveal whichever of the two pieces of information she

³A modeling issue arises at $p = 0$. What if the Firm, after some history of transfers and disclosures, assigns probability 0 to $\omega = 1$, but the Agent then discloses this proof? Fortunately, our purpose is to identify the best equilibrium, not to characterize the set of all equilibria; hence, this issue is irrelevant for our purpose. The equilibrium that we describe remains an equilibrium if it is required that players cannot switch away from probability 1 beliefs, and remains the best equilibrium if this requirement (not imposed by perfect Bayesian equilibrium) is dropped.

actually possesses, if any, the Firm would not update his belief to p' after all.) This allows us to conduct the analysis entirely in terms of beliefs, and to abstract from issues that pertain to the type of information that is being disclosed, which leaves open some interesting questions (for instance, in which order should information be released?)

The Agent does not care about the Firm's decision *per se*. All she seeks to do is to maximize the sum of the net transfers she receives during the K rounds. The Firm seeks to maximize the payoff from his decision after the K rounds, net of the payments that he made. There is neither discounting, nor any other type of friction during the K rounds. In particular, there is no cost to disclosing information.

More formally, a history of length k is a sequence

$$h_k = \{(t_{k'}^A, t_{k'}^F, \iota_{k'})\}_{k'=0}^{k-1},$$

where $(t_{k'}^A, t_{k'}^F, \iota_{k'}) \in \mathbb{R}_+^2 \times \mathcal{I}$. The set of all such histories is denoted by H_k (set $H_0 := \emptyset$). Given some final history h_K (which does not include the Firm's final action), the Agent's payoff is simply the sum of all net transfers over all rounds:

$$V(h_K) = \sum_{k=0}^{K-1} (\tau_k^F - \tau_k^A).$$

Given state ω , the Firm's overall payoff results from his action, as well as from the sum of net transfers. If the Firm chooses the safe action, he gets

$$W(\omega, h_K, N) = \sum_{k=0}^{K-1} (\tau_k^A - \tau_k^F).$$

If instead the Firm decides to invest, he receives

$$W(\omega, h_K, I) = \sum_{k=0}^{K-1} (\tau_k^A - \tau_k^F) + 1 \cdot \mathbf{1}_{\omega=1} - \gamma \cdot \mathbf{1}_{\omega=0},$$

where $\mathbf{1}_A$ denotes the indicator function of the event A .

A prior belief p_0 (that the state is 1) and a strategy profile $\sigma := (\sigma^F, \sigma^A)$ define a distribution over (ω, h_K, a) , and we let $V(\sigma), W(\sigma)$, or simply (V, W) , denote the expected payoffs of the Agent and the Firm, respectively, with respect to this distribution. When the strategy profile is understood, we also write $V(h_k), W(h_k)$ for the players' continuation payoffs, given history h_k . We further write $V_0(\sigma)$ or $V_1(\sigma)$, or simply (V_0, V_1) , for the payoff of the Agent, when these payoffs are conditional on the state $\omega = 0, 1$. That is, $V_1(\sigma)$ is the payoff of type-1 Agent, when the strategy profile is σ .

Our focus is to identify the (perfect Bayesian) equilibria σ that maximize the payoff of the type-1 Agent. More precisely, we are interested in the limit of this equilibrium payoff as the number of rounds K becomes arbitrarily large.⁴ We also discuss other equilibria (see Lemma 1).

One possible motivation for focusing on this equilibrium is that the probability of the state 1 is itself the result of some effort by the Agent. For instance, it could be that the action consists in buying or not an asset at a given price and the two states correspond to the Agent having noisy signals about the value of the asset to the Firm. It may be socially desirable to reward those of the Agent's activities that increase the accuracy of her signals. We embed our model in such a framework in Section 5. In such a framework, this equilibrium is best in an *ex ante* sense, which might help players to coordinate on it.

As noted at the beginning of this section, that the Agent's information about the state is perfect is irrelevant for our analysis as long as it is optimal for the Firm to invest in the information if she is of type 1 and optimal for him to not invest if she is of type 0. What matters then is the Firm's belief about the Agent's type. His outside option, as a function of this belief, has the same features as in our simpler model: for low enough beliefs (that the Agent is of type 1), it is constant (at a level that can be normalized to zero), and it is increasing linearly for stronger beliefs.

We conclude this subsection with a more formal description of strategies and equilibrium.

A behavior strategy σ^F for the Firm is a collection $(\{\tau_k^F\}_{k=0}^{K-1}, \alpha^F)$, where τ_k^F is a probability transition $\tau_k^F := H_k \rightarrow \mathbb{R}_+$, specifying a transfer $t_k^F := \tau^F(h_k)$ as a function of the history so far,

⁴The equilibrium that we will obtain is also an equilibrium of the infinite-horizon, undiscounted game, but using limits allows us to uniquely pin down the limiting strategy profile.

as well as an action (a probability transition as well), $\alpha^F : H_K \rightarrow \{I, N\}$ after the K -th round.⁵ A behavior strategy σ^A for the Agent is a collection $\{\tau_k^A, \iota_k^A\}_{k=0}^{K-1}$, where $\tau_k^A := \mathcal{P}(\mathcal{I}) \times H_k \rightarrow \mathbb{R}_+$ is a probability transition specifying the transfer $t_k^A := \tau^A(h_k)$ in round k given the history so far and given the information she has, and $\iota_k^A := \mathcal{P}(\mathcal{I}) \times H_k \times \mathbb{R}_+^2 \rightarrow \mathcal{I}$ is a probability transition specifying the information that is released in round k , as a function of the state, the history up to the current round, and the transfers that were made in the round.⁶ As the Agent can only reveal what she knows, we impose $\iota_k^A(\mathcal{I}_0, h_k, t_k^F, t_k^A) \in \mathcal{I}_0$ for all (h_k, t_k^F, t_k^A) and all \mathcal{I}_0 (that is, message sets are type-dependent). Further, the possible sets of information are assumed to be sufficiently rich that, from every prior belief p that the state is 1, there exists a piece of information whose disclosure would lead to a posterior belief p' , for all $p' \in [p, 1]$.⁷

The solution concept is perfect Bayesian equilibrium, as defined in Fudenberg and Tirole (1991, Definition 8.2).⁸

4 Equilibrium Analysis

4.1 Preliminaries

For the sake of simplicity, we restrict attention here to pure strategies. However, we prove all the results in the Appendix, without imposing this restriction.

A pure strategy calls for the Agent to disclose a specific piece of information at each round (revealing nothing being a special case). Of course, this is only possible if the Agent possesses

⁵Formally, \mathcal{I} is some topological space, $\tau^F(h_k)$ is a probability distribution on \mathbb{R}_+ , and for each Borel set $A \in \mathbb{R}_+$, $\tau^F(\cdot)[A]$ is a measurable function of h_k .

⁶The Agent's information consists of a subset \mathcal{I}_0 of \mathcal{I} (i.e., an element of $\mathcal{P}(\mathcal{I})$, such that the set of all subsets of \mathcal{I}) represents all the pieces of information she has. The type-1 agent has all the information, i.e. $\mathcal{I}_0 = \mathcal{I}$.

⁷Formally, the description of the game includes a Borel measure ν on $\mathcal{P}(\mathcal{I})$, representing the prior distribution over the sets of pieces of information owned by the type-0 Agent, with the property that, for every sequence $\{\iota_1, \dots, \iota_k\}$ such that $\nu(\mathcal{I}_0 : \{\iota_1, \dots, \iota_k\} \in \mathcal{I}_0) = (1-p)p_0/(p(1-p_0))$, and for every $p' \in [p, 1]$, there exists $\iota_{k+1} \in \mathcal{I}$ such that $\nu(\mathcal{I}_0 : \{\iota_1, \dots, \iota_k, \iota_{k+1}\} \in \mathcal{I}_0) = (1-p')p_0/(p'(1-p_0))$. For example, we can pick $\mathcal{I} = [0, 1]^{\mathbb{N}}$, along with the Lebesgue measure on each unit interval. Note that this is formally a game of incomplete information with more than two types (in the sense of Harsanyi, 1967/68), because there are as many types as possible subsets of pieces of information that the Agent might hold.

⁸Fudenberg and Tirole define perfect Bayesian equilibria for finite multistage games with observed actions only. Here in contrast, both the type space and the action sets are infinite. The natural generalization of their definition is straightforward and therefore not given here.




Figure 1: A feasible action

this piece of information. She does if she is of type 1, but she might not if she is of type 0.

This disclosure implies that, from the Firm's point of view, and ignoring the uninteresting case in which the Agent is assumed to reveal nothing, the posterior probability that he assigns to state 1 will take one of two values: it will either (i) jump from p_0 up to some p' , if the piece of information $\iota(p_0, p')$ is indeed revealed, or (ii) fall to zero. This is illustrated in Figure 1 below. The two arrows indicate the two possible posterior probabilities that the Firm assigns to state 1. Note that, as a stochastic process, and viewed from the Firm's perspective, this belief must follow a martingale: the Firm's expectation of his posterior belief must be equal to his prior belief. However, this is not the case from the Agent's point of view. Given her knowledge of the state, she assigns different probabilities to the event that the Firm's posterior belief takes each of these two values than the Firm. If she is a type-1 Agent, she knows for sure that the probability will not decrease over time, because she possesses all pieces of information (the Firm's belief process is then a submartingale, relative to her information). Conditional on state 0, the posterior probability is below p_0 (the process is then a supermartingale). However, depending upon whether or not the type-0 Agent happens to possess the specific piece of information whose disclosure the equilibrium calls for, she knows for sure whether the posterior probability will be p' or 0.

Note that the joint payoffs to the players cannot exceed the surplus in the game. This is a feasibility restriction. If the probability of state 1 is p , given the history h_k , then the expected surplus (assuming that the Firm takes an optimal eventual decision) is $p \cdot 1 + (1 - p) \cdot 0 = p$.

This means that the continuation payoffs must satisfy

$$pV_1(h_k) + (1 - p)V_0(h_k) + W(h_k) \leq p. \quad (1)$$

The equilibrium is *efficient* if this constraint is binding, that is, if the type-1 Agent discloses all her information eventually, on the equilibrium path.

In fact, if there is an equilibrium with payoffs (V_0, V_1) to the Agent, there exists an efficient equilibrium with these payoffs, because the Agent can always disclose the state in the last period *on the equilibrium path*. This cannot weaken the incentives for the players to carry out the planned transfers (because it can only increase the payoff from following the specified equilibrium actions), but it guarantees that the correct action is taken eventually. That being so, we may focus on efficient equilibria.

There are further constraints on equilibrium payoffs, for example, those imposed by individual rationality. From any history onward, the Agent can secure a payoff of zero, independently of her type:

$$V_1(h_k) \geq 0, V_0(h_k) \geq 0.$$

The Firm, on the other hand, can secure a higher continuation payoff. If he receives no further information, he receives his outside option

$$w(p) = (p - \gamma(1 - p))^+. \quad (2)$$

Given that additional information cannot hurt the Firm, this is a lower bound on $W(h_k)$.

This game admits many equilibria. For instance, there is an equilibrium in which no transfers are ever made and no information is ever released. This equilibrium achieves the lower bounds on the players' payoffs, thereby providing a useful threat from any history onward, but it is clearly inefficient.

There also exists an efficient equilibrium in which no transfers are ever made, and the type-1 Agent reveals the state in the last period, so that the posterior probability that she assigns to state 1 is 1 with probability p , and 0 otherwise. This yields a payoff of zero to the Agent, and a payoff of p to the Firm, which, given feasibility and individual rationality, is an upper bound to

his payoff.

From the Agent's point of view, we already know that her expected payoff $pV_1 + (1 - p)V_0$ cannot exceed $p - w(p)$. Is there an equilibrium in which this payoff is achieved? One round is enough to establish this: the Firm pays this amount upfront, and the Agent reveals the state. If the Firm fails to pay, no information is disclosed. Note that the Firm is indifferent to paying or not paying, given the punishment for failing to do so, and the Agent is willing to release this information.

Given some equilibrium yielding payoffs (V_0, V_1, W) in the game with K rounds, we also claim that

$$V_1 + W \leq \bar{V}_1 + w(p),$$

where \bar{V}_1 denotes the highest equilibrium payoff for the type-1 Agent. Otherwise, by simply starting from the equilibrium that yields V_1 to the type-1 Agent and W to the Firm, and by increasing the initial transfer that the Firm is asked to make by an amount $W - w(p)$, we would obtain another equilibrium in which the type-1 Agent gets a payoff strictly above \bar{V}_1 , which is a contradiction. Therefore, the equilibrium that maximizes the type-1 Agent's payoff cannot leave any surplus to the Firm, and it also maximizes the sum of the Firm's and type-1 Agent's payoffs.

Last, because any efficient equilibrium must satisfy

$$pV_1(h_k) + (1 - p)V_0(h_k) + W(h_k) = p, \tag{3}$$

we claim that any equilibrium that maximizes V_1 also maximizes $V_1 - V_0$. To see this, note that we can assume without loss of generality that the former is efficient (by specifying full disclosure in the last round), so that, dropping arguments,

$$V_0 = \frac{p(1 - V_1) - W}{1 - p}, \text{ and so } V_1 - V_0 = \frac{V_1 + W - p}{1 - p}.$$

Therefore, maximizing the payoff difference $V_1 - V_0$ is equivalent to maximizing the sum $V_1 + W$. However, as we have already remarked, this is in turn equivalent to maximizing V_1 only.

We summarize these observations in the following Lemma.

Lemma 1 *Given any history h_k , $k \leq K - 1$, the following holds for the continuation game:*

1. *There exists an equilibrium in which both players are held to their minmax payoffs:*

$$W = w(p), V_1 = V_0 = 0.$$

2. *There exists an equilibrium in which the Firm gets all the surplus:*

$$W = p, V_0 = V_1 = 0.$$

3. *There exists an equilibrium in which the Agent receives all the surplus, net of the Firm's outside option:*

$$V = p - w(p), W = w(p).$$

4. *If the Agent receives V as a continuation equilibrium payoff, there exists an equilibrium in which the Firm receives all the residual surplus:*

$$W = p - V.$$

5. *Any equilibrium that maximizes V_1 , the type-1 Agent's payoff, also maximizes $V_1 + W$, the sum of the Firm's and type-1 Agent's payoffs, as well as $V_1 - V_0$, the difference between the two Agent's types' payoffs.*

6. *The set of equilibrium payoffs is non-decreasing in K , the number of rounds.*

The last conclusion is an immediate consequence of the fact that players can always choose not to make transfers, nor to disclose any information, in the first round. Note also that, due to the fact that the type-1 Agent can always mimic the type-0 Agent, her payoff must be at least as high as the type-0's payoff. This implies that the maximal equilibrium payoff for the type-0's Agent is the one that maximizes the Agent's *ex ante* payoff, as described above.

This still leaves open our main concern: what is the highest equilibrium payoff \bar{V}_1 of the type-1 Agent?

4.2 Benefits of Splitting Information

We now turn our focus to the equilibrium that maximizes the payoff of the type-1 Agent. We already know that it is possible for that Agent to appropriate some of the value of her information, but the question is whether she can get more than $p - w(p)$, which is just as much as the type-0 Agent gets in the equilibrium we constructed so far. We show that the answer is negative; *unless* the Agent can reveal the information slowly:

Proposition 1 *If $K = 1$, the highest equilibrium payoff to the type-1 Agent is equal to $p_0 - w(p_0)$. Moreover, in any equilibrium, both Agent's types must receive the same payoff: $V_0 = V_1$.*

Proof: With one round of communication, the payoff of the Agent can come only from the payment in the first (and only) round. Therefore, the payoffs of the two types of Agent must be the same.

If, in equilibrium, the set of possible posterior probabilities that the Firm assigns to state 1 is $\{0, p_1\}$ (recall that mixed strategies are dealt with in the Appendix), then the most the Firm is willing to pay is

$$\mathbb{E}_F [w(p')] - w(p_0),$$

where the subscript to the expectation refers to the fact that this is relative to the Firm's belief, and p' is the the Firm's posterior belief—a random variable with values in $\{0, p_1\}$. Because beliefs must follow a martingale from the Firm's point of view, it must be that the probability that the posterior probability that the Firm assigns to state 1 is p_1 is p_0/p_1 , because

$$p_0 = \frac{p_0}{p_1} \cdot p_1 + \frac{p_1 - p_0}{p_1} \cdot 0.$$

This means that the additional value from this information, relative to what the Firm can secure, is

$$\frac{p_0}{p_1} w(p_1) - w(p_0).$$

This is increasing in p_1 and so maximized at $p_1 = 1$, yielding a payment $p_0 \underbrace{w(1)}_{=1} - w(p_0)$. As we have already observed, this payoff can be supported as an equilibrium payoff. \square




Figure 2: Revealing information in two steps

The intuition is simple: with one round to go, the highest payment is achieved when the Firm's next outside option, i.e., the payoff from the action it takes, is maximized, and this requires full information disclosure on the equilibrium path. Note that, when $p_0 \leq p^*$, the highest payoff in one round that the type-1 Agent can get in equilibrium is simply the prior probability p_0 .

It turns out that, with as few as two rounds, the type-1 Agent can earn more. The key is to reveal information gradually. Figure 2 represents one possible strategy for the disclosure of information. In the first step in the first round, the Agent discloses the piece of information that leads to the Firm holding a posterior belief of p^* (or 0, if she fails to do so), and information is then fully disclosed in the second step in the first round. No payment is made in the first round. In the second round, the equilibrium of the one-round game is played. What is the resulting payoff to the type-1 agent? Note that the payoff of the two-round game, obtained in the second round, is now

$$p^* - w(p^*) = p^* > p_0.$$

This argument is summarized in the following proposition.

Proposition 2 *Suppose that $p_0 < p^*$ and $K \geq 2$. Then the highest equilibrium payoff of the type-1 Agent is at least p^* .*

This illustrates the benefits of splitting information and revealing it slowly over time. Is the splitting that we described optimal with two rounds to go? As it turns out, for $p_0 < (p^*)^2$, it is. Yet there are many other ways of splitting information with two periods to go that improve

upon the one-round equilibrium, and among them, splits that also improve over the one-period equilibrium when $p_0 > p^*$.

Allowing additional rounds will further improve what a type-1 Agent can achieve. This can be understood by considering a graphical representation. Consider Figure 3. As shown on the left panel, information is revealed here in three steps. First, the belief is split into 0 and p^* . Second, at p^* (assuming this belief is reached), it is split in 0 and p' . Finally, at p' , it is split in 0 and 1. The right panel shows how to determine the type-1 Agent's payoff graphically. The two solid (red) segments represent the maximal payments that the type-1 Agent can demand at each round for the information that is being released in the second and third rounds. (In the first round, no payment can be demanded, because if future payments drive down the Firm's continuation payoff from the second round onward to his outside option, his continuation payoff is zero whether his posterior belief goes up or down). Thus, the sum of the lengths of the solid red segments is the payoff of the type-1 Agent. In contrast, in the equilibrium involving two rounds only, in which information is fully disclosed once the belief reaches p^* , the payment to the Agent is only equal to the distance of the vertical segment between the outside option w at p^* and the chord connecting $(0, 0)$ and $(1, 1)$ evaluated at p^* (i.e., the lower segment, plus the dotted segment). It is clear that the profit with three rounds exceeds the profit with only two, because the chords from the origin to the point $(p, w(p))$ become steeper as p increases.

4.3 The highest payoff for the type-1 Agent

As we saw, the type-1 Agent may earn a higher payoff than the type-0 Agent if it is possible to disclose information gradually. We now solve for the highest equilibrium payoff for the type-1 Agent, for every number of rounds K and every prior belief p_0 .

The best equilibrium involves a pure strategy by the Agent: the equilibrium is characterized by a sequence of pieces of information that she releases one by one, if she can. That is, the type-1 Agent discloses these pieces sequentially and reveals the state in the last period by disclosing all remaining pieces of information (which the type-0 Agent cannot do). If the Agent is a type-0 Agent, she discloses these pieces sequentially as well, until she can do so no longer. When this situation arises, the Firm updates the probability that he assigns to state $\omega = 1$ to zero, because




Figure 3: Revealing information in three steps: evolution (left) and payoff (right)

she fails to disclose the requisite piece of information.

This equilibrium, then, can be summarized by a sequence of posterior beliefs, representing the Firm's belief after each round that the state is 1, conditional on all pieces of information having been exhibited up to that round. This sequence $\{p_0, \dots, p_K\}$ starts at the Firm's prior belief, p_0 , and ends up at $p_K = 1$.

Of course, an equilibrium must also specify transfers, as well as how players behave off the equilibrium path. The equilibrium is such that, from every round onward, and for every history up to this round, the Firm is held to his outside option. Therefore, if the Firm's belief in the next round is either p_{k+1} , or 0, given the current belief p_k , he is willing to pay

$$\mathbb{E}_F[w(p')] - w(p_k),$$

where p' is the random belief in the next round, with possible values 0 and p_{k+1} . The Agent does not make any transfers. In other words, the Agent extracts the maximum payment she can hope for from the Firm at every round. This sounds intuitive, but, as we shall see, this will no longer be optimal when a slightly more general class of mechanism is considered.

If either player deviates from the specified course of actions, play reverts to the worst equilib-




Figure 4: Revealing information in many steps (left); foregone profit at each step (right)

rium, in which no further transfer is made and no information is ever disclosed again (the Firm's beliefs are still updated according to Bayes' rule, whenever possible, so they remain constant after the deviation).

It is intuitively clear from the analysis of the previous subsections that splitting the disclosure of information over many rounds is beneficial. The left panel of Figure 4 illustrates the total payoff that results from a splitting that involves many small steps (which is the sum of all vertical segments). The Reader might be tempted to surmise that, in the limit as $K \rightarrow \infty$, the type-1 Agent will be able to extract the full value of the information. The right panel explains why this conjecture is incorrect. As the Firm's belief goes from $p - dp$ to p , his outside option increases from $w(p - dp)$ to $w(p)$, yet the type-1 Agent only charges a fraction of this, giving up $w(p)dp/p$ in this process. This loss, or foregone profit, need not be large when the step size dp is small, but then again, the smaller the step size, the larger the number of steps that the disclosure policy involves. As a result, the type-1 Agent cannot avoid giving up a fraction of the value of the information. Note that this sacrificed profit does not benefit the Firm, who is always charged his full willingness to pay. Therefore, it benefits the type-0 Agent, whose profit does not tend to zero, even as the number of rounds goes to infinity.

It turns out that the optimal strategy can be computed explicitly for every K . Given that the

equilibrium strategies depend only on the number of remaining rounds and on the Firm's belief, we write $V_{1,K}(p_0)$ for the payoff of the type-1 Agent, given that there are K rounds remaining and the Firm's belief is p_0 (assuming that this history is on the equilibrium path).

The following proposition describes the optimal sequence of beliefs $\{p_0, \dots, p_K\}$, and the payoff to the type-1 Agent, as a function of the number of rounds and the prior belief p_0 . Here, $(x)^- := -\min\{0, x\} \geq 0$.

Proposition 3 *The maximal equilibrium payoff of the type-1 Agent with K rounds, given the Firm's prior belief p_0 , is given recursively by*

$$V_{1,K}(p_0) = \begin{cases} K\gamma(1 - p_0^{1/K}) - (p_0 - \gamma(1 - p_0))^- & \text{if } p_0 \geq (p^*)^{\frac{K}{K-1}}, \\ V_{1,K-1}(p^*) & \text{if } p_0 < (p^*)^{\frac{K}{K-1}}, \end{cases}$$

for $K > 1$, with $V_{1,1}(p_0) = \gamma(1 - p_0) - (p_0 - \gamma(1 - p_0))^-$. On the equilibrium path, in the initial round, the type-1 Agent reveals a piece of information that leads to the Firm assigning a posterior probability to state 1 of

$$p_1 = \begin{cases} p_0^{\frac{K-1}{K}} & \text{if } p_0 \geq (p^*)^{\frac{K}{K-1}}, \\ p^* & \text{if } p_0 < (p^*)^{\frac{K}{K-1}}, \end{cases}$$

after which the play proceeds as in the best equilibrium with $K - 1$ rounds, given prior p_1 .

Proof: The proof is by induction on the number of rounds. The argument presented here is for pure strategies. See Appendix A for the proof that there is no mixed-strategy equilibrium with a greater payoff to the type-1 Agent.

Our induction hypothesis is that, with $k \geq 1$ periods to go, and a prior belief $p = p_0$, the best equilibrium involves setting the next (non-zero) posterior belief, p_1 , equal to $p_1 = p^{\frac{k-1}{k}}$ if $p^{\frac{k-1}{k}} \geq p^*$ (i.e. if $p \geq (p^*)^{\frac{k}{k-1}}$ for $k \geq 2$), and equal to p^* otherwise.⁹ Further, the type-1 Agent's maximal payoff with k rounds to go is equal to

$$V_{1,k}(p) = k\gamma(1 - p^{1/k}) - (p - \gamma(1 - p))^- \text{ if } p \geq (p^*)^{\frac{k}{k-1}}, \text{ and } V_{1,k}(p) = V_{1,k-1}(p^*) \text{ if } p < (p^*)^{\frac{k}{k-1}}.$$

⁹In this proof, when we say that the equilibrium involves setting the posterior belief p_1 , we mean that, from the type-1 Agent's point of view, the posterior belief will be p_1 , while from the point of view of the Firm, the posterior belief will be a random variable p' with possible values $\{0, p_1\}$.

Note that this claim implies that

$$V_{1,k}(p^*) = k\gamma(1 - (p^*)^{1/k}).$$

Finally, as part of our induction hypothesis, we claim the following. Given some equilibrium, let $X \geq 0$ denote the payoff of the Firm, net of its outside option, with k rounds left. That is, $X := W_k(p) - w(p)$, where $W_k(p)$ is the Firm's payoff given the history leading to the equilibrium belief p with k rounds to go. Let $V_{1,k}(p, X)$ be the maximal payoff of the type-1 Agent over all such equilibria, with associated belief p , and excess payoff X promised to the Firm (set $V_{1,k}(p, X) := -\infty$ if no such equilibrium exists). Then we claim that

$$V_{1,k}(p, X) \leq V_{1,k}(p) + X.$$

We first verify this with one round. As we saw in Proposition 1, if $K = 1$, it is optimal to set the posterior p_1 equal to 1, which is $p^{\frac{K-1}{K}}$, the relevant specification given that $p_1^0 = 1 \geq p^*$. The payoff to the type-1 Agent is

$$V_{1,1}(p) = p - (p - \gamma(1 - p))^+ = \gamma(1 - p) - (p - \gamma(1 - p))^-,$$

as was to be shown. Note that this equilibrium is efficient. This implies that $V_{1,1}(p, X) \leq V_{1,1}(p) + X$, for all $X \geq 0$, because any additional payoff to the Firm must come as a reduction of the net transfer from the Firm to the Agent.

Assume that this holds with k rounds to go, and consider the problem with $k + 1$ rounds. Of course, we do not know (yet) whether, in the continuation game, the Firm will be held to its outside option.

Note that the Firm assigns probability p/p_1 to the event that its posterior belief p' will be p_1 , because, by the martingale property, we have

$$p = \mathbb{E}_F[p'] = \frac{p}{p_1} \cdot p_1 + \frac{p_1 - p}{p_1} \cdot 0.$$

This implies that, with $k + 1$ rounds, the Firm is willing to pay at most

$$\bar{t}_{k+1}^F := \frac{p}{p_1} (w(p_1) + X') - w(p),$$

where X' is the excess payoff of the Firm with k rounds to go, given posterior belief p_1 . Therefore, the payoff to the type-1 Agent is at most

$$V_{1,k+1}(p) \leq \bar{t}_{k+1}^F + V_{1,k}(p_1; X') \leq \frac{p}{p_1} (w(p_1) + X') - w(p) + V_{1,k}(p_1) - X',$$

where the second inequality follows from our induction hypothesis. Note that, since $p/p_1 < 1$, this is a decreasing function of X' : it is best to hold the Firm to his outside option when the next round begins. Therefore, we maximize

$$\frac{p}{p_1} w(p_1) + V_{1,k}(p_1).$$

Note first that, given the induction hypothesis, all values $p_1 \in [p, (p^*)^{\frac{k}{k-1}}]$ yield the same payoff, because for any such p_1 , $V_{1,k}(p_1) = V_{1,k-1}(p^*)$. The remaining analysis is now a simple matter of algebra. Note that, for $p_1 \in [(p^*)^{\frac{k}{k-1}}, p^*]$ (which obviously requires $p < p^*$), the objective becomes (using the induction hypothesis)

$$V_{1,k}(p_1) = k\gamma(1 - (p_1)^{1/k}) + (p_1 - \gamma(1 - p_1)),$$

which is increasing in p_1 , so that the only candidate value for p_1 in this interval is $p_1 = p^*$. Consider now picking $p_1 \geq p^*$. Then we maximize

$$\frac{p}{p_1} (p_1 - \gamma(1 - p_1)) + k\gamma(1 - p_1^{1/k}),$$

which admits a unique critical point $p_1 = p^{\frac{k}{k+1}}$, achieving a payoff equal to $(k+1)\gamma(1 - p^{1/(k+1)}) + p - \gamma(1 - p) = (k+1)\gamma(1 - p^{1/(k+1)})$. Note, however, that this critical point satisfies $p_1 \geq p^*$ if and only if $p \geq (p^*)^{\frac{k+1}{k}}$.

Therefore, the unique candidates for p_1 are $\{p^*, \max\{p^*, p^{\frac{k}{k+1}}\}, 1\}$. Observe that setting the posterior belief p_1 equal to $\max\{p^*, p^{\frac{k}{k+1}}\}$ does at least as well as choosing either p^* or 1. This

establishes the optimality of the strategy, and the optimal payoff for the type-1 Agent, with $k + 1$ rounds to go.

Finally, we must verify that $V_{1,k+1}(p; X) \leq V_{1,k+1}(p) - X$. Given that we have observed that it is optimal to set $X' = 0$ in any case, any excess payoff to the Firm with $k + 1$ rounds to go is best obtained by a commensurate reduction in the net transfer from the Firm to the Agent in the first round (among the $k + 1$ rounds). This might violate individual rationality for some type of the Agent, but even if it does not, it still yields a payoff $V_{1,k+1}(p; X)$ no larger than $V_{1,k+1}(p) - X$ (if it does violate individual rationality, $V_{1,k+1}(p; X)$ must be lower). \square

Note that, fixing $p_0 < p^*$, $p_0 < (p^*)^{\frac{k}{k-1}}$ for all k large enough, so that, with enough rounds ahead, it is optimal to set $p_1 = p^*$ in the first round, then to follow the sequence of posterior beliefs $(p^*)^{\frac{k-1}{k}}, (p^*)^{\frac{k-2}{k}}, \dots, 1$. The payoff to the type-1 Agent from doing so tends to

$$\lim_{k \rightarrow \infty} V_{1,k}(p^*) = -\gamma \ln p^*, \quad (4)$$

and the sequence of posterior beliefs that are used successively becomes dense in $[p^*, 1]$. Therefore, with sufficiently many rounds to allow progressive disclosure, the equilibrium involves progressive disclosure of information, with a first big step leading to the posterior belief p^* , given the prior belief $p_0 < p^*$, followed by a succession of very small disclosures, leading the probability that the Firm assigns gradually up all the way to 1. This is illustrated in Figure 5 below.

Here is an alternative, heuristic derivation of the formula in (4). Note that, for $p \geq p^*$, the payment that the type-1 Agent can extract from the Firm if the following posterior belief is $p' \in \{0, p + dp\}$ is

$$\begin{aligned} \frac{p}{p + dp} w(p + dp) - w(p) &= \\ \frac{p}{p + dp} ((p + dp) - \gamma(1 - p - dp)) - (p - \gamma(1 - p)) &= \gamma \frac{dp}{p} + O(dp^2). \end{aligned}$$

If the entire interval $[p^*, 1]$ is divided in this fashion, the resulting payoff then tends to

$$\int_{p^*}^1 \gamma \frac{dp}{p} = \gamma(\ln 1 - \ln p^*) = -\gamma \ln p^*,$$




Figure 5: Revealing information in many steps

which is the value that we found. This also illustrates that the limiting payoff is independent of the precise way in which the information is divided up over time, as long as the mesh of the partition tends to zero. The following corollary records the limiting value.

Corollary 1 *As $K \rightarrow \infty$, the optimal payoff tends to, for $p_0 < p^*$,*

$$V_{1,\infty}(p_0) = -\gamma \ln p^*.$$

Note that this payoff is independent of p_0 (for $p_0 < p^*$), and the first piece of information, which leads to a posterior belief of p^* , is given away for free, because it does not affect the Firm's outside option. All later, very small releases of information do affect this outside option, and are priced accordingly.

Figure 6 illustrates how the payoff to the type-1 Agent varies with the number of rounds and the Firm's prior belief.

4.4 More General Payoffs

How do our results depend on our assumptions about the outside option? While the piecewise linear structure of the Firm's payoff proves quite convenient for explicit formulas, the main results of Subsection 4.3 generalize to more general specifications.

Suppose that the payoff of the Firm (gross of any transfers) as a function of his posterior




Figure 6: Value functions given the number of rounds left ($\gamma = 1$)

belief p after the K rounds is a non-decreasing continuous function $w(p)$, and normalize $w(0) = 0$. This payoff can be thought as the reduced form of some decision problem that the Firm faces, as in our baseline model. In that case, w must be convex, but due to the fact that we are taking w as a primitive here, we do not assume such a property here. Then it is again in the interest of the type-1 Agent to split information as finely as possible for any prior belief p_0 if and only if the function w is *star-shaped*, i.e., if and only if the average $w(p)/p$ is a strictly increasing function of p .¹⁰ If in a given round the Firm's belief goes from p to either $(p + dp)$ or 0, the Agent can charge up to

$$\frac{p}{p + dp} w(p + dp) - w(p) = (w'(p) - w(p)/p) dp + O(dp^2)$$

for it.¹¹ Given the Firm's prior belief p_0 , the type-1 Agent's payoff becomes then (in the limit,

¹⁰This condition, which is weaker than strict convexity, has already appeared in the economics literature in the study of risk (see Landsberger and Meilijson, 1990). We thank Arthur Robson for pointing this out.

¹¹In case $w(p)$ is not differentiable, then $w'(p)$ is the right-derivative, which is well-defined in case w is star-shaped.

as the number of rounds K goes to infinity)

$$\int_{p_0}^1 [w'(p) - w(p)/p] dp = w(1) - w(p_0) - \int_{p_0}^1 w(p) dp/p,$$

which generalizes the formula that we have seen for the special case $w(p) = (p - (1 - p)\gamma)^+$. That is, the type-1 Agent's payoff is the area between the marginal payoff of the Firm and its average payoff. To see that this splitting of information as finely as possible is optimal, consider some arbitrary interval of beliefs $[\underline{p}, \bar{p}]$. By having the posterior belief of the Firm jump from \underline{p} to \bar{p} , the payoff in that round is given by

$$\frac{\bar{p}}{\underline{p}} w(\bar{p}) - w(\underline{p}).$$

If instead this interval of beliefs is split as finely as is possible, the payoff over this range is

$$w(\bar{p}) - w(\underline{p}) - \int_{\underline{p}}^{\bar{p}} \frac{w(p)}{p} dp.$$

Hence, splitting is better if and only if

$$\frac{1}{\bar{p} - \underline{p}} \int_{\underline{p}}^{\bar{p}} \frac{w(p)}{p} dp \leq \frac{w(\bar{p})}{\bar{p}},$$

which is satisfied if the average $w(p)/p$ is increasing. Conversely, if the average is decreasing over some range $[\underline{p}, \bar{p}]$, this argument shows that it is better to have the belief jump from \underline{p} to \bar{p} than to split it as finely as possible.¹² If the average is decreasing over some range, what determines the jump? Note that, as mentioned, the payoff from a jump is $\underline{p}w(\bar{p})/\bar{p} - w(\underline{p})$, while the marginal benefit from splitting information disclosures finely at any given belief p (in particular, at \bar{p} and

¹²If the average is constant over some interval (as in our example over the range $[0, p^*]$), then in the limit it is irrelevant whether the belief jumps or not, but the limiting procedure we adopted will exclude any splitting over this range, because for any finite K , such a split would correspond to a “wasted” round.




Figure 7: Splitting information with an arbitrary outside option

\underline{p}) is $w'(\underline{p}) - w(\underline{p})/\underline{p}$. Setting the marginal benefits equal at \underline{p} and \bar{p} , respectively, yields that

$$\frac{w(\bar{p})}{\bar{p}} = \frac{w(\underline{p})}{\underline{p}} \text{ and } w'(\bar{p}) = \frac{w(\bar{p})}{\bar{p}}.$$

See Figure 7. The left panel illustrates how having two rounds improves on one round. Starting with a prior belief p_0 , the highest equilibrium payoff the type-1 Agent can receive in one round is given by the dotted black segment. If instead information is disclosed in two steps, with an intermediate belief p_1 , the type-1 Agent's payoff becomes the sum of the two solid (red) segments, which is strictly more, because $w(p)/p$ is strictly increasing. The right panel illustrates the jump in beliefs that occurs over the relevant interval when $w(p)/p$ is not strictly increasing, as occurs in our leading example for $p < p^*$.

Another extreme feature of our model is that how much information the type-0 Agent possesses is irrelevant to the Firm's decision. In some cases, this seems like a reasonable assumption: having interesting ideas about how to build an electric bulb does not amount to much if some key steps are missing. On the other hand, presenting positive yet inconclusive evidence that a rogue regime has weapons of mass destruction might affect the decision-maker's opinion, and hence, his decision, even when it is common knowledge that this evidence is inconclusive.

To model this, we may assume that, if the Firm's belief is p , but the Agent fails to disclose any further information, the posterior belief falls to $\lambda(p) < p$, where λ is some non-decreasing function (and there is no more revelation of information). We have focused our attention so far on the special case in which $\lambda(p) = 0$ for all $p < 1$. We could as well have picked an arbitrary (non-decreasing) function λ . In that case, following the reasoning presented in the preceding paragraphs, splitting information over some range of beliefs $[\underline{p}, \bar{p}]$ would lead to revenues

$$w(\bar{p}) - w(\underline{p}) - \int_{\underline{p}}^{\bar{p}} \frac{w(p) - w(\lambda(p))}{p - \lambda(p)} dp,$$

while jumping from belief \underline{p} to belief \bar{p} in one round would yield

$$w(\bar{p}) - w(\underline{p}) - (\bar{p} - \underline{p}) \frac{w(\bar{p}) - w(\lambda(\bar{p}))}{\bar{p} - \lambda(\bar{p})}.$$

Hence, it is clear that splitting information at p is desirable if and only if the function w is star-shaped at $\lambda(p)$, i.e. if and only if $(w(p) - w(\lambda(p)))/(p - \lambda(p))$ is strictly increasing. Note that this is automatically satisfied if w is strictly convex, independently of the function λ .

5 Inducing Effort

We now provide an example that shows why equilibria that maximize $V_1 - V_0$ may be of special interest when our game of selling information is embedded into a larger game in which there has been previous investment in obtaining information.

Suppose that there are two underlying states of Nature: $s_N \in \{L, H\}$. Nature chooses first s_N with probability $\rho = \Pr[s_N = H]$, which is common knowledge.

The Agent then decides privately how much effort to exert to obtain an informative private signal $\omega \in \{0, 1\}$ (which is the type of the Agent in the model above). In particular, if she puts in effort $c(\alpha)$, for some $\alpha \in [0, 1)$, then the conditional probability distribution over ω is

$$\mathbb{P}[\omega = 0 | s_N = 0] = 1,$$

$$\mathbb{P}[\omega = 1 | s_N = 1] = \alpha.$$

Hence, assuming that the Agent obtains signal 1, the probability of the state being H is 1, whereas assuming that the Agent obtains signal 0, the probability of the state being H is

$$\mathbb{P}[s_N = H | \omega = 0] = \frac{\rho(1-\alpha)}{\rho(1-\alpha) + (1-\rho)} = \rho \frac{1-\alpha}{1-\alpha\rho} < \rho.$$

In words, obtaining signal $\omega = 1$ can be interpreted as finding evidence that the state of Nature is $s_N = H$ and obtaining signal $\omega = 0$ means a lack of such evidence, but a lack of evidence implies that the probability expressed in the posterior belief is lower than that expressed in the prior. (The signals are mapped to the distribution over the set \mathcal{I} , which for simplicity depends on the realized signal, but not on the chosen effort).

Ex ante, if the Firm expects the Agent to exert effort $c(\alpha)$, he assigns prior probability

$$p_0 = \rho\alpha$$

to the Agent having the $\omega = 1$ signal.

Suppose that, once the Agent learns her private signal, players play our original game, with K rounds of communication and transfers, followed by the Firm's investment decision. Suppose that "Invest" yields payoff 1 if $s_N = H$ and payoff $-\hat{\gamma}$ if $s_N = L$. "Not Invest" yields a safe payoff 0, as before. These payoffs map onto expected payoffs conditional on the Agent's signal/type as follows. If the Agent's type is $\omega = 1$, the expected return to investment is 1. If the Agent's type is 0, the expected return to investment is

$$\rho \frac{1-\alpha}{1-\alpha\rho} - \hat{\gamma} \left(1 - \rho \frac{1-\alpha}{1-\alpha\rho} \right) = -\hat{\gamma} + (1 + \hat{\gamma}) \rho \frac{1-\alpha}{1-\alpha\rho}.$$

Redefine $\gamma := \hat{\gamma} - (1 + \hat{\gamma}) \rho \frac{1-\alpha}{1-\alpha\rho}$ and assume that $\rho < \frac{\hat{\gamma}}{\hat{\gamma}+1}$. That yields $\gamma > 0$, as in our

original model. If the Firm learns that $\omega = 0$, the optimal action is “Not Invest,” whereas if the Firm learns that $\omega = 1$, the optimal action is to “Invest” and $w(p) = (p - \gamma(1 - p))^+$. Finally, straightforward algebra shows that $\rho < \frac{\hat{\gamma}}{\hat{\gamma}+1}$ implies that $p_0 < p^* = \frac{\gamma}{\gamma+1}$ for any α . This concludes the description of how to map this extension onto our original model, with the interpretation that the Agent’s type is not the relevant state of Nature, but rather her informative signal about it.

Assume now that $c(\alpha)$ is differentiable, strictly increasing, and convex. The first-best effort level (i.e., the level that maximizes total expected surplus) maximizes

$$\rho\alpha - c(\alpha),$$

and so the first-best effort α solves

$$c'(\alpha_{FB}) = \rho.$$

In turn, if in equilibrium the Agent expects to receive payoff V_1 if the signal is 1, and payoff V_0 if the signal is 0, she will choose an effort level α that maximizes

$$\rho\alpha V_1 + (1 - \rho\alpha)V_0 - c(\alpha),$$

which is maximized at

$$c'(\alpha_A) = \rho(V_1 - V_0).$$

In (a pure-strategy) equilibrium, the Firm’s beliefs over the chosen α have to be consistent with the Agent’s choice. Due to the fact that $V_1 - V_0 < 1$, no equilibrium can support the efficient effort level. Moreover, we have shown that with only one round of communication, $V_0 = V_1$; hence, the Agent would not put any effort!

What is the second-best effort level α that is sustainable in equilibrium for a general K ? Note that $V_1 - V_0$ depends on the Firm’s expectation about α . If we maximize $V_1(p_0) - V_0(p_0)$ for every candidate α , then every α such that

$$c'(\alpha) \leq \rho \max \{V_1(p_0) - V_0(p_0)\}$$

can be supported in an equilibrium of our extended game (in which the maximization is over equilibria with a given K and p_0). Due to the fact that the total surplus is concave in α , the largest (and most efficient) effort level that can be supported in equilibrium requires a continuation equilibrium that maximizes $V_1(p_0) - V_0(p_0)$.

Of course, the equilibrium that we constructed is not the unique one that maximizes incentives to exert effort. Due to the fact that in our equilibrium $V_0(p_0) > 0$, there exists a whole continuum of equilibria for which $V_1(p_0) - V_0(p_0)$ is the same as in our equilibrium, but that differ in the payoff achieved by the Firm. In particular, any equilibrium in which at the beginning of the communication stage the Agent pays the Firm some $X \in [0, \bar{X}]$ is an equilibrium with the same difference $V_1(p_0) - V_0(p_0)$, where the bound

$$\bar{X} := \frac{p - (p + \gamma(1 - p))^+ + \gamma p \min\{\ln p^*, \ln p\}}{1 - p}$$

is the payoff of the type-0 Agent in the equilibrium we constructed above.

The model for moral hazard we describe here is about incentives to obtain information. An analogous model can be constructed in which the Agent observes the state s_N for sure, but by exerting effort she can affect the probability of the state being 1 (i.e. she can increase ρ).

6 Intermediaries

6.1 The role of an intermediary

We have seen how the type-1 Agent, by selling information progressively, can obtain a payoff above the expected value p_0 . Indeed, for $p_0 < p^*$, with sufficiently many rounds, this payoff approaches $-\gamma \ln p^*$. Nevertheless, this still falls short of the full value of information in this event. Without it, the Firm obtains a payoff of 0. With it, the Firm would obtain a payoff of 1. Therefore, the actual value of information is one, but of course the Firm does not know it, and if he did, he would no longer be willing to pay for it. Could more general mechanisms help the type-1 Agent to appropriate this value?

Note that, so far, the posterior belief of the Firm was either above its prior belief, or equal

to zero. This is a consequence of the focus on pure strategies. If, in equilibrium, the type-1 Agent randomized over disclosing one or the other piece of information, the Firm's posterior belief, conditional on seeing one of these pieces, could fall short of the prior belief yet be strictly positive. For this to happen, it would suffice that the likelihood that type-1 Agent will disclose the other piece of information exceeds a certain level.

Yet, as we have already mentioned, there is no mixed-strategy equilibrium that yields a higher payoff to the type-1 Agent. Very roughly, this is because she would have to be indifferent over revealing either piece of information, so her payoff could not exceed that which she would obtain from disclosing the piece of information that leads to the Firm having a stronger posterior belief.¹³

Suppose now that the Agent has access to a trusted, disinterested intermediary –a mediator, in the game-theoretic sense. This intermediary, upon receiving the information ι , could decide to randomize between releasing this piece of information or not. Therefore, if the piece of information is not disclosed, the Firm cannot draw a definite conclusion from this. He will not know whether the information was not disclosed because the Agent did not possess it or because it was censored by the intermediary. In this case, the Firm's posterior belief could fall below his prior belief, yet remain positive. Note that the Agent may not be able to perform this garbling herself, because she need not be indifferent over the two resulting outcomes.

Alternatively, our intermediary might embody a richer set of signals. In particular, the intermediary can be replaced by tests that the Agent can choose to publicly perform, but whose outcome is correlated only imperfectly with the information that the Agent owns. For instance, the outcome of some supervised exploratory drilling that is authorized by the Agent is certainly correlated with the size of the underlying oil field, but it can affect the Firm's belief in either direction and to various degrees.

We will show that such elementary obfuscation of outcomes, which leads to only two possible posterior beliefs, is sufficient to achieve the best equilibrium asymptotically. However, one can conceive of more complicated methods of obfuscation. For instance, the Agent might be expected

¹³Introducing a public randomization device would not help either, because making the randomization observable only helps to make the equilibrium payoff set convex, which is of no use here. Similarly, allowing for cheap talk between the Agent and the Firm would not change the findings of our baseline model either, because the two types of Agent have perfectly aligned interests.

to transmit two pieces of information to the intermediary, and, depending upon which pieces of information it obtains, the intermediary could randomize over disclosing none, one or the other, or both of these pieces of information. This would lead to four possible posterior beliefs. Finally, the Agent could reveal her signal to the intermediary fully, which in turn would send a signal that was correlated arbitrarily with the Agent's type.

6.2 An Illustration

Consider the simple example in which $\gamma = 1$, so that $p^* = 1/2$. The right panel of Figure 8 illustrates one of the procedures that the intermediary may follow, starting from a given belief $p_0 = 1/3$. Here, the intermediary sends one of two messages, low or high. The high message makes the Firm more optimistic, with a corresponding posterior belief of $1/2$. The low message makes the Firm more pessimistic, with a posterior belief of $1/6$. Due to the fact that the Firm's belief is a martingale, and because $p_0 = 1/3$ is the mean of $1/2$ and $1/6$, the two messages must be equally likely from the Firm's point of view.

How likely is each message from the type-1 Agent's point of view? Note that the low posterior belief, $1/6$, is half as high as the prior belief, $1/3$. This means that, from the Firm's point of view, the low message is half as likely to be observed as the high message when the state is $\omega = 1$. Due to the fact that he assigns an unconditional probability of $1/2$ to the low message, he must then assign probability $1/2 \cdot 1/2 = 1/4$ to this low message conditional on the Agent being of type 1. This is then the probability that the type-1 Agent must assign to this low message.

The left panel of Figure 8 depicts the three continuation payoffs in the best equilibrium that we characterized in Section 4.3, without an intermediary, starting from a Firm's belief $1/6$. The type-1 Agent receives $-\gamma \ln p^* = \ln 2$, the Firm receives $w(1/6) = 0$, yet the sum of all three payoffs must equal the surplus $p = 1/6$, so that the type-0's Agent payoff can be read off the y -axis as shown. Note that $0 < V_0 \leq V_1$.¹⁴

Consider then the following scheme when there are many rounds. In the second round, it

¹⁴Note also that, because the intermediary is trusted, the payoff V_0 need no longer be an expectation over the payoff of the type-0 Agent's payoff as a function of the specific information she owns, because the intermediary does not need to release information to the Firm. Thus, in this section, we shall always think of the type-0 Agent as receiving a payoff that is not contingent on the specific amount of incomplete information that she possesses.




Figure 8: The Role of an Intermediary

is understood that the Agent will make a payment of V_0 to the Firm if and only if the realized message is low (in particular, there is no payment by the Agent to the Firm if the realized message is high). Aside from this one-time, conditional payment from the Agent to the Firm, all payments by the Firm to the Agent and all disclosures of information by the Agent to the Firm, occur from the second round onward, as in the equilibrium without an intermediary (which is obviously possible even with an intermediary), given the realized message.

In the initial round, before the message is sent, the Firm must pay the difference between his expected continuation payoff and his current outside option, 0. If he fails to do so, no further messages are sent. How much is the Firm willing to pay? Note that, if there was no payment from the Agent to the Firm conditional on a low message, he is not willing to pay anything, because his outside option after either message is still 0. Nevertheless, because he expects to receive V_0 if an event occurs whose probability is $1/2$ from his point of view, he is willing to pay up to $V_0/2$ upfront in this scheme. How much is this scheme worth to the type-1 Agent? Her expected payoff is

$$\frac{1}{2}V_0 + \frac{1}{4}(V_1 - V_0) + \frac{3}{4}V_1 = V_1 + \frac{1}{4}V_0 > V_1.$$

To see this, note that she gets $V_0/2$ upfront, $V_1 - V_0$ in the event that the message is low (an

event to which she assigns probability $1/4$), and V_1 in the event that the message is high. As a result, with this scheme, her payoff with a prior belief of $1/3$ is strictly larger than V_1 , which is her maximal payoff without an intermediary.

Observe first that such a scheme is not possible without an intermediary, because the type-1 Agent is not indifferent over realized messages. She strictly prefers the high message to obtain, so that such a scheme cannot be replicated by mixed strategies without an intermediary. Second, note that the payment that the Agent makes if a low message occurs is not informative *per se*. This is because this payment is no larger than V_0 and the continuation payoffs of the Agent is at least as much, independently of her type. Higher payments would not work, because the type-0 Agent would not be willing to pay given the continuation equilibrium, so the occurrence of a payment or not would convey information about the Agent's type. From the left panel of Figure 8, it is clear that, the closer the expected payoff pV_1 of the type-1 Agent is to the total surplus p , the smaller is the resulting V_0 ; hence, the smaller the scope for such a scheme becomes. Yet as long as V_0 remains strictly positive, such schemes remain possible.

This scheme is nothing but a bet, or a trade, between two agents whose beliefs differ as to the probability of some event's occurring. The type-1 Agent attaches probability $1/4$ to the event that the Firm's posterior belief will be $1/6$, while the Firm attaches probability $1/2$ to this event. Therefore, there is room for a profitable trade, the only bound on which is that the bet cannot exceed the type-0 Agent's continuation payoff. Note that the type-0 Agent loses from this scheme (as compared to our original equilibrium), because she is the one who assigns a high probability to the event that the posterior belief is $1/6$. Still, her payoff remains positive and she has no choice but to go along, because her payoff is equal to the Firm's payment in the initial period.

It is then clear how an intermediary might help. Without an intermediary, it was possible to drive the Firm's payoff to his outside option. With an intermediary, we can further drive the type-0 Agent's payoff down; indeed, as we shall see, we can drive it all the way down to her minmax payoff, namely zero. In this fashion, the intermediary allows the type-1 Agent to extract the full surplus. In the next subsection, we shall state this insight more formally.

6.3 The Optimization Problem

In what follows, there is not much to gain from continuing to use our leading example. The outside option w of the Firm, then, is an arbitrary continuous function with $w(0) = 0, w(1) = 1$, with the feature that full information disclosure is efficient, i.e. $w(p) \leq p$ for all $p \in [0, 1]$.

The procedure used by the intermediary can be summarized by a distribution $F_k(\cdot|p)$ over the Firm's posterior beliefs, given the prior belief p , and given the number of rounds k .¹⁵ Due to the fact that this distribution is known, the Firm's belief must be a martingale, which means that, given p ,

$$\int_0^1 p' dF_k(p'|p) = p, \text{ or } \int_0^1 (p' - p) dF_k(p'|p) = 0. \quad (5)$$

To put it differently, $F(\cdot|p)$ is a mean-preserving spread of the Firm's prior belief p .

Given such a distribution, and some equilibrium to be played in the continuation game for each resulting posterior belief p' , how much is the Firm willing to pay up front? Again, this must be the difference between his continuation payoff and his outside option, namely

$$\bar{t}_k^F := \int_0^1 (w(p') + X(p')) dF_k(p'|p) - w(p),$$

where, as before, $X(p')$, or X' for short, denotes the Firm's payoff, net of the outside option, in the continuation game, given that the posterior belief is p' .

Assume that the distribution $F(\cdot|p)$ assigns probability q to some posterior belief p' . This means that the Firm attaches probability q to his next posterior belief turning out to be p' . What is the probability q_1 assigned to this event by the type-1 Agent? This must be qp'/p , because

$$p' = \mathbb{P}[\omega = 1|p'] = \frac{pq_1}{q},$$

where the first equality from the definition of the event p' , and the second follows from Bayes' rule, given the prior belief p . This is an obvious generalization of the calculation that we have seen in the simple example of Subsection 6.2.

¹⁵As mentioned, we can think of the Agent as disclosing all her information to the intermediary initially; then, in every round, the intermediary releases a noisy signal correlated with the true state of the world in such a way that the distribution of the posterior belief is F_k .

Therefore, the maximal payoff that the type-1 Agent expects to receive from the next round onward is

$$\int_0^1 V_{1,k-1}(p', X') \frac{p'}{p} dF_k(p'|p),$$

where, as before, $V_{1,k-1}(p', X')$ denotes the maximal payoff of the type-1 Agent, with $k-1$ rounds to go, given that the Firm's payoff, net of his outside option, is X' and his belief is p' .

Combining these two observations, we obtain that the payoff of the type-1 Agent is at most

$$\int_0^1 (w(p') + X') dF_k(p'|p) - w(p) + \int_0^1 V_{1,k-1}(p', X') \frac{p'}{p} dF_k(p'|p), \quad (6)$$

and our objective is to maximize this expression, for each p , over all distributions $F_k(\cdot|p)$, as well as mappings $p' \mapsto X' = X(p')$ (subject to (5) and the feasibility of X').

6.4 The Optimal Transfers

As a first step in the analysis of this problem, we will establish the following.

Lemma 2 *Fix the prior belief p and the number of remaining rounds k . The best equilibrium payoff of the type-1 Agent, as defined by (6), is achieved by setting, for each $p' \in [0, 1]$, the Firm's net payoff in the continuation game defined by p' equal to*

$$X(p') = \begin{cases} X^*(p') & \text{if } p' < p, \\ 0 & \text{if } p' \geq p, \end{cases}$$

where

$$X^*(p') := \frac{p'(1 - V_{1,k-1}(p')) - w(p)}{1 - p'}.$$

The type-1 Agent's continuation payoff is then given as

$$V_{1,k-1}(p', X^*(p')) = V_{1,k-1}(p') - X^*(p').$$

Proof: See Appendix B. □

This lemma formalizes the intuition from the example that we used in Subsection 6.2: it is best to promise as high a rent as possible to the Firm if the posterior belief is lower than the prior belief, and as low as possible if it is higher. The function X^* describes this upper bound. As in the example, this bound turns out to be the entire continuation payoff of the type-0 Agent in the best equilibrium for the type-1 Agent with $k - 1$ periods to go. We can express this bound in terms of the Firm's belief and the type-1 Agent's continuation payoff, given that the equilibrium is efficient. Of course, it is possible to give even higher rents to the Firm, provided that the equilibrium that is played in the continuation game gives the type-0 Agent a higher payoff than the equilibrium that is best for the type-1 Agent. The proof of this lemma establishes that what is gained in the initial period by considering higher rents is more than offset by what must be relinquished in the continuation game, in order to generate a high enough payoff to the type-0 Agent.

The key intuition here is that the type-1 Agent assigns a higher probability to the event that the posterior belief will be $p' > p$ than does the Firm and conversely, a lower probability to the event that $p' < p$, because she knows that the state is 1. Therefore, the type-1 Agent wants to offer the Firm an extra continuation payoff in the event that $p' < p$ (and collect extra money for it now), and offer as small a continuation payoff as possible in the event that $p' > p$. Given that the Agent and the Firm have different beliefs, there is room for profitable bets, in the form of transfers whose odds are actuarially fair from the Firm's point of view, but profitable from the point of view of the type-1 Agent. Such bets were not possible without the intermediary (at least in pure strategies), because, at the only posterior belief lower than p , namely $p' = 0$, there was no room for any further transfer in this event (because there was no further information to be sold).

6.5 The Value of an Intermediary

Having solved for the optimal transfers, we may now focus on the issue of identifying the optimal distribution $F_k(\cdot|p)$. Plugging in our solution for X' into (6), we obtain that

$$V_{1,k}(p) = \sup_{F_k(\cdot|p)} \int_0^1 v_{k-1}(p'; p) dF_k(p'|p) - w(p), \quad (7)$$

where

$$v_{k-1}(p'; p) := \begin{cases} w(p') + \frac{p-p'}{p} X^*(p') + \frac{p'}{p} V_{1,k-1}(p') & \text{for } p' < p, \\ w(p') + \frac{p'}{p} V_{1,k-1}(p') & \text{for } p' \geq p, \end{cases}$$

and the supremum is taken over all distributions $F_k(\cdot|p)$ that satisfy (5), namely, $F_k(\cdot|p)$ must be a distribution with mean p .

This optimality equation cannot be solved explicitly. Nevertheless, the associated operator is monotone and bounded above. Therefore, its limiting value as we let k tend to infinity, using the initial value $V_{1,0}(p) = 0$ for all p , converges to the smallest (positive) fixed point of this operator. This fixed point gives us the limiting payoff of the type-1 Agent as the number of rounds grows without bound.

It turns out that we can guess this fixed point. One of the fixed points of (7) is

$$V_1(p) = \frac{p - w(p)}{p}$$

Recall that this value is the upper bound on $V_{1,k}(p)$ that we derived earlier, so it is the highest payoff that we could have hoped for. In fact, we have

Theorem 2 *The function defined by, for all p ,*

$$V_1^{int}(p) := \frac{p - w(p)}{p}$$

is the limiting value of the game to the type-1 Agent, as the number of rounds tends to infinity.

Proof: See Appendix C. □

Figure 9 provides a side-by-side comparison between the payoff to the type-1 Agent in the best equilibrium with and without intermediary. In this case, this value is

$$V_1^{int}(p) = \begin{cases} 1 & \text{for } p \leq p^*, \\ \gamma \frac{1-p}{p} & \text{for } p \geq p^*. \end{cases}$$




Figure 9: Payoffs with or without an intermediary ($\gamma = 1$, V_{int} is the payoff with an intermediary)

As is clear from the proof, the intermediary manages to extract the full value for the information through bets similar to the one discussed in Subsection 6.2 and very small changes in the Firm’s belief. That is, from any belief p onward, it assigns equal probability to the posterior beliefs $p - \varepsilon$ and $p + \varepsilon$, for some small $\varepsilon > 0$, accompanied by an upfront payment by the Firm that includes a maximal payback by the Agent to the Firm in case the posterior belief decreases. The maximal payback is determined by the continuation value of the type-0 Agent. The full value of the information is then extracted asymptotically as ε tends to 0 and the number of rounds tends to infinity. This may not be the only way that the intermediary can achieve this maximum. In our leading example, for instance, in which $w(p) = (p - (1 - p)\gamma)^+$, whenever the belief is $p > p^*$, this belief can be split into two posterior beliefs that are (arbitrarily close to) p^* and 1. However, in this case, as in our proof for the general case, the number of “bets” that are required to extract the full value of the information is arbitrarily large.

7 Final Remarks

The model can be extended in many ways. For example:

1. Suppose that there is discounting with every round of communication and that the Firm can decide to make his decision as to whether or not to invest in the venture before the K rounds are over. Then it is no longer true that adding another round of communication will strictly increase the Agent's maximal equilibrium payoff. This is because (a) the Agent faces a trade-off between collecting more money overall and collecting it earlier, and (b) the Firm will ultimately prefer to take his outside option rather than wait for another period, once the possible benefits from waiting become too small. Hence, in the best equilibrium, the number of rounds in which communication actually takes place is bounded, so that the exact number of rounds available will be of no importance, provided that there is a large number of them. However, as long as the players are not impatient, the best equilibrium still involves a gradual release of information and the number of rounds of active communication increases with the discount factor. While this version with discounting does not lend itself to closed-form formulas, it is easy to see that, as the discount factor approaches 1, the payoff to the type-1 Agent must tend to the payoff in the undiscounted game. Furthermore, in our leading example, numerical simulations show that this convergence occurs at a geometric rate.
2. Suppose that the Agent cares to some extent that the Firm takes the correct action (say, *ceteris paribus*, her payoff increases by some small $\varepsilon > 0$). Then, in the one-shot game, it is dominant for the Agent to reveal all her information. As a result, the Firm will not make any payment. This logic clearly extends to the finite horizon game, no matter how long the horizon is. On the other hand, this unraveling argument does not extend to the infinite-horizon game (say, with little but positive discounting), and it is possible to construct equilibria in our leading example in which the Agent is paid for a gradual release of information. Of course, the value of ε does put bounds on how extreme the Firm's posterior belief can become before the Agent discloses all information. Nevertheless, our

results are robust, inasmuch as the maximal equilibrium payoff to the type-1 Agent will be continuous at $\varepsilon = 0$.

3. A strong assumption of our model is that the Firm knows the set of relevant actions. Yet in many cases of interest, the Firm does not know what these are. How does one invest in an invention before knowing what this invention is about? Of course, if the Firm is entirely clueless about what to do until all the information is disclosed, it is fairly clear what the type-1 Agent's optimal strategy involves: the Agent should reveal to the Firm all but a small detail of the invention without any delay, thereby strengthening as much as possible the Firm's posterior belief p that the invention is valuable, and then sell the information for p . (This strategy does not work in our model because once the Firm's belief is strong enough, the Firm can get a high payoff by investing without any further information.) A more interesting model would allow the Firm to know *a priori* and privately some of the key parts of the invention: in that case, an Agent who reveals all but a few elements would risk the Firm already knowing the missing pieces. In that case, we conjecture that again it would be optimal for the successful inventor to release information more gradually.

References

- [1] Admati, A.R. and M. Perry (1991). "Joint Projects without Commitment," *The Review of Economic Studies*, **58**, 259–276.
- [2] Admati, A. R. and P. Pfleiderer (1990). "Direct and Indirect Sale of Information," *Econometrica*, **58**, 901–928.
- [3] Admati, A. R. and P. Pfleiderer (1988). "Selling and Trading on Information in Financial Markets," *The American Economic Review*, *P&P*, **78**, 96–103.
- [4] Anton, J. and D. Yao (1994). "Expropriation and Inventions," *American Economic Review*, **84**, 190–209.

- [5] Anton, J. and D. Yao (2002). “The Sale of Ideas: Strategic Disclosure, Property Rights, and Contracting,” *Review of Economic Studies*, **69**, 513–531.
- [6] Arrow, K. J. (1959). “Economic Welfare and the Allocation of Resources for Invention,” Rand Research Paper P-1856.
- [7] Aumann, R. J. and S. Hart (2003). “Long Cheap Talk,” *Econometrica*, **71**, 1619–1660.
- [8] Aumann, R. J. and M. B. Maschler (1995). *Repeated Games with Incomplete Information*, The MIT Press.
- [9] Che, Y.-K. and J. Sákovics (2004). “A Dynamic Theory of Holdup,” *Econometrica*, **72**, 1063–1103.
- [10] Compte, O. and P. Jehiel (2004). “Gradualism in Bargaining and Contribution Games,” *The Review of Economic Studies*, **71**, 975–1000.
- [11] Eső, P. and B. Szentes (2007). “The price of advice,” *The Rand Journal of Economics*, **38**, 863–1179.
- [12] Forges, F. (1990). “Equilibria with Communication in a Job Market Example,” *The Quarterly Journal of Economics*, **105**, 375–398.
- [13] Forges, F. and F. Koessler (2008). “Long Persuasion Games,” *Journal of Economic Theory*, **143**, 1–35.
- [14] Fudenberg, D. and J. Tirole (1991). *Game Theory*, MIT Press, Cambridge.
- [15] Harsanyi, J.C. (1967/68). “Games with Incomplete Information Played by Bayesian Players,” parts I-III, *Management Science*, **14**, 159–182, 320–334, 486–502.
- [16] Kamenica, E. and M. Gentzkow (2009). “Bayesian Persuasion,” working paper, University of Chicago.

- [17] Goldwasser, S. S. Micali and C. Rackoff (1985). “The knowledge complexity of interactive proof-systems,” In *Proc. 17th ACM Symp. on Theory of Computing*, Providence, ACM, 291–304.
- [18] Gul, F. (2001). “Unobservable Investment and the Hold-Up Problem,” *Econometrica*, **69**, 343–376.
- [19] Landsberger, M. and I. Meilijson (1990). “Lotteries, insurance and star-shaped utility functions,” *Journal of Economic Theory*, **52**, 1–17.
- [20] Marx, L. M. and S. A. Matthews (2000). “Dynamic Voluntary Contribution to a Public Project,” *The Review of Economic Studies*, **67**, 327–358.
- [21] Rosenberg, D., Solan, E. and N. Vieille (2009). “Strategic Information Exchange,” working paper, HEC Paris.

Appendix

A. Proof of Proposition 3

We provide here the missing details for the proof of Proposition 3. In Section 3, the result has been established for pure strategies. What remains to be shown is that there exists no mixed-strategy equilibrium yielding a higher payoff V_1 . Consider then a candidate mixed-strategy equilibrium. This strategy profile can be summarized by a distribution $F_{k+1}(\cdot|p)$ that is used by the Agent (on the equilibrium path) with $k + 1$ rounds left, given belief p , and the continuation payoffs $W_k(\cdot)$ and $V_k(\cdot)$. As before, we may assume that the equilibrium is efficient, and so we can assume that, given that the Firm obtains a net payoff of X_k (i.e., given that $W_k = (p - (1 - p)\gamma)^+ + X_k$), the type-1 Agent receives $V_{1,k}(p, X_k)$, the highest payoff to this type given that the Firm receives at least a net payoff of X_k . Since $V_{1,k}$ maximizes the sum of the Firm's and type-1 Agent's payoff (see Lemma 1), we have that, for all k, p and $X \geq 0$,

$$V_{1,k}(p, X) \leq V_{1,k}(p) - X.$$

The payoff $V_{1,k+1}(p)$ of the type-1 Agent is at most, with $k + 1$ rounds to go,

$$\sup_{F_{k+1}(\cdot|p)} \int_0^1 \left[(p' - (1 - p')\gamma)^+ + X_k(p') + V_{1,k}(p', X_k(p')) \frac{p'}{p} \right] dF_{k+1}(p'|p) - (p - (1 - p)\gamma)^+,$$

where the supremum is taken over all distributions $F_{k+1}(\cdot|p)$ that satisfy

$$\int_{[0,1]} (p' - p) dF_{k+1}(p'|p) = 0,$$

i.e. such that the belief of the Firm follows a martingale. To emphasize the importance of the posterior $p' = 0$, we alternatively write this constraint as $\int_0^1 (p' - p) dF_{k+1}(p'|p) = pF_{k+1}(0|p)$, where $\int_0^1 dF_{k+1}(p'|p) := 1 - F_{k+1}(0|p)$.

If the type-1 Agent randomizes, he must be indifferent between all elements in the support of his mixed action, that is, for all $p' > 0$ in the support of $F_{k+1}(\cdot|p)$, $V_{1,k}(p', X') = \underline{V}_k$, for some \underline{V}_k independent of p' . (The existence of such a mixed-strategy equilibrium would also impose

additional restrictions, but since our purpose is to show that such an equilibrium cannot improve upon the pure-strategy equilibrium, we disregard them here.) Relaxing possibly the constraints (which can only increase the value of the objective), assume then that, for all such values of p' , and $X \geq 0$,

$$V_{1,k}(p', X) = V_{1,k}(p') - X.$$

By substitution, we obtain that $V_{1,k+1}(p)$ is at most equal to, for some \underline{V}_k ,

$$\begin{aligned} & \sup_{F_{k+1}(\cdot|p)} \int_0^1 \left[(p' - (1-p')\gamma)^+ + V_{1,k}(p') - \underline{V}_k + \underline{V}_k \frac{p'}{p} \right] dF_{k+1}(p'|p) - (p - (1-p)\gamma)^+ \\ &= \sup_{F_{k+1}(\cdot|p)} \int_0^1 \left[(p' - (1-p')\gamma)^+ + V_{1,k}(p') \right] dF_{k+1}(p'|p) + F_{k+1}(0|p)\underline{V}_k - (p - (1-p)\gamma)^+. \end{aligned}$$

Because $\underline{V}_k \leq V_{1,k}(p')$ for all $p' > 0$ in the support of $F_{k+1}(\cdot|p)$, it is then clear that this is largest when $\underline{V}_k = \min_{p' \in \text{supp } F_{k+1}(\cdot|p)} V_{1,k}(p')$. We claim that, along with the structure of the continuation payoff, this observation is sufficient to ensure that the optimal distribution $F_{k+1}(\cdot|p)$ assigns positive probability to only two posterior beliefs, namely 0 and some $p' > p$. This will establish that attention can be restricted to pure strategies.

The claim is proved by induction. With only round left, the result is obvious, because the Firm cannot pay more than $p - (p - (1-p)\gamma)^+$, and this is what the Firm pays in the pure-strategy equilibrium.

Assume that, defining Z_k by $Z_k(p) := V_{1,k}(p) + (p - (1-p)\gamma)^+ - p$, all p , we have

$$Z_k(p) = V_{1,k-1}(p^*) - p \text{ if } p < p_k := p^{\frac{k}{k-1}}, \text{ and } Z_k(p) = \gamma[k(1 - p^{1/k}) - (1-p)] \text{ otherwise.}$$

These are the payoffs with k rounds to go, in the pure-strategy equilibrium, and our induction hypothesis is that these are also the highest equilibrium payoffs. Given our earlier observation,

$Z_{k+1}(p)$ is at most

$$\begin{aligned} & \sup_{F_{k+1}(\cdot|p)} \int_0^1 [p' + W_k(p')] dF_{k+1}(p'|p) + F_{k+1}(0|p) \min_{p' \in \text{supp } F_{k+1}(\cdot|p)} (Z_k(p') + p' - (p' - (1 - p')\gamma)^+) \\ &= \sup_{F_{k+1}(\cdot|p)} \int_0^1 W_k(p') dF_{k+1}(p'|p) + p + F_{k+1}(0|p) \min_{p' \in \text{supp } F_{k+1}(\cdot|p)} (Z_k(p') + p' - (p' - (1 - p')\gamma)^+). \end{aligned}$$

Note that the function W_k is single-peaked (as is $V_{1,k}$). Therefore, the $F_{k+1}(\cdot|p)$ maximizing this expression can assign positive probability to only one strictly positive value of p' . Thus, the equilibrium with the highest payoff to the type-1 Agent with $k + 1$ periods to go involves a pure strategy, and it must be the one identified in the proof of Proposition 3, in the text (which gives the desired formula for the payoff Z_{k+1}). \square

B. Proof of Lemma 2

First of all, we must understand the function $V_{1,k}(p, X)$. Note that, as observed earlier, we can always assume that the equilibrium is efficient: take any equilibrium, and assume that, in the last round, on the equilibrium path, the type-1 Agent discloses the state. This modification can only relax any incentive (or individual rationality) constraint. This means that payoffs must satisfy (3), which provides a rather elementary upper bound on the maximal payoff to the type-1 Agent: in the best possible case, the payoffs X and $V_{0,k}(p, X)$ are zero, and hence we have

$$V_{1,k}(p) \leq \frac{p - w(p)}{p}.$$

Our observation from Lemma 1 that the equilibrium that maximizes the type-1 Agent's payoff also maximizes the sum of the Firm's and type-1 Agent's payoffs is obviously true here as well. Hence, any increase in X must lead to a decrease in $V_{1,k}(p, X)$ of at least that amount. As long as X is such that $V_{0,k}(p, X)$ is positive, we do not need to decrease $V_{1,k}(p, X)$ by more than this amount, because it is then possible to simply decrease the net transfer made by the Firm to the Agent in the initial period by as much. Therefore, either $V_{1,k}(p, X) = V_{1,k}(p) - X$, if X is smaller than some threshold value $X_k^*(p)$ (X^* for short), or $V_{0,k}(p, X) = 0$. By continuity, it must be

that, at $X = X^*$,

$$p(V_{1,k}(p) - X^*) + X^* + w(p) = p, \text{ or } X^* = \frac{p(1 - V_{1,k}(p)) - w(p)}{1 - p}.$$

Therefore, for values of X below X^* , we have that $V_1(p, X) = V_{1,k}(p) - X$, and this payoff is obtained from the equilibrium achieving the payoff $V_{1,k}(p)$ to the type-1 Agent, by reducing the net transfer from the Firm to the Agent in the initial round by an amount X . For values of X above X^* , we know that $V_{0,k}(p, X) = 0$, so that

$$V_{1,k}(p, X) \leq 1 - \frac{w(p) + X}{p}.$$

We may now turn to the issue of the optimal net payoff to grant the Firm in the continuation round. This can be done pointwise, for each posterior belief p' . The previous analysis suggests that, to identify what the optimal value of X' is, it is convenient to break down the analysis into two cases, according to whether or not X' is above X^* . Consider some posterior belief p' in the support of the distribution $F_k(\cdot|p)$. From (6), the contribution to the type-1 Agent's payoff from this posterior is equal to

$$w(p') + X' + V_{1,k-1}(p', X') \frac{p'}{p} \begin{cases} = w(p') + X' + (V_{1,k-1}(p') - X') \frac{p'}{p} & \text{if } X' \leq X^*(p'), \\ \leq w(p') + X' + \left(1 - \frac{w(p') + X'}{p'}\right) \frac{p'}{p} & \text{if } X' > X^*(p'). \end{cases}$$

Note that, for $X' > X^*(p')$, the upper bound to this contribution is decreasing in X' , and since this upper bound is achieved at $X' = X^*(p')$, it is best to set $X' = X^*(p')$ in this range. For $X' \leq X^*(p')$, this depends on p' : if $p' > p$, it is best to set X' to zero, while if $p' < p$, it is optimal to set X' to $X^*(p')$. To conclude, the optimal choice of X' is

$$X(p') = \begin{cases} X^*(p') & \text{if } p' < p, \\ 0 & \text{if } p' \geq p, \end{cases}$$

as claimed. □

C. Proof of Theorem 2

Recall that the function to be maximized is

$$\begin{aligned} & \int_0^p \left[w(p') + V_{1,k-1}(p') \frac{p'}{p} + \frac{p-p'}{p} \frac{p'(1-V_{1,k-1}(p')) - w(p')}{1-p'} \right] dF_k(p'|p) \\ & + \int_p^1 \left[w(p') + V_{1,k-1}(p') \frac{p'}{p} \right] dF_k(p'|p) - w(p), \end{aligned}$$

or re-arranging,

$$\int_0^p \left[\frac{1-p}{p} \frac{p'w(p') + p'V_{1,k-1}(p')}{1-p'} + \frac{(p-p')p'}{p(1-p')} \right] dF_k(p'|p) + \int_p^1 \left[w(p') + V_{1,k-1}(p') \frac{p'}{p} \right] dF_k(p'|p) - w(p).$$

Let us define $x_k(p) := p - w(p) - pV_{1,k}(p)$, and so multiplying through by p , and substituting, we get

$$\begin{aligned} p - w(p) - x_k(p) &= \int_0^p \left[\frac{1-p}{1-p'} (p'w(p') + p' - w(p') - x_{k-1}(p')) + \frac{(p-p')p'}{1-p'} \right] dF_k(p'|p) \\ &+ \int_p^1 [pw(p') + p' - w(p') - x_{k-1}(p')] dF_k(p'|p) - pw(p), \end{aligned}$$

or re-arranging,

$$\begin{aligned} x_k(p) &= p - w(p) - \int_0^p \left[\frac{1-p}{1-p'} ((p'-1)w(p') - x_{k-1}(p')) + p' \right] dF_k(p'|p) \\ &- \int_p^1 [p' - (1-p)w(p') - x_{k-1}(p')] dF_k(p'|p) + pw(p). \end{aligned}$$

This gives

$$x_k(p) = (1-p) \int_0^p \frac{x_{k-1}(p')}{1-p'} dF_k(p'|p) + \int_p^1 x_{k-1}(p') dF_k(p'|p) + (1-p) \int_0^1 (w(p') - w(p)) dF_k(p'|p).$$

Note that the operator mapping x_{k-1} into x_k , as defined by the minimum over $F_k(\cdot|p)$ for each p , is a monotone operator. Note also that $x = 0$ is a fixed point of this operator (consider $F_k(\cdot|p) = \delta_p$, the Dirac measure at p). We therefore ask whether this operator admits a larger

fixed point. So we consider the optimality equation, which to each p associates

$$x(p) = \min_{F(\cdot|p)} \left\{ (1-p) \int_0^p \frac{x(p')}{1-p'} dF(p'|p) + \int_p^1 x(p') dF(p'|p) + (1-p) \int_0^1 (w(p') - w(p)) dF(p'|p) \right\}.$$

It is standard to show that x is continuous on $(0, 1)$. Further, consider the feasible distribution $F(\cdot|p)$ that assigns probability $1/2$ to $p - \varepsilon$, and $1/2$ to $p + \varepsilon$, for $\varepsilon > 0$ small enough. This gives as upper bound

$$x(p) \leq \frac{1}{2} \cdot \frac{1-p}{1-p+\varepsilon} x(p-\varepsilon) + \frac{1}{2} \cdot x(p+\varepsilon) + (1-p) \left(\frac{w(p+\varepsilon) + w(p-\varepsilon)}{2} - w(p) \right), \quad (8)$$

or

$$\begin{aligned} x(p) + (1-p)w(p) &\leq \frac{1}{2} \cdot \frac{1-p}{1-p+\varepsilon} (x(p-\varepsilon) + (1-p+\varepsilon)w(p-\varepsilon)) \\ &\quad + \frac{1}{2} (x(p+\varepsilon) + (1-p-\varepsilon)w(p+\varepsilon)) + \varepsilon w(p+\varepsilon) \\ &= \frac{1}{2} (x(p-\varepsilon) + (1-p+\varepsilon)w(p-\varepsilon)) + \frac{1}{2} (x(p+\varepsilon) + (1-p-\varepsilon)w(p+\varepsilon)) \\ &\quad + \varepsilon \left(w(p+\varepsilon) - w(p-\varepsilon) - \frac{x(p-\varepsilon)}{1-p+\varepsilon} \right). \end{aligned}$$

Suppose that $x(p) > 0$ for some $p \in (0, 1)$. Then, since x is continuous, $x > 0$ on some interval I . Because w is continuous, the last summand is then negative for all $p \in I$, for $\varepsilon > 0$ small enough. This implies that the function $z : p \mapsto x(p) + (1-p)w(p)$ is convex on I , and therefore differentiable a.e. on I . Re-arranging our last inequality, we have

$$2 \left(w(p-\varepsilon) - w(p+\varepsilon) + \frac{x(p-\varepsilon)}{1-p+\varepsilon} \right) + \frac{z(p) - z(p-\varepsilon)}{\varepsilon} \leq \frac{z(p+\varepsilon) - z(p)}{\varepsilon}.$$

Integrating over I , taking limits as $\varepsilon \rightarrow 0$ and using the a.e. differentiability of z gives

$$\int_I \frac{x(p)}{1-p} \leq 0.$$

Because x is positive and continuous, it must be equal to zero on I . Because I is arbitrary, it follows that $x = 0$ on $(0, 1)$.

Because x is the largest fixed point of the optimality equation, and because the map defined by the optimality equation is monotone, it follows that the limit of the iterations of this map, applied to the initial value $x_0 : x_0(p) := p - w(p) - pV_{1,0}(p)$, all $p \in (0, 1)$, is well-defined and equal to 0. Given the definition of x , the claim regarding the limiting value of $V_{1,k}$ follows. \square