

SELLING INFORMATION

By

Johannes Hörner and Andrzej Skrzypacz

December 2009

Revised: August 2011

This Revision: November 2012

COWLES FOUNDATION DISCUSSION PAPER NO. 1743RR



**COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281**

<http://cowles.econ.yale.edu/>

Selling Information*

Johannes Hörner[†] and Andrzej Skrzypacz[‡]

November 11, 2012

Abstract

An Agent who owns information that is potentially valuable to a Firm bargains for its sale, without commitment and certification possibilities, short of disclosing it. We propose a model of gradual persuasion and show how gradualism helps mitigate the hold-up problem (that the Firm would not pay once it learns the information). An example illustrates how it is optimal to give away part of the information at the beginning of the bargaining, and sell the remainder in dribs and drabs. The Agent can only appropriate part of the value of information. Introducing a third-party allows her to extract the maximum surplus.

Keywords: value of information, dynamic game.

JEL codes: C72, D82, D83

*We thank Daron Acemoglu, Robert Aumann, Hector Chade, Alex Frankel, Drew Fudenberg, David Kreps, Max Kwiek, R.Vijay Krishna, Sofia Moroni, Romans Pancs, Arthur Robson, Rann Smorodinsky and seminar participants at the Barcelona JOCS, the Collegio Carlo Alberto, Turin, Essex University, the European University Institute, Harvard-MIT, University of Illinois at Urbana-Champaign, Northwestern University, Oxford University, Simon Fraser University, Stanford University, University of British Columbia, University of Western Ontario, UC San Diego, Yale, X-HEC Paris, Paris Sciences Economiques, Toulouse School of Economics, the European Summer Symposium in Economic Theory at Gerzensee 2012, the SED 2009 meeting, the Stony Brook 2010 game theory conference, and the 2011 Asian meeting of the Econometric Society 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.

1 Introduction

In the late 1960s, Italian inventor Aldo Bonassoli and Count Alain de Villegas contacted Elf, the French public oil company, to sell their discovery: a device that could detect oil fields from the air. Early tests were spectacularly successful. Contracts worth 200m Swiss Francs (in 1976), and 250m Swiss Francs (in 1978) were signed as tests proceeded, with the agreement of both French president Valery Giscard d'Estaing and Prime Minister Raymond Barre.

Unfortunately, the device was a hoax, exposed in 1979.¹ The story of the “Great Oil Sniffer Hoax” was made public in December 1983. Elf never completed paying for the final contract, but nevertheless had spent over \$150m.

But one cannot blame the “inventors” to be careful. Who has not heard of Robert Kearns, the inventor of the intermittent windshield wiper systems, who was a little too forthcoming with information about his invention with engineers of Ford?

This dilemma, known as the Arrow information paradox (1959) is the subject of this paper. The potential buyer of information needs to know its value before purchasing it since otherwise she may end up paying for a hoax. But often the only way to verify the information is to transmit it fully and once the seller has this detailed knowledge, she has no incentives to pay for it. This problem has obvious economic implications, as rewarding innovative activity is key in encouraging it. It is a problem not just for inventors, but owners of information in general: hedge funds claim to have special investment techniques which, of course, they cannot disclose; similarly, experts have confidential sources of information. Scientists and engineers claim to have superior and valuable knowledge –again, which cannot be disclosed.

We provide a game-theoretic analysis of the interaction between a buyer (Firm) and a seller (Agent) and examine when and how information should be transmitted, and payments made. In doing so, we determine how much of the information value can be appropriated by the seller, and how this problem is mitigated if sufficiently elaborate ways of transmitting information can

¹The early tests were successful because the inventors had an inside contact at Elf, who gave them maps of the oil fields that were being surveyed. The alleged map displayed by the device’s screen was simply a photograph. This was also how the hoax was uncovered: the two inventors routinely photocopied objects that were then placed behind a wall. Magically, the image of the object behind the wall appeared on the screen. Until an engineer bent the object, a ruler, right before it was placed behind the wall.

be applied, or a third party, such as a trusted intermediary, can be involved.

The game of information sale that we consider is dynamic. Within a round, the players make voluntary payments and then the Agent can disclose some information to the Firm. We assume that information is *verifiable* and *divisible*. In particular, in our model the Agent has one of two types (*i.e.*, she is knowledgeable or not) and the information transmission is modeled as *tests to verify the Agent's information*. Verifiability of information means that each test has a known difficulty: the competent Agent can always pass it (in the baseline model), but the incompetent Agent passes it with a probability commensurate to its difficulty. Easier tests have a higher probability of being passed by an incompetent Agent. Divisibility of information means that there is a rich set of tests with varying difficulties. There is no commitment on either side.

We construct tight bounds on the limits of the competent Agent's payoff as the number of possible communication rounds grows to infinity (they also establish bound on the difference of the competent and incompetent payoffs, which might be the relevant measure for those applications in which incentives to acquire information are explicitly taken into account). We characterize three such bounds: when we consider only pure-strategy equilibria (in which a competent Agent always passes the test), when we allow for mixed-strategy equilibria, and finally when we allow for tests so hard that even a competent Agent may occasionally fail. The latter case is equivalent to allowing for a trusted intermediary.

Lack of commitment creates a hold-up problem: since the Agent is selling information, once the Firm learns it, it has no reason to pay for it. Therefore, it seems at first difficult to make the Firm pay different amounts to different types, since such screening would inform the Firm about the Agent's type and lead it to renege on payments. Although we can make the Firm pay for a piece of information, it is necessary that it pays before it learns it.

That leads to our first main result that "splitting information" generally increases the competent Agent's payoff. That is, the competent Agent's payoff is higher in equilibria in which she takes two tests in a sequence (and is paid for each separately) than if she takes both of them at once (which is equivalent to taking one harder test). That intuition underlies the structure of the best equilibrium in pure strategies in our leading example: first, an initial chunk of information is given away for free that makes the Firm very uncertain about the Agent's competence. Then

the Agent sells information in dribs and drabs and gets paid a little for each bit. Although the expression for the limit payoff depends on the assumption that there is a rich set of tests and arbitrarily many rounds of communication, the benefit of splitting does not depend on either assumption.

Second, we show how randomization can help improve the performance of the contract. In the best pure-strategy equilibrium the incompetent Agent collects on average a non-trivial amount of payments, which leaves room for improvement. We first show that using (non-observable) mixed strategies can help by taking advantage of the fact that the competent Agent and the Firm may have different (endogenous) risk attitudes (more precisely, that the sum of their continuation payoffs need not be concave). An important practical implication is that the competent Agent can gain from the possibility that the Firm's belief that she is competent *might go down* during the bargaining process.

Performance can be further improved if the players have access to tests that both Agents' types can fail with positive probability, or alternatively, if parties have access to a trusted intermediary that can "noise up" the test's outcome. In fact, we prove that, with the help of such a third-party, the Agent can appropriate the maximum surplus (Theorem 3).

Our finding that selling information gradually is beneficial to the seller should (in terms of providing the highest incentives to acquire information) come as no surprise to anyone who was ever involved in consulting. The free first consultation is also reminiscent of standard business practice. The further benefits of intermediation might be more surprising. Yet it is indeed common practice to hire third parties to evaluate the value of information. This third-party structure is used as a "buffer" to ensure that the buyer does not have access to any unnecessary confidential information about the seller at any point during the sales process.²

Most of the paper analyzes the information sales problem for the specific payoff structure that arises in a simple example in which the Agent's information is decisive for a Firm's optimal decision that is explicitly modeled. However, there is nothing particular about this example. We generalize our results to arbitrary specifications of how the Firm's payoff varies with its

²We thank Rann Smorodinsky for sharing his experience in this respect. As a seller of software, the sale involved no less than three third parties specialized in this kind of intermediation –Johnson-Laird, Inc., Construx Software and NextGeneration Software Ltd.

belief about the Agent's competence. This specification could arise from decision problems that are more complicated than the binary one considered in the example. We prove that selling information in small bits is profitable as long as this payoff function is star-shaped, that is, as long as its average is increasing in the belief. Moreover, we show that, with rich enough tests, the type-1 Agent can extract the entire expected value quite generally.

The paper is related to the literature on hold-up, for example Gul (2001) and Che and Sákovics (2004). One difference is that in our game what is being sold is information and hence the value of past pieces sold depends on the realization of value of additional pieces. Moreover, we assume that there is no physical cost of selling a piece of information and hence the Agent does not care *per se* about how much information the Firm gets or what action it takes. In contrast, in Che and Sákovics (2004) each piece of the project is costly to the Agent and the problem is how to provide incentives for this observable effort rather than unobservable effort in our model. Finally, our focus is on the different ways of information transmission, which is not present in any of these papers.

The formal maximization problem, and in particular the structural constraints on information revelation, 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. There are important differences between our paper and the motivation of these papers, however. In particular, unlike in that literature, payoff-relevant actions are taken before information disclosure is over, since the Firm pays the Agent as information gets 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, as the Agent is constrained in the messages she can send by the information she actually owns. Cheap-talk (*i.e.* the possibility to send 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, before her type is known.

In our case, this is trivial and does not deliver differential payoffs to the Agent's types (*i.e.*, a higher payoff to the competent type).

The martingale property is distinctive of information, and this is a key difference between our set-up and other models in which gradualism appears. In particular, the benefits of gradualism are well known in games of public goods provision (see Admati and Perry, 1991, Compte and Jehiel, 2004 and Marx and Matthews, 2000). Contributions are costly in these games, whereas information disclosure is not costly *per se*. In fact, costlessness is a second hallmark of information disclosure that plays an important role in the analysis. (On the other hand, the specific order of moves is irrelevant for the results, unlike in contribution games.) The opportunity cost of giving information away is a function of the equilibrium to be played. So, unlike in public goods game, the marginal (opportunity) cost of information is endogenous. Relative to sales of private goods, the marginal value of information cannot be ascertained without considering the information as a whole, very much as for public goods.

Our focus (proving one's knowledge) and instrument (tests that imperfectly discriminate for it) are reminiscent of the literature on zero-knowledge proofs, which also stresses the benefits of repeating such tests. This literature that starts with the paper of Goldwasser, Micali and Rackoff (1985) is too large to survey here. A key difference is that, in that literature, passing a test conveys information about the type without revealing anything valuable (factoring large numbers into primes does not help the tester factoring numbers himself). In many economic applications, however, it is hard to convince the buyer that the seller has information without giving away some of it, which is costly –as it is in our model.

Indeed, unlike in public goods games, or zero-knowledge proofs, splitting information is not always optimal. As mentioned, this hinges on a (commonly satisfied) property of the Firm's payoff, as a function of its belief about the Agent's type.

Less related are some papers in industrial organization. Our paper is complementary to Anton and Yao (1994 and 2002) in which an inventor tries to obtain a return to his information in the absence of property rights. In Anton and Yao (1994) the inventor has the threat of revealing information to competitors of the Firm and it allows him to receive payments even after she gives the Firm all information. In Anton and Yao (2002) some contingent payments are allowed

and the inventor can use them together with competition among firms to obtain positive return to her information. In contrast, in our model, there are no contingent payments and we assume that only one Firm can use the information.

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 (2011) apply ideas similar to Aumann and Maschler (1995) to study optimal information disclosure policy when the Agent does not have private information about the state of the world, but cares about the Firm’s action.

2 The Main Example

We start with a simple example in which we explicitly model how the Firm’s value for the Agent’s information arises from a decision problem. Later sections take this value as exogenous data.

2.1 Set-Up

There is one Agent (she) and a Firm (it). The Agent is of one of two possible types: $\omega \in \Omega := \{0, 1\}$, she is either competent (1) or not (0). We also refer to these as type-1 and type-0. The Agent’s type is private information. The Firm’s prior belief that $\omega = 1$ is $p_0 \in (0, 1)$.

The game is divided into K rounds of communication (we focus on the limit as K grows large), followed by an action stage. In the action stage the Firm must choose either action I (“Invest”) or N (“Not Invest”). Not Investing yields a payoff of 0 independently of the Agent’s type. Investing yields a payoff of 1 if $\omega = 1$ and $-\gamma < 0$ if $\omega = 0$. That is, investing is optimal if the Agent is competent, as such an Agent has the know-how (or information) to make the investment thrive; however, if the Agent is incompetent, it is safer to abstain from investing.

In each of the K rounds of communication timing is as follows. First, the Firm and the Agent choose a monetary transfer to the other player, t_k^A and t_k^F , respectively.³ Second, after these

³The Reader might wonder why we allow the Agent to pay the Firm. After all, it is the Agent who owns the unique valuable good, information. Such payments will turn out to be irrelevant when only pure strategies are considered, but will play a role in the second part of the paper with more complex communication.

simultaneous transfers, the Agent chooses whether to reveal some information by undergoing a test.⁴ We propose the following concrete model of gradual persuasion/communication using tests. We assume that for every $m \in [0, 1]$, there exists a test that the competent Agent passes for sure, but that the incompetent Agent passes with probability m . The level of difficulty, m , is chosen by the Agent and observed by the Firm.⁵ If the Firm's prior belief about the Agent being competent is p and a test of difficulty m is chosen and passed, the posterior belief is

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

Thus, the range of possible posterior beliefs as m varies is $[p, 1]$ (if the test is passed). An uninformative test corresponds to the case $m = 1$. If the Agent fails the test, then the Firm correctly updates its belief to zero. To allow for rich communication, tests of any desired precision m are available at each of the K rounds, and their outcomes are conditionally independent.

In words, by disclosing information, the Agent affects the Firm's belief that she is competent. Persuasion can be a gradual process: after the Agent discloses some information, the Firm's posterior belief p' can be arbitrary, provided the prior belief p is not degenerate. But the Firm uses Bayesian updating. Viewed as a stochastic process whose realization depends on the disclosed information, the sequence of posterior beliefs is a martingale from the Firm's point of view.⁶

For now, we do not allow the competent Agent to flunk the test on purpose, nor do we consider tests so difficult that even the competent Agent might fail them. Describing strategies in terms of martingales (see footnote 6), this means that we restrict attention to processes whose sample paths are either non-decreasing, or absorbed at zero. We discuss these richer communication

⁴Nothing hinges on this timing. Payments could be made sequentially rather than simultaneously, and occur after rather than before the test is taken.

⁵Because the level of difficulty is determined in equilibrium, it does not matter that the Agent chooses it rather than the Firm.

⁶More abstractly, as in the literature on repeated games with incomplete information, we can think of an Agent's strategy as a choice of a martingale –the Firm's beliefs– given the consistency requirements imposed by Bayes' rule. For concreteness, we model this as the outcome of a series of tests whose difficulty can be varied. Alternatively, we may think of information as being divisible, and the Agent choosing how much information to disclose at each round; the incompetent Agent might be a charlatan who might be lucky or skilled enough to produce some persuasive evidence, as in the Great Oil Sniffer Hoax. While the particular implementation of information transmission is irrelevant for the purpose of equilibrium analysis, it is somewhat simpler to describe the game using tests, as it does not require a continuum of types –charlatans of varying skill levels.

possibilities in the second part of the paper.

2.2 Histories, Strategies and Payoffs

More formally, a (public) history of length k is a sequence

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

where $(t_{k'}^A, t_{k'}^F, m_{k'}, r_{k'}) \in \mathbb{R}_+^2 \times [0, 1] \times \{0, 1\}$. Here, m_k is the difficulty of the test chosen by the Agent in stage k and r_k is the outcome of that test (which is either positive, 1, or negative, 0). The set of all such histories is denoted H_k (set $H_0 := \emptyset$).

A (behavior) strategy σ^F for the Firm is a collection $(\{\tau_k^F\}_{k=0}^{K-1}, \alpha^F)$, where (i) 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 (public) history so far, as well as (ii) an action (a probability transition as well), $\alpha^F : H_K \rightarrow \{I, N\}$ after the K -th round. A (behavior) strategy σ^A for the Agent is a collection $\{\tau_k^A, \mu_k^A\}_{k=0}^{K-1}$, where (i) $\tau_k^A : \Omega \times H_k \rightarrow \mathbb{R}_+$ 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, (ii) $\mu_k^A : \Omega \times H_k \times \mathbb{R}_+^2 \rightarrow [0, 1]$ is a probability transition specifying the difficulty of the test (*i.e.*, the value of m), as a function of the Agent's type, the history up to the current round, and the transfers that were made in the round. All choices are possibly randomized.

These definitions imply that there is no commitment on either side: the Firm (and the Agent) can stop making payments at any time, and nothing compels the Agent to disclose information if she prefers not to.

In terms of payoffs, we assume there is neither discounting nor any other type of frictions during the K rounds (for example, taking the tests is free). Absent any additional information revelation, the Firm's optimal action is to invest if and only if its belief p that her type is 1 satisfies

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

Hence, its payoff from the optimal action is given by

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

where $x^+ := \max\{0, x\}$. While our analysis covers both the case in which the prior belief p_0 is below or above p^* , we have in mind the more interesting case in which p_0 is smaller than p^* . The payoff $w(p)$ is the Firm's *outside option*. Here, given the motivating investment decision, the specific outside option reduces to a call option. This is the distinguishing feature of our main example. In our general results we cover a richer class of outside option specifications.

The Agent does not value the knowledge that she potentially holds, nor does she care about the Firm's investment decision. All she cares about is getting paid. The Firm cares about the payoff from the investment decision, net of any payments to the Agent. Given some final history h_K (which does not include the Firm's final action to invest or not), type- ω Agent's realized payoff is the sum of all net transfers over all rounds, independently of her type:

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

Given type ω , the Firm's overall payoff results from its action, as well as from the sum of net transfers. If the Firm chooses the safe action, it gets

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

If instead the Firm decides to invest, it 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 and a strategy profile $\sigma := (\sigma^F, \sigma^A)$ define a distribution over $\Omega \times H_K \times \{I, N\}$, 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, V_1 , for the payoff to the Agent, when we condition on the type $\omega = 0, 1$.

The solution concept is perfect Bayesian equilibrium, as defined in Fudenberg and Tirole (1991, Definition 8.2).⁷ We assume that players have access to a public randomization device at the beginning of each round (a draw from a uniform distribution), as this facilitates an argument in a proof. The best equilibrium that we identify (whether in pure or mixed strategies, or with a mediator) turns out not to take advantage of this device, so that results do not depend on it.

A central assumption of our model is that information revealed by successful tests is valuable: if the Firm decides to stop making payments once its belief reaches some level p , its expected payoff is given by $w(p)$ and that is increasing (at least over some range) in p . There are cases in which tests can be conducted whose outcome reveals nothing valuable. (As is the case with “zero-knowledge proofs” which could be captured in our model by taking $w(p)$ equal to 0 for all $p < 1$ and $w(1) = 1$.) In practice, however, it is difficult to think of demonstrations (blueprints, prototypes, etc.) that do not involve some valuable information leakage. A competent Agent might still be useful to the Firm after she produces sufficient evidence to convince the Firm that she is competent. There is no problem for her in getting compensated for the value that she might retain in this way; our interest is in the value that she might give away in the process of convincing the Firm. For example, while proving that the Agent knows how to solve a particular problem may not give the Firm the full solution, it would give the Firm confidence that the problem is solvable and even some hints about the direction it would need to follow to find the solution on its own.

2.3 Preliminary Remarks

This game admits a plethora of equilibria, but our focus is on identifying the best equilibrium for the competent Agent. It is not difficult to motivate our interest in this equilibrium. After all, rewarding agents for their expertise is socially desirable if acquiring this knowledge is costly.

⁷Fudenberg and Tirole define perfect Bayesian equilibria for finite multistage games with observed actions only. Here instead, both the type space and the action sets are infinite. The natural generalization of their definition is straightforward and omitted.

Clearly, there are many ways to model this information acquisition stage prior to the game. Our working paper develops a particular example. As such, the equilibria we construct can be interpreted as the best relational contracts for the competent Agent subject to different restrictions on the communication technology and self-enforcing constraints.

As usual, how good payoffs can be sustained on the equilibrium path depends on the worst punishment payoffs that are consistent with a continuation equilibrium. In our game, after every history there is a “babbling” equilibrium in which the Agent never undergoes a test (*i.e.*, chooses $m = 1$ in each period), and neither the Agent nor the Firm make payments. This gives the Agent a payoff of 0, and the Firm a payoff of $w(p)$, its outside option. This equilibrium achieves the lower bound on the payoffs of all the participants simultaneously, so it is the most potent punishment available. This implies that it is without loss of generality that we can restrict attention to equilibria in which any observable deviation triggers reversion to this equilibrium (the Firm getting then its outside option $w(p)$ given its belief once the deviation occurs). To induce compliance, it suffices to make sure that all players receive at least their minmax payoff (0 and $w(p)$) at any time.

If the Firm assigns probability p to the type-1 Agent, then, from its point of view, the expected total surplus is at most $p \cdot 1 + (1 - p) \cdot 0 = p$ (this is in the best possible scenario in which it eventually takes the right investment decision). Hence, given some equilibrium, any history h_k and resulting belief p , continuation payoffs must satisfy

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

With only one round of communication, $K = 1$, both types of the Agent have to receive the same payoff in any equilibrium so $V_1 \leq p - w(p)$ in this case. By (1), $p - w(p)$ is also the upper bound on the average, or *ex ante* payoff of the Agent.

How much can gradual communication improve V_1 ? By (1), given that $W(h_k) \geq w(p)$ and $V_0(h_k) \geq 0$, the type-1 Agent cannot receive more than $1 - w(p)/p$. Clearly $p - w(p) < 1 - w(p)/p$ whenever $w(p) < p$, so the upper bound is strictly larger than the maximum *ex ante* payoff. Can we improve on the latter?

Before we present the analysis that answers this question, we make two observations.

First, it is worth pointing out that, in some cases, maximizing the incentives to acquire information is not about maximizing the type-1 Agent's payoff V_1 , but the difference $V_1 - V_0$. But the two objectives coincide. This can be seen in three steps: first, in terms of the Agent's equilibrium payoffs (V_0, V_1) , there is no loss of generality in assuming that the equilibrium achieving this payoff is *efficient*, *i.e.*, that it satisfies (1) with equality: disclosing the type in the last period on the equilibrium path does not affect the Agent's payoff and only makes compliance with the equilibrium strategy more attractive to the Firm. Second, if (1) holds as an equality, then

$$V_1 - V_0 = \frac{V_1 + W - p_0}{1 - p_0}.$$

Hence, maximizing the difference in the types' payoffs amounts to maximizing the sum $V_1 + W$. Third, maximizing V_1 is equivalent to maximizing $V_1 + W$. This is because payoffs between the principal and the Agent can be transferred one-to-one via the first payment that the Firm makes: if $W > w(p_0)$, we can decrease W and increase V_1 by the same amount by requiring the Firm to make a larger payment upfront. Hence, in maximizing $V_1 + W$ over all equilibria, there is no loss in assuming that $W = w(p_0)$, a fixed quantity, and so in maximizing V_1 only.

Second, while we characterize the equilibrium maximizing V_1 (or $V_1 - V_0$), one may be interested in other equilibrium payoffs. A partial characterization is as follows. The Firm cannot hope for more than p , the entire surplus, and there is a trivial equilibrium that achieves this upper bound: on path, no payment is ever made, and the Agent reveals her type in the last period ($m = 0$). There is an equally simple equilibrium that achieves the maximum *ex ante* payoff of the Agent, $pV_1 + (1 - p)V_0$: the Firm is expected to pay the difference $p_0 - w(p_0)$ in the initial period, and the agent reveals her type ($m = 0$) if and only if the Firm makes this payment. In this game, the strategy that maximizes this *ex ante* payoff of the informed player is trivial, unlike in standard games with incomplete information (see Aumann and Maschler, 1995). Because the type-1 Agent can always mimic the type-0 Agent in the choice of m (and once a test is failed Agent's payoff is 0), the competent type payoff must be at least as large as the incompetent's. This implies that the best equilibrium for the type-0 Agent maximizes the Agent's *ex ante* payoff.

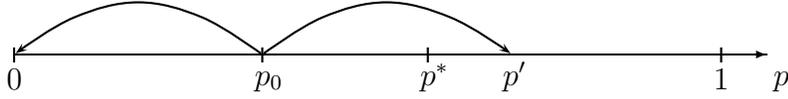


Figure 1: A feasible action

3 Pure Strategies

We now turn to the focus of the analysis: what equilibrium maximizes the payoff of the type-1 Agent, and how much of the surplus can she appropriate? Note that this maximum payoff is non-decreasing in K , the number of rounds: players can always choose not to make transfers or disclose any information in the first round. Hence, for any p_0 , the highest equilibrium payoff for the type-1 Agent has a well-defined limit as $K \rightarrow \infty$ that we seek to identify.

In this section we assume that the type-1 Agent always passes any test she takes. She is not allowed to “flunk” the test on purpose, a possibility that we will allow in Section 4: there we enrich the description of the game to allow the agent to choose whether to pass the test after she chooses the difficulty m . In that richer game the analysis in this section is equivalent to restricting attention to pure strategies.

From the Firm’s point of view, its posterior will take one of two values: either it jumps from p_0 up to some p' , if the test is successful. Or it jumps down to zero. This is illustrated in Figure 1. The two arrows indicate the two possible posterior beliefs. As mentioned, viewed from the Firm’s perspective, this belief must follow a martingale: the Firm’s expectation of its posterior belief must be equal to its prior belief. This is not the case from the Agent’s point of view. Given her knowledge of the type, she assigns different probabilities to these posterior beliefs than the Firm. If she is competent, she knows for sure that the belief will not decrease over time. If she is incompetent, the expectation of the posterior belief is below p_0 , as she does not know whether she will be lucky in taking the test (the process is then a supermartingale).

Instead of describing the information part of an equilibrium outcome by the tests taken so far

$\{m_k\}$ and their results, we can equivalently describe it by *martingale splitting*, *i.e.* the sequence of Firm's beliefs that the type is 1, conditional on all tests so far. As long as the Agent passes the tests, the Firm's equilibrium beliefs follow a non-decreasing sequence $\{p_0, \dots, p_{K+1}\}$ starting at the Firm's prior belief, p_0 , and ending at $p_{K+1} = 1$ (assuming, without loss, that the equilibrium is efficient). If the Agent fails a test, the posterior drops to zero.

An equilibrium must also specify payments. It turns out that the type-1 Agent's payoff decreases if the Firm is ever granted any payoff in excess of its outside option. On the one hand, the Agent could demand more in earlier rounds by promising to leave some surplus to the Firm in later rounds. On the other hand, the willingness-to-pay of the Firm for this future surplus is lower than the cost of such a promise to the type-1 Agent. The reason is that the Firm assigns a lower probability than the type-1 Agent to the posterior increasing (and payments once the posterior drops to zero are not individually rational). Finally, it is not hard to see that there is no point here in having the Agent make any payments. To sum up, if the Firm's belief in the next round is either p_{k+1} or 0, given the current belief p_k , then the equilibrium specifies that the Firm pays her willingness-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} , and $\mathbb{E}_F[\cdot]$ is the expectation operator for the Firm.

This leaves us with the determination of the sequence of posterior beliefs.⁸

3.1 The Main Example: A Geometric Analysis

We already know that it is possible for the Agent to appropriate some of the value of her information, but the question is whether she can get more than $p_0 - w(p_0)$, which is just as much as the type-0 Agent gets in the equilibrium we constructed so far.

Consider first the case $K = 1$. In this case, the highest equilibrium payoff to the type-1 Agent is indeed equal to $p_0 - w(p_0)$. Suppose that a successful test takes the posterior to $p_1 \geq p_0$. Using

⁸In addition, an equilibrium must also specify how players behave off the equilibrium path. As discussed, the most effective punishment for deviations is reversion to the babbling equilibrium, and this is assumed throughout.

the martingale property, it must be that the probability that the posterior is p_1 is p_0/p_1 , because

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

hence, the Firm is willing to pay

$$\mathbb{E}_F[w(p')] - w(p_0) = \frac{p_0}{p_1}w(p_1) - w(p_0) \leq p_0 - w(p_0),$$

where the inequality follows from $w(p_1) \leq p_1$. Setting p_1 to 1 is best, as it makes the inequality tight: with one round, revealing all information is optimal.

Note that, when $p_0 \leq p^*$, $w(p_0) = 0$, and the highest payoff in one round that the type-1 Agent can achieve is the prior p_0 , which is increasing in $p_0 \leq p^*$. This suggests one way to improve on the payoff with two rounds. In the first round, the Agent takes a test whose success leads to a posterior of p^* *for free*. Indeed, the Firm is not willing to pay for a test that does not affect its outside option. In the second round, the equilibrium of the one-round game is played, given the belief p^* . This second and only payment yields

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

This is illustrated in the right panel of Figure 2. The lower kinked line is the outside option w , the upper straight line is total surplus, p . Hence, the payment in the second round is given by the length of the vertical segment at p^* in the right panel, which is larger than the payment with only one round, given by the length of the vertical segment at p_0 .

To sum up: the Agent gives away a chunk of information for free, making the Firm really unsure whether investing is a good idea. Then she charges as much as she can for the disclosure of all her information.

Is the splitting that we described optimal with two periods to go? As it turns out, it is so if and only if $p_0 < (p^*)^2$. But 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 $p_0 > p^*$. The optimal strategy is given at the end of this

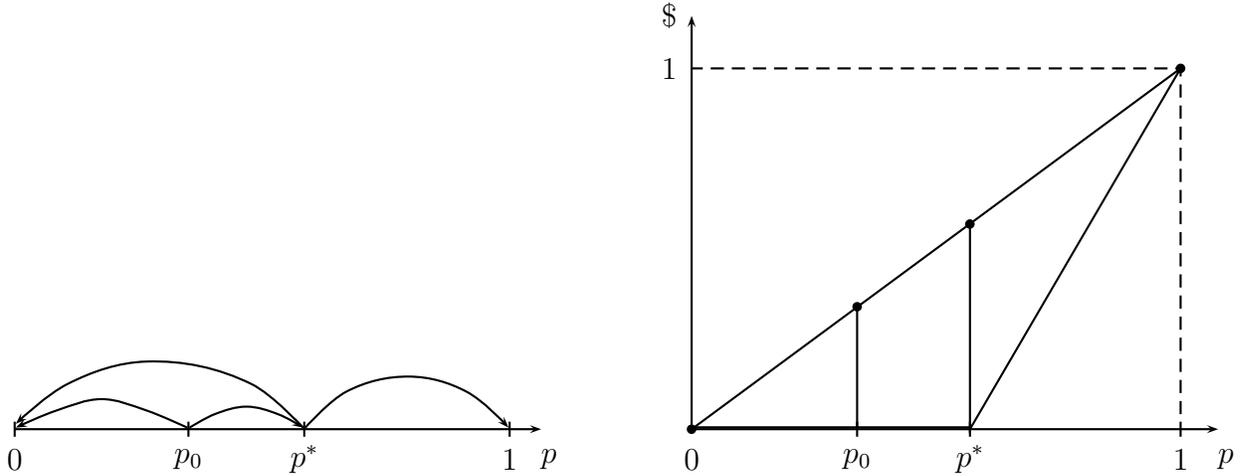


Figure 2: Revealing information in two steps

subsection.

Can we do better with more rounds? Consider Figure 3. As shown on the left panel, information is revealed in three stages. 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. The two solid (red) segments represent the maximum payments that can be demanded at the second and third stage. (No payment is made in the first, as the splitting does not affect the Firm's outside option.) Their added lengths measures the type-1 Agent's payoff. Compare with our two-stage equilibrium, in which all information is disclosed once the belief reaches p^* : the payment is equal to the length 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 red segment, plus the dotted segment). The payoff with three stages is larger, as the chords from the origin to the point $(p, w(p))$ become steeper as p increases. We could go on: information splitting is beneficial. Figure 4 illustrates the total payoff that results from a splitting that involves many small steps (which is the sum of all vertical segments).

Does it follow that the competent Agent extracts the maximum value of information as $K \rightarrow \infty$? Unfortunately, no: see the right panel of Figure 4. As the Firm's belief goes from $p - dp$ to p , the Firm must pay (using the martingale property, the test must be passed with

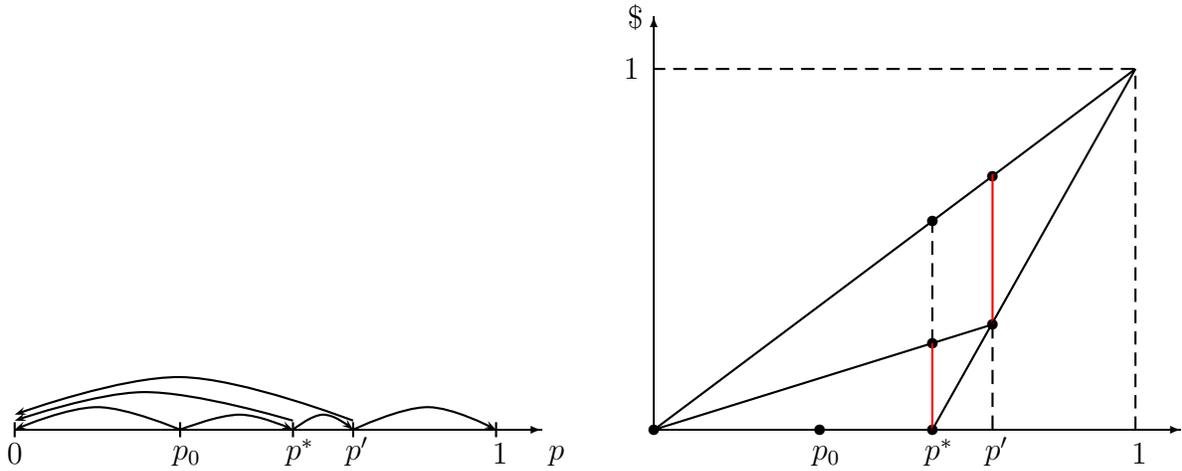


Figure 3: Revealing information in three steps: evolution (left) and payoff (right)

probability $\frac{p-dp}{p}$)

$$\frac{p-dp}{p}w(p) - w(p-dp),$$

yet its outside increases by $w(p) - w(p-dp)$. The type-1 Agent gives up the difference $w(p)dp/p$ in the process. This foregone payoff 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 involved. Note that this foregone payoff does not benefit the Firm, which is always charged its full willingness-to-pay. The type-0 Agent reaps this payoff. As a result, her payoff does not vanish, even as $K \rightarrow \infty$.

What does the maximum payoff converge to as $K \rightarrow \infty$? Plugging in the specific form of w from our leading example, the payment for a splitting of p into $p' \in \{0, p+dp\}$ is

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

where $O(x) < M|x|$ for some constant M and all x . If the entire interval $[p^*, 1]$ is divided in this

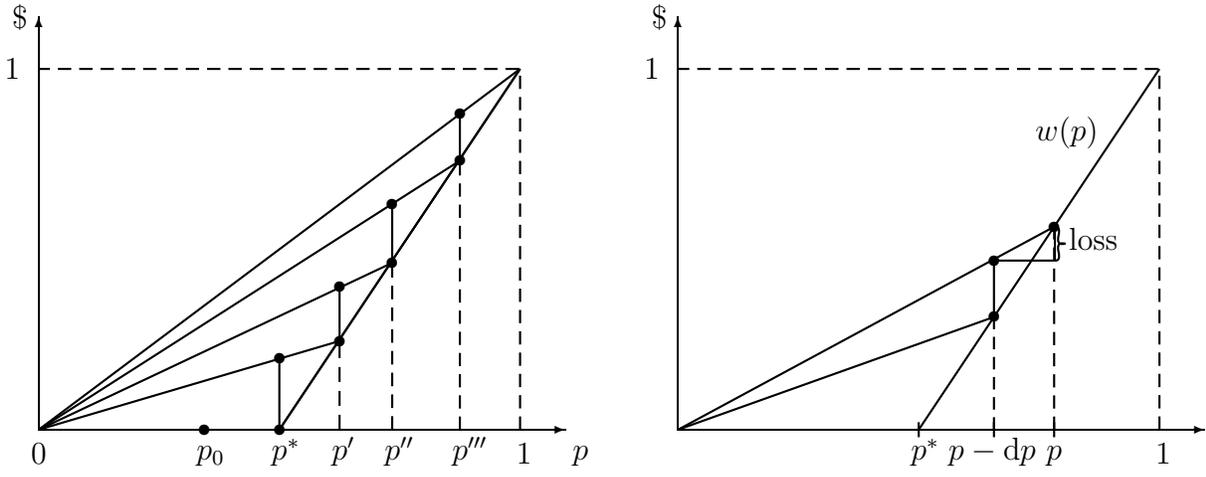


Figure 4: Revealing information in many steps (left); Foregone profit at each step (right)

fashion into smaller and smaller intervals, the resulting payoff to the competent Agent tends to

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

This suggests that the limiting payoff is independent of the exact way in which information (above p^*) is divided up over time, as long as the mesh of the partition tends to zero.

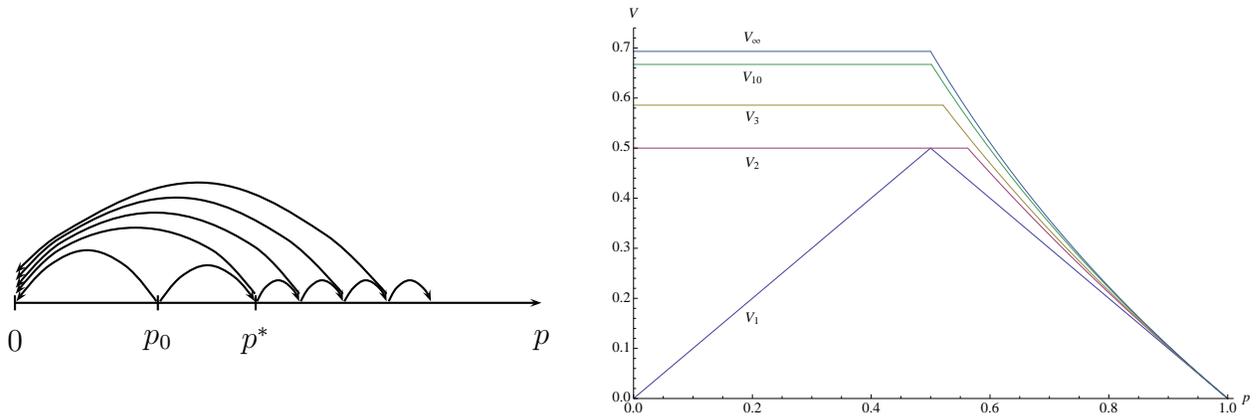


Figure 5: Revealing information in many steps (left); Payoff as a function of K (right).

Lemma 1 *As $K \rightarrow \infty$, the maximum payoff to the type-1 Agent in pure strategies tends to, for $p_0 < p^*$,*

$$V_1^p(p_0) := -\gamma \ln p^*.$$

This lemma follows as immediate corollary from the next one. Note that this payoff is independent of p_0 (for $p_0 < p^*$). Indeed, the first chunk of information, leading to a posterior belief of p^* if K is large enough, is given away for free. It does not affect the Firm's outside option, but it makes the Firm as unsure as can be about the right decision. From that point on, the Agent starts selling information in excruciatingly small bits, leaving no surplus whatsoever to the Firm, as in the left panel of Figure 5.

We conclude this subsection by the explicit description of the equilibrium that achieves the maximum payoff of the type-1 Agent, as a function of the number of rounds and the prior belief p_0 . Here, $(x)^- := -\min\{0, x\} \geq 0$.

Lemma 2 *The maximum equilibrium payoff of the type-1 Agent in a game with K rounds is recursively given 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 leading to a posterior belief 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 .

Note that, fixing $p_0 < p^*$, and letting $K \rightarrow \infty$, it holds that $p_0 < (p^*)^{\frac{K}{K-1}}$ for all K large enough, so that, with enough rounds ahead, it is optimal to set $p_1 = p^*$ in the first, and then to follow the sequence of posterior beliefs $(p^*)^{\frac{K-1}{K}}, (p^*)^{\frac{K-2}{K}}, \dots, 1$, and the sequence of posteriors

successively used becomes dense in $[p^*, 1]$. Therefore, with sufficiently many rounds, 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 Firm's belief gradually up all the way to one. The right panel of Figure 5 shows how the payoff varies with K .

Note also that, for any K and any equilibrium, if p and $p' > p$ are beliefs on the equilibrium path, then $V_0(p') - V_1(p') \leq V_0(p) - V_1(p)$, as long as only the Firm makes payments. Indeed, going from p to p' , the type-1 Agent forfeits the payments that the Firm might have made over this range of beliefs (hence $V_1(p') < V_1(p)$), while the type-0 Agent only forfeits them in the event that she is able to pass the test: hence she loses less, and might even gain (for instance, she might not have been able to pass the first free test at $p < p^*$). As a result, the type-1 Agent has a preference for lower beliefs, relative to the type-0 Agent. Having to give away information is more costly to an Agent who knows that she owns it. This plays an important role once mixed strategies are considered.

3.2 General Outside Options

Assuming that the outside option is given by a call option, as in our main example, leads to closed-form expressions. However, the analysis can be generalized. Suppose that the payoff of the Firm (gross of any transfers) as a function of its posterior belief p after the K rounds is a non-decreasing continuous function $w(p)$, and normalize $w(0) = 0$, $w(1) = 1$. We further assume that $w(p) \leq p$, for all $p \in [0, 1]$, for otherwise full information disclosure is not socially desirable. 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 since it is a primitive here, we do not assume so.

Recall that the best equilibrium with many rounds called for a first burst of information released for free (assuming $p < p^*$), after which information is disclosed in dribs and drabs. One might wonder whether this is a general phenomenon.

The answer, as it turns out, depends on the shape of the outside option. It is 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 (strictly) *star-shaped*, *i.e.*, if and only if the average, $w(p)/p$, is a strictly increasing function of p .⁹ More generally, if a function is star-shaped on some intervals of beliefs, but not on others, then information will be sold in small bits at a positive price for beliefs in the former type of interval, and given away for free as a chunk in the latter. In our main example, w is not star-shaped on $[0, p^*]$, as the average value $w(p)/p$ is constant (and equal to zero) over this interval. However, it is star-shaped on $[p^*, 1]$. Hence our finding.

Let us first consider a star-shaped outside option. 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.¹⁰ Given the Firm's prior belief p_0 , the type-1 Agent's payoff becomes then (in the limit, 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 $w(p) = (p - (1 - p)\gamma)^+$.¹¹ 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 splitting information as finely as possible is best in that case, fix some arbitrary interval of beliefs $[\underline{p}, \bar{p}]$, and consider the alternative strategy under which the posterior belief of the Firm jump from \underline{p} to \bar{p} , the payment from the Firm to the Agent in that round is given by

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

⁹This condition already appears in the economics literature in the study of risk (see Landsberger and Meilijson, 1990). It is weaker than convexity: the function $p \mapsto p^\alpha$ is star-shaped for $\alpha > 1$, but only convex for $\alpha \geq 2$.

¹⁰In case $w(p)$ is not differentiable, then $w'(p)$ is the right-derivative, which is well-defined in case w is star-shaped.

¹¹In our main example, w is (globally) *weakly* star-shaped: that is, the function $p \mapsto w(p)/p$ is only weakly increasing. The formula for the maximum payoff in the limit $K \rightarrow \infty$ is the same whether there is a jump in the first period or not. But for any finite K , splitting information disclosures over the range $[p_0, p^*]$ is suboptimal, as it is a "wasted period," whose cost only vanishes in the limit.

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

$$\frac{1}{\bar{p} - \underline{p}} \int_{\underline{p}}^{\bar{p}} \frac{w(p)}{p} dp \leq \frac{w(\bar{p})}{\bar{p}}, \quad (2)$$

which is satisfied if the average $w(p)/p$ is increasing.

Equation (2) also explains why splitting information finely is not a good idea if the average outside option is strictly decreasing over some range $[\underline{p}, \bar{p}]$, as the inequality is reversed in that case. What determines the jump? As mentioned, the payoff from a jump is $\underline{p}w(\bar{p})/\bar{p} - w(\underline{p})$, while the marginal benefit from finely splitting information disclosures at any given belief p (in particular, at \bar{p} and \underline{p}) is $w'(p) - w(p)/p$. 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 6. 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, since $w(p)/p$ 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 $p < p^*$.

There is a simple way to describe the maximum resulting payoff. Given a non-negative function f on $[0, 1]$, let

$$\text{sha } f$$

denote the largest weakly star-shaped function that is smaller than f . In light of the previous discussion (see right panel of Figure 6), the following result should not be too unexpected.

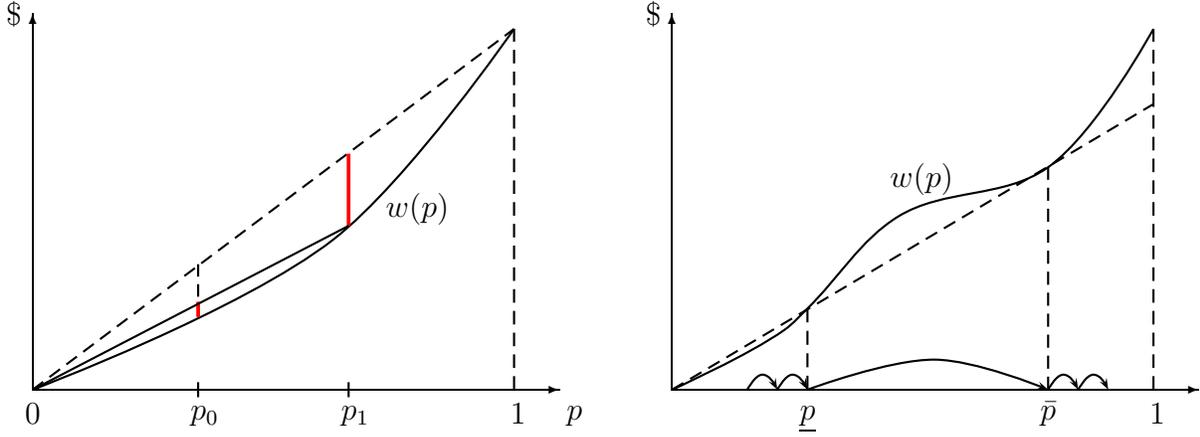


Figure 6: Splitting information with an arbitrary outside option

Theorem 1 *The maximum equilibrium payoff to the type-1 Agent in pure strategies tends to, as $K \rightarrow \infty$,*

$$V_1^p(p_0) = 1 - sha w(\hat{p}_0) - \int_{\hat{p}_0}^1 sha w(p) dp/p,$$

where $\hat{p}_0 := \min \{p \in [p_0, 1] : w(p) = sha w(p)\}$.

That is, the same formula as in the case of a star-shaped function applies, provided one applies it to the largest weakly star-shaped function that is smaller than w . In words, the maximum payoff to the type-1 Agent is the area between the marginal and the average outside option of the Firm, after “regularizing” this outside option by considering the largest weakly star-shaped function below it.¹²

The proof also elucidates the structure of the optimal information disclosure policy, at least in the limit. Let

$$I_w := \text{cl} \{p \in [0, 1] : sha w(p) = w(p) \text{ and } w(p)/p \text{ is strictly increasing at } p\}.$$

¹²There is an obvious analogy with results in the literature in marginal pricing, in which, not surprisingly, star-shaped functions play an important role as well.

In our main example, $w(p) = w(p)$ for all p , but $I_w = [1/2, 1]$. Then the set of on-path beliefs as $K \rightarrow \infty$ held by the firm is contained, and dense, in I_w if $I_w \neq \emptyset$. If $I_w = \emptyset$, any policy is optimal.

Note that this result immediately implies that the highest payoff to the type-1 Agent is higher, the lower the outside option w . That is, if we consider two functions w, \tilde{w} such that $w \geq \tilde{w}$, then the corresponding payoffs satisfy $V_1^p \leq \tilde{V}_1^p$. The “favorite” outside option for the Agent is $w(p) = 0$ for all $p < 1$, and $w(1) = 1$ (though this does not quite satisfy our maintained continuity assumptions). In that case, the type-1 Agent appropriates the entire surplus. This is the case considered in the literature on “zero-knowledge proofs:” the revision in the Firm’s belief that successive information disclosures entail does not affect its willingness-to-pay.

4 Mixed Strategies and Mediation

So far, we have assumed that the competent Agent always passes the test, which implies that the Firm’s posterior belief is either non-decreasing, or absorbed at zero.

There are two reasons why even the competent Agent may fail. First, she may be able to choose to flunk the test (it turns out that such option may improve upon the equilibria considered so far). In practice, it is hard to see what prevents an Agent from failing intentionally a given test: software can be crippled or slowed down, prototypes can be damaged or impaired, imprecise or even incorrect answers can be given. To model this possibility, we add a third dimension to the Agent’s strategy; namely, in every round, after a test has been privately performed, the Agent has the choice, in case of a success, to report a failure. As further notation is not needed, we refer the interested Reader to the working paper for a formal definition. Because the model considered in Section 3 corresponds to the special case in which the competent Agent always passes the test—the only interesting pure strategy in the extended model—we refer to this version as the model with mixed strategies.

A second reason for why a competent Agent might fail a test is simply that the test might be noisy, or very hard. One might devise procedures that are so difficult that even knowledgeable agents might be occasionally unsuccessful; not many recognized experts provide correct

predictions every time.

There is an important difference between these two cases. In the first case, a competent Agent who fails the test must be willing to fail. In the second case, she might just not be able to pass it. Hence, in the first case, equilibrium imposes more stringent requirements than in the second. Clearly, we can model the second case by allowing for a more general technology, *i.e.*, tests that are parameterized by two probabilities, (m_0, m_1) , where m_ω is the probability with which the type- ω Agent passes the test. From a game-theoretic point of view, this is equivalent to allowing for a (disinterested) mediator in the baseline model: the competent Agent always passes the test, whose outcome is observed by the mediator, but not by the Firm. Then, the mediator chooses whether to report whether the test was successful or not to the Firm. Our description follows the second approach, and we refer to this version as the model with mediation.

While the “game-theoretic” mediator is an abstraction that does not require a third-party to be involved, but merely the necessary technology (a trustworthy noisy channel whose output depends on the outcome of the test), it is worth stressing that such intermediaries are actually being involved in sales of intellectual property. As mentioned in the introduction, there are law firms, consulting firms and specialized companies that are hired for the purpose of estimating and certifying the value of intellectual property and facilitating technological transfers.

While mixed strategies turn out to be less valuable than mediators, the fundamental principle for why lower posterior beliefs can be useful is the same in both cases. The next subsection provides an illustration.

4.1 The Value of Lower Posteriors: An Illustration

Consider the main example, in which the outside option is a simple call option, and consider $\gamma = 1$ and the limiting case $K = \infty$. Using the best pure-strategy equilibrium (for the type-1 Agent) as a benchmark, the type-1 Agent has a payoff function given by $-\ln p$ for $p > p^*$, and $-\ln p^*$ for $p \leq p^*$.

Suppose that the Firm and the Agent agree to the following (self-enforcing) scheme. If the test fails, the posterior belief falls to $p - \Delta$, for some $\Delta > 0$. If the test succeeds, the posterior belief jumps to $p + \Delta$. Pick Δ such that $p^* < p - \Delta < p + \Delta < 1$. Such posterior beliefs are

achieved by mixing by the type-1 Agent (or by a mediator on her behalf), given that the type-0 Agent will disclose that the outcome of the test is a success whenever she is lucky. Because the possible posterior beliefs are symmetric around p , the two events (that information gets disclosed or not) must be equally likely from the Firm's point of view.

The new twist is that, in the event that the posterior belief drops to $p - \Delta$, the Agent is expected to pay the Firm an amount $X > 0$. No payment is made by the Agent if the posterior belief increases to $p + \Delta$. Because both posterior beliefs are equally likely, the Firm is willing to pay $X/2$ upfront in exchange for this contingent future payment, and the equilibrium calls for the Firm to make this payment in addition to the familiar term that corresponds to the variation in its expected outside option.

Such a side-payment is neutral from the point of the view of the Firm: after all, the upfront payment is fair, given the odds that the posterior goes up or down. But it is not fair from the Agent's point of view: because the posterior belief is more likely to go down if the Agent is incompetent, by definition of the posterior belief, this implies that the incompetent Agent is more likely to have to pay back than the competent Agent. In this fashion, some payoff gets shifted from the incompetent to the competent Agent.

There are two constraints on the size of this payment X . First, it cannot exceed the continuation payoff of the type-0 Agent, for otherwise she would renege on the back payment in case she fails the test. That is, $X \leq V_0(p - \Delta)$, where V_0 is her continuation payoff. Second, in the case the mixing is performed by the (type-1) Agent, rather than by a mediator, it must be that the Agent is actually indifferent between passing or failing the test. In this case, assuming that after this payment play resumes according to the best pure strategy equilibrium described above, the continuation payoffs after this payment are $-\ln(p + \Delta)$ and $-\ln(p - \Delta)$ respectively; hence, we must set X so as to exactly offset this difference in continuation payoffs, *i.e.*, $X = \ln(p + \Delta) - \ln(p - \Delta)$. This certainly satisfies $X < V_0(p - \Delta)$ if Δ is small enough. As mentioned, because $V_0 - V_1$ (the difference in payoffs in the best equilibrium) is increasing in p , this implies that the type-0 Agent is happy to claim she passes the test whenever she is lucky. The left panel of Figure 7 illustrates how the mixing works, starting from a given belief $p > p^*$.

Given that the Firm pays $X/2$ upfront, and that, by construction, the continuation payoff of

the type-1 Agent is the same whether the posterior belief goes up or down (namely, $\ln(p - \Delta)$), her expected payoff is

$$\frac{\ln(p + \Delta) - \ln(p - \Delta)}{2} + \ln(p - \Delta) = -\frac{\ln(p + \Delta) + \ln(p - \Delta)}{2} > -\ln p,$$

where the strict inequality follows from Jensen's inequality. Hence, we have just improved on our limit payoff $V_1(p) = -\ln p$.

What is the key to this improvement, and how much can such schemes improve on the competent Agent's payoff? It turns out to depend on the curvature of the *sum* of the Firm's and competent Agent's payoffs. Let $V_0^m(p)$ and $V_1^m(p)$ denote the limiting payoffs as $K \rightarrow \infty$ in the best equilibrium that uses mixed (or pure) strategies and define $h(p) := V_1^m(p) + w(p)$. if $V_0^m(p) = 0$ for some p , the incompetent Agent would no longer make any payments; by (1), this implies that $h(p) = \bar{h}(p) := 1 - (1 - p)w(p)/p$ (\bar{h} is the bound from (1) and $V_0 \geq 0$). This would yield the highest possible payoff to the competent Agent, given the Firm's outside option. So suppose that $h < \bar{h}$ on some interval around p , and for the sake of contradiction, assume that h is not concave on this interval, *i.e.* there exists $p_1 < p < p_2$ such that

$$h(p) < \frac{p_2 - p}{p_2 - p_1}h(p_1) + \frac{p - p_1}{p_2 - p_1}h(p_2).$$

We generalize the previous scheme to this case: the agent pays $V_1^m(p_1) - V_1^m(p_2)$ to the principal if and only if the posterior drops to p_1 , and play reverts then (or if the posterior belief turns out to be p_2) to the equilibrium that achieves V_1^m . The type-1 Agent is indifferent between both posterior beliefs, and so is willing to randomize. Given her assessment of the likelihood of each of these events, the Firm is willing to pay upfront

$$\frac{p_2 - p}{p_2 - p_1}[w(p_1) + V_1^m(p_1) - V_1^m(p_2)] + \frac{p - p_1}{p_2 - p_1}w(p_2) - w(p),$$

as this is the difference between its expected continuation payoff and its current outside option. The type-1 Agent's payoff $\hat{V}_1(p)$ consists then of this payment and her continuation payoff $V_1^m(p_2)$,

so that, adding up,

$$\begin{aligned} h(p) \geq \hat{V}_1(p) + w(p) &= \frac{p_2 - p}{p_2 - p_1} [w(p_1) + V_1^m(p_1) - V_1^m(p_2)] + \frac{p - p_1}{p_2 - p_1} w(p_2) + V_1^m(p_2) \\ &= \frac{p_2 - p}{p_2 - p_1} h(p_1) + \frac{p - p_1}{p_2 - p_1} h(p_2). \end{aligned}$$

Note that the participation constraint for the incompetent Agent, $V_0^m(p_1) > V_1^m(p_1) - V_1^m(p_2)$ is always satisfied if p_1, p_2 are close enough to p and $V_0^m(p_1) > 0$, and so h must be locally concave at any p at which $V_0^m(p) > 0$.¹³

The concavity of the sum of the payoffs of the competent Agent and the Firm in the best equilibrium should not be surprising: if it were convex, a lottery could increase their joint payoff, at the expense of the incompetent Agent. The upfront payment by the Firm, followed by the contingent payment by the Agent is the simplest way of implementing such a lottery.

To summarize: using contingent payments in the way described improves the competent Agent's payoff, and this can be done as long as the type-0 Agent's payoff is not zero, and, in case the type-1 Agent is actually required to perform the randomization herself, as long as h is not locally concave. Equilibrium imposes additional constraints on the type-1 Agent's payoff, which is the subject of the next subsection.

4.2 Maximum Payoff with Mixed Strategies and Mediation

4.2.1 Mixed Strategies

First, consider the case of mixed strategies. Two constraints have been derived on the limiting value of h , the sum of the payoffs of the Firm and the competent agent. First, it must be less than $\bar{h} = 1 - (1 - p)w(p)/p$, as implied by feasibility given that the type-0 Agent's payoff is non-negative. Second, on any interval on which $h < \bar{h}$, the function h must be locally concave. There is a third constraint on h that is rather obvious: h must exceed w , the outside option of the Firm, as the type-1 Agent's payoff is non-negative.

¹³This hinges on continuity of V_1^m and V_0^m ; V_1^m is continuous because it is always possible to use the same disclosure strategy starting at p_2 as the continuation strategy given p_1 would specify from the first posterior belief above p_2 onward; the first payment must be adjusted, but the continuity in payoffs as $p_1 \rightarrow p_2$ then follows from the continuity of w . Continuity of V_0^m follows from the continuity of V_1^m .

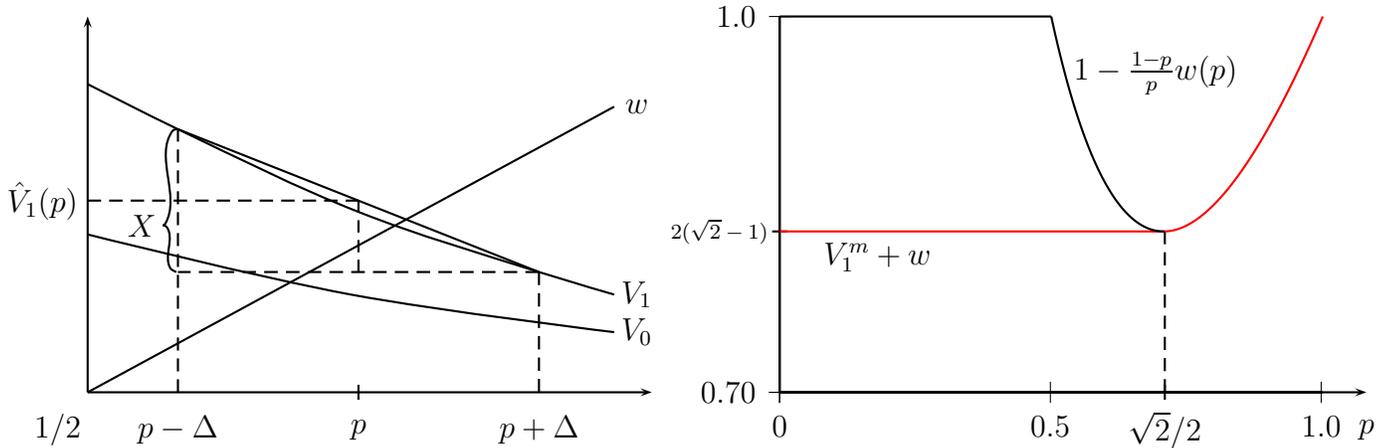


Figure 7: Construction of the scheme (left); Maximum limit payoff $V_1^m + w$, $\gamma = 1$ (right).

Finally, the basic splitting of Section 3 delivers one more restriction, namely, the function h must be no steeper than $w(p)/p$. We can always split the prior belief p_0 into the posterior beliefs in $\{0, p_1\}$, $p_1 > p_0$. The Firm is willing to pay $p_0 w(p_1)/p_1 - w(p_0)$ for such a test, so that, at the very least,

$$V_1^m(p_0) \geq \frac{p_0}{p_1} w(p_1) - w(p_0) + V_1^m(p_1),$$

or

$$\frac{h(p_1) - h(p_0)}{p_1 - p_0} \leq \frac{w(p_1)}{p_1}. \quad (3)$$

If h were known to be differentiable, this would reduce to the requirement that $h'(p)$ be smaller than $w(p)/p$. More generally, chords connecting points $(p_0, h(p_0))$ and $(p_1, h(p_1))$ must be flatter than the ray with slope $w(p_1)/p_1$.

As it turns out, equilibrium imposes no additional restriction on h , as we show in Appendix.¹⁴ What is the smallest function that satisfies these four requirements?¹⁵ In our main example, some

¹⁴Roughly, any function satisfying these properties cannot be improved upon with one more round, even with mixed strategies. Because the payoff of the type-1 Agent is increasing in her continuation payoff, this means that the highest limiting payoff must be below this function. Conversely, the limiting payoff must satisfy these properties. Hence, it follows that this lowest function is the limiting payoff.

¹⁵One might wonder why the *smallest* function h satisfying the requirements is the appropriate one; this is because, starting from the highest equilibrium payoff with one round, and applying the two schemes that we have

algebra gives that

$$V_1^m(p) = \begin{cases} 2\sqrt{\gamma}(\sqrt{1+\gamma} - \sqrt{\gamma}) - w(p) & \text{if } p < p^m := \sqrt{p^*}, \\ 1 - w(p)/p & \text{if } p \geq p^m. \end{cases}$$

The smallest function h^m is shown on the right panel of Figure 7 in the case $\gamma = 1$. The following corollary records the limiting value for prior beliefs below p^* .

Lemma 3 *As $K \rightarrow \infty$, the maximum payoff to the type-1 Agent in mixed strategies tends to, for $p_0 < p^*$,*

$$V_1^m(p_0) = 2\sqrt{\gamma}(\sqrt{1+\gamma} - \sqrt{\gamma}) < 1.$$

That is, full extraction occurs for high enough ($p \geq p^m$, in which case $V_0^m(p) = 0$) but not for low beliefs. Still, even for $p < p^m$, this is a marked improvement upon pure strategies. Because the competent Agent gains from mixed strategies, and the Firm does not lose from them, it must be that the type-0 Agent loses. For $p \leq p^m$, her payoff function is given by $V_0(p) = 1 + (p \ln p)/(1 - p)$ (for $\gamma = 1$).

How about more general outside options? The logic is robust: let h^m be the smallest function satisfying the four requirements above (which is well-defined, as the lower envelope of functions satisfying the requirements satisfies them as well). The following theorem elucidates the role of h^m .

Theorem 2 *Assume that w is weakly star-shaped. As $K \rightarrow \infty$, the maximum payoff to the type-1 Agent in mixed strategies tends to:*

$$V_1^m(p_0) = h^m(p_0) - w(p_0).$$

To emphasize, the result does not *assume* that only tests or schemes that we have described so far can be used. It shows that, at least as the number of rounds is sufficiently large, these suffice.

described, we recursively obtain higher values for h as the number of rounds increases, but we cannot “overtake” the smallest function that satisfies the four requirements.

4.2.2 Mediation

It turns out that a similar reasoning can be used to characterize the maximal V_1 in case a mediator can send noisy messages based on the test results (or the Agent has access to noisy tests). The only difference is that the scheme that involves payments by the Agent in case the posterior drops is no longer constrained by the indifference of the competent Agent, which imposed that h was locally concave whenever h fell short of the upper bound \bar{h} . So we are left with the other three restrictions on the function h . It turns out, as before, that the solution is given by the smallest function satisfying these requirements. The maximum payoff has a particularly simple expression, and the result does not require w to be star-shaped.

As the next theorem states, the type-1 Agent can extract all the surplus from the type-0 Agent as well as the Firm, up to its outside option.

Theorem 3 *As $K \rightarrow \infty$, the maximum payoff to the type-1 Agent with an intermediary tends to:*

$$V_1^{int}(p_0) = 1 - \frac{w(p_0)}{p_0}.$$

In our main example, this means that, for $p_0 < p^*$, the maximum payoff of the competent Agent is 1 –and there is nothing left to improve upon.

5 Final Remarks

This paper describes self-enforcing contracts based on gradual persuasion to facilitate sale of information. Clearly, in real-life applications, the mechanism that we describe is limited by the extent to which information is divisible, or tests are available. On the other hand, it can be facilitated by repeated interactions and reputation-building. Although it may not appear that way at first glance, we claim that the logic of our result is quite robust in several dimensions, discussed below.

1. We have assumed that taking a test entails no cost to the Agent. If tests are costly, our results would continue to hold as long as taking a test is contractible. Otherwise, the

standard hold-up logic applies: in the last round, the Agent would not take the test. Hence the Firm would not pay, and the equilibrium would unravel. Che and Sákovics (2004) suggest the following solution: if the equilibrium concept is epsilon-equilibrium and easier tests are cheaper (for example, if we interpret harder tests as taking many easier tests at once), then gradualism would restore the desired outcome. Here, gradualism helps resolve the problem of the Agent holding up the Firm, while our paper shows how gradualism helps resolving the opposite hold-up.

2. Suppose that the Agent cares to some extent that the Firm takes the correct action (say, her payoff increases by some small $\varepsilon > 0$). Then, in the one-shot game, the Agent will reveal all her information, and so the Firm has no incentive to pay. In this case as well, there is unravelling. But this unraveling argument does not extend to the infinite-horizon game, and it is possible to construct equilibria in our leading example in which the Agent is paid for a gradual release of information. The value of ε restricts how extreme the Firm's posterior belief can become before the Agent discloses all information. Nevertheless, the maximum equilibrium payoff to the competent Agent is continuous at $\varepsilon = 0$.
3. If discounting took place with every round of communication, then there would be no benefit in having arbitrarily many rounds. This is because the Agent faces a trade-off between collecting more money overall and collecting it earlier, and because the Firm ultimately prefers taking its outside option rather than waiting for another period, once the benefits from waiting become small. Hence, in the best equilibrium, the number of rounds in which communication actually takes place is bounded. However, as long as the players are not too impatient, the best equilibrium still involves a gradual release of information, and the number of rounds of active communication increases with the discount factor. It is easy to see that, as discounting vanishes, the payoff to the competent Agent must tend to her payoff in the undiscounted game. In our leading example, numerical simulations show that for the pure-strategy case this convergence occurs at a geometric rate.
4. As mentioned in the introduction, our mechanism is reminiscent of zero-knowledge proofs. But gradualism is a technological constraint in this literature. There is no counterpart

to the Firm's outside option, and the only objective is to convince the other party that the Agent holds the information. It is as if $w(p) = 0$ for $p < 1$, and $w(1) = 1$, in which case it is optimal to reveal all details but the "last key," increasing the Firm's posterior close to 1, and then to sell just that remaining piece. Gradualism arises in our mechanism precisely because the Firm's outside option depends on its belief, as is plausible in most economic applications. In fact, often the buyer has private information as well, and an inventor always risks making herself obsolete by revealing additional information to the Firm. Considering such a model, in which both parties hold private information, is left for future research.

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," *The American Economic Review*, **84**, 190–209.
- [5] Anton, J. and D. Yao (2002). "The Sale of Ideas: Strategic Disclosure, Property Rights, and Contracting," *The 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] Bruckner, A. M. and E. Ostrow (1962). “Some function classes related to the class of convex functions,” *Pacific Journal of Mathematics*, **12**, 1203–1215.
- [10] Che, Y.-K. and J. Sákovics (2004). “A Dynamic Theory of Holdup,” *Econometrica*, **72**, 1063–1103.
- [11] Compte, O. and P. Jehiel (2004). “Gradualism in Bargaining and Contribution Games,” *The Review of Economic Studies*, **71**, 975–1000.
- [12] Esó, P. and B. Szentes (2007). “The price of advice,” *The Rand Journal of Economics*, **38**, 863–1179.
- [13] Forges, F. (1990). “Equilibria with Communication in a Job Market Example,” *The Quarterly Journal of Economics*, **105**, 375–398.
- [14] Forges, F. and F. Koessler (2008). “Long Persuasion Games,” *Journal of Economic Theory*, **143**, 1–35.
- [15] Fudenberg, D. and J. Tirole (1991). *Game Theory*, MIT Press, Cambridge.
- [16] Kamenica, E. and M. Gentzkow (2011). “Bayesian Persuasion,” *American Economic Review*, **101**, 2590–2615.
- [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.

A Proofs

Because mediation imposes one fewer constraints on the payoff function to be determined than mixed strategies, as explained in Section 4.2.2, we prove the three theorems in the following order: first, Theorem 1 (pure strategies), then Theorem 3 (mediation) and then Theorem 2 (mixed strategies).

A.1 Proof of Lemma 2 and Lemma 1

The proof of Lemma 2 is by induction on the number of rounds. Lemma 2 immediately implies Lemma 1

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.¹⁶ Further, the type-1 Agent’s maximal payoff with k rounds to go is equal to

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

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. Clearly,

¹⁶In 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 $\{0, p_1\}$.

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 $p^{\frac{0}{1}} = 1 \geq p^*$. 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 its 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).

A.2 Proof of Theorem 1

Given a function f , the average function of f is denoted

$$f^a(x) := f(x)/x.$$

Given a non-negative function f on $[0, 1]$, let $\text{sha } f$ denote the largest weakly star-shaped function that is smaller than f . This function is well-defined, because (i) if f_1, f_2 are two weakly star-shaped functions lower than f , the pointwise maximum g (*i.e.* $g(p) := \max\{f_1(p), f_2(p)\}$) is star-shaped as well,¹⁷ and (ii) the limit of a convergent sequence of star-shaped functions is star-

¹⁷Given $p_1 < p_2$, let $g(p_1) = f_i(p_1), g(p_2) = f_j(p_2)$. Then $g^a(p_2) = f_j^a(p_2) \geq f_i^a(p_2) \geq f_i^a(p_1) = g^a(p_1)$.

shaped (Thm. 2, Bruckner and Ostrow, 1962), who also show that a star-shaped function must be non-decreasing.

The theorem claims that the equilibrium payoff, given w , and $\hat{w} := \text{sha } w$, is given by

$$V_1^p(p_0) = 1 - \hat{w}(\hat{p}_0) - \int_{\hat{p}_0}^1 \hat{w}^a(p) dp,$$

where $\hat{p}_0 := \min \{p \in [p_0, 1] : w(p) = \text{sha } w(p)\}$. Further, letting

$$I_w = \text{cl } \{p \in [0, 1] : \text{sha } w(p) = w(p) \text{ and } w^a \text{ is strictly increasing at } p\},$$

we show that the set of beliefs held by the firm is contained, and dense, in I_w if $I_w \neq \emptyset$. If $I_w = \emptyset$, any policy is optimal.

Let us start by showing that this payoff can be achieved asymptotically (*i.e.*, as $K \rightarrow \infty$). Let J_w denote the complement of I_w , which is a union of disjoint open intervals. Let $\{(p_n^-, p_n^+)\}_{n \in \mathbf{N}}$ denote an enumeration of its endpoints. Finally, let $\check{p}_0 := \min \{p \in I_w, p \geq p_0\}$. Note that, for all n , by continuity of w (using that $\frac{\hat{w}(p_n^+)}{p_n^+} = \frac{\hat{w}(p_n^-)}{p_n^-}$ by definition of (p_n^-, p_n^+)),

$$\hat{w}(p_n^+) - \hat{w}(p_n^-) - \int_{p_n^-}^{p_n^+} \hat{w}^a(p) dp = p_n^- (w^a(p_n^+) - w^a(p_n^-)) = 0.$$

Similarly, if $\hat{p}_0 < \check{p}_0$,

$$\hat{w}(\check{p}_0) - \hat{w}(\hat{p}_0) - \int_{\hat{p}_0}^{\check{p}_0} \hat{w}^a(p) dp = 0.$$

Fix any sequence of finite subsets of points $P^K = \{p_k^K : k = 0, \dots, K\} \subseteq I_w \cap [p_0, 1]$ (where p_k^K is strictly increasing in k), for $K \in \mathbf{N}$, with $p_0^K = \check{p}_0$, $p_K^K = 1$, such that P^K becomes dense in I_w as $K \rightarrow \infty$. Consider the pure strategy according to which, in the first period, if $\check{p}_0 > p_0$, the type-1 Agent gives away the information for free that leads to a posterior \check{p}_0 ; afterwards, the price paid in each period given that the posterior is supposed to move from p_k^K to p_{k+1}^K is given by

the maximum amount $p_k^K (w^a(p_{k+1}^K) - w^a(p_k^K))$. Failure to pay leads to no further disclosure, and failure to disclose leads to no further payment. Given K , the payoff of following this pure strategy is (by considering Riemann sums and using the equality from the previous equation)

$$\sum_{k=0}^{K-1} p_k^K (w^a(p_{k+1}^K) - w^a(p_k^K)) \rightarrow 1 - \hat{w}(\check{p}_0) - \int_{I_w \cap [\check{p}_0, 1]} \hat{w}^a(p) dp = 1 - \hat{w}(\hat{p}_0) - \int_{\hat{p}_0}^1 \hat{w}^a(p) dp.$$

Conversely, we show that (i) for any K , the best payoff given w is the same as for some weakly star-shaped function smaller than w , and (ii) if $w \geq \tilde{w}$, then $V_1 \leq \tilde{V}_1$. The result follows.

Note that the payoff from the sequence of beliefs $p_1, p_2, \dots, p_{K-1}, p_K = 1$, starting from p_0 is given by

$$\begin{aligned} & p_0(w^a(p_1) - w^a(p_0)) + p_1(w^a(p_2) - w^a(p_1)) + \dots + p_{K-1} \cdot (w^a(1) - w^a(p_{K-1})) \\ &= 1 - w(p_0) - (1 - p_{K-1})w^a(1) - \dots - (p_1 - p_0)w^a(p_1), \end{aligned}$$

so that

$$V_{1,K}(p_0) + w(p_0) = 1 - \sum_{k=0}^{K-1} (p_{k+1} - p_k)w^a(p_{k+1}).$$

Note that maximizing $V_{1,K}(p) + w(p)$ and maximizing $V_{1,K}(p)$ are equivalent, so this amounts to finding the sequence that maximizes the sum

$$1 - \sum_{k=0}^{K-1} (p_{k+1} - p_k)w^a(p_{k+1}),$$

with $p_0 = p$. Because $w \leq \tilde{w}$ implies $w^a \leq \tilde{w}^a$, we have just established the following.

Lemma 4 *Suppose that $\tilde{w} \geq w$ pointwise. Then, for every K , and every prior belief p_0 ,*

$$\tilde{V}_{1,K}(p_0) \leq V_{1,K}(p_0),$$

where $\tilde{V}_{1,K}(p_0)$ and $V_{1,K}(p_0)$ are the type-1 Agent's payoffs given outside option \tilde{w} and w , respectively.

To every sequence of beliefs $p_0, p_1, \dots, p_K = 1$, we can associate the piecewise linear function w_K on $[p_0, 1]$ that obtains from linear interpolation given the points

$$(p_0, w(p_0)), (p_1, w(p_1)), \dots, (1, 1).$$

Lemma 5 *For all K , p_0 , the optimal policy is such that the function w_K is weakly star-shaped.*

Proof: This follows immediately from the payoff from the formula for the price of a jump from p_1 to p_2 ,

$$p_1 (w^a(p_2) - w^a(p_1)).$$

Indeed, if p_1, p_2, p_3 are consecutive jumps, it must be that doing so dominates skipping p_2 , *i.e.*

$$p_1 (w^a(p_2) - w^a(p_1)) + p_2 (w^a(p_3) - w^a(p_2)) \geq p_1 (w^a(p_3) - w^a(p_1)),$$

or $w^a(p_3) \geq w^a(p_1)$. A similar argument applies to the first jump. \square

Note finally that the payoff from the sequence $\{p_1, \dots, p_K\}$ given w is the same as given w_K . The result follows. The asymptotic properties of the optimal policy follow as well.

We start with the theorem, which implies the lemma by a straightforward computation.

A.3 Proof of Theorem 3

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 . Due to the fact that this distribution is known, the Firm's belief must be a martingale, which means that, given p ,

$$\int_{[0,1]} p' dF_k(p' | p) = p, \text{ or } \int_{[0,1]} (p' - p) dF_k(p' | p) = 0. \quad (4)$$

To put it differently, $F_k(\cdot | p)$ is a mean-preserving spread of the Firm's prior belief p .¹⁸

¹⁸The notation $[0, 1]$ for the domain of integration emphasizes the possibility of an atom at 0. This, however, plays no role for payoffs, as there is no room for transfers once the prior drops to zero, and $w(0) = 0$, and we will then revert to the more usual notation.

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 its continuation payoff and its 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 $F_k(\cdot | p)$ assigns probability q to some posterior belief p' . This means that the Firm attaches probability q to its 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 .

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 its outside option, is X' and its belief is p' .

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

$$\int_0^1 (w(p') + X') dF_k(p' | p) - w(p) + \int_0^1 V_{1,k-1}(p', X') \frac{p'}{p} dF_k(p' | p), \quad (5)$$

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 (4) and the feasibility of X').

A.3.1 The Optimal Transfers

As a first step in the analysis, we prove the following.

Lemma 6 *Fix the prior belief p and the number of remaining rounds k . The best equilibrium payoff of the type-1 Agent, as defined by (5), 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: First of all, we must derive some properties of 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 her type. This modification can only relax any incentive (or individual rationality) constraint. This means that payoffs must satisfy (1) with equality, 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 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 (5), 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. □

The intuition behind this lemma is that to maximize V_1 , because the type-1 Agent assigns a smaller probability to the posterior decreasing than the type-0 Agent, 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 her type 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).

A.3.2 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 (5), we obtain that

$$V_{1,k}(p) = \sup_{F_k(\cdot|p)} \int_0^1 v_{k-1}(p'; p) dF_k(p' | p) - w(p), \quad (6)$$

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 (4), 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 (6) 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. We may now finally prove the theorem.

Proof of Theorem 3: 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,

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

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

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

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

A.4 Proof of Lemma 3 and Theorem 2

We adapt the arguments from the proof of Theorem 3. Recall that w is assumed to be weakly star-shaped (in particular, non-decreasing). Consider a mixed-strategy equilibrium. In terms of beliefs, such an equilibrium 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 $W_k = w(p) + X_k$), 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, it holds 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,

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

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, she must be indifferent between all elements in the support of its 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' . Assume (as will be verified) that in all relevant arguments, p' and $X \geq 0$ are such that it holds that

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

Recall that this is always possible if X is small enough, cf. Lemma 6. Furthermore, for the type-0 Agent to go along, we must verify that $V_{0,k} \geq X$. By substitution, we obtain that $V_{1,k+1}(p)$ is at most equal to

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

So let V_1^* denote the smallest fixed point larger than 0 of the map T given by

$$T(V_1)(p) = \sup_{F(\cdot|p)} \int_0^1 [w(p') + V_1(p')] dF(p' | p) + F_{k+1}(0 | p) \min_{p' \in \text{supp } F(\cdot|p), p' > 0} V_1(p') - w(p),$$

for which $V_1^*(1) + w(1) = 1$. The function V_1^* , and hence h^* is continuous by standard arguments.

As argued in the text, either $h^* := V_1^* + w$ is equal to \bar{h} at p , or it is locally concave at p . Indeed, for any $0 < p_1 < p < p_2 \leq 1$,

$$V_1^*(p) + w(p) \geq \frac{p_2 - p}{p_2 - p_1}(V_1^*(p_1) + w(p_1)) + \frac{p - p_1}{p_2 - p_1}(V_1^*(p_2) + w(p_2)),$$

and by choosing p_1, p_2 close to p , the constraint (that X is small enough) is satisfied. Clearly, also, h^* is no steeper than $p \mapsto w(p)/p$ (given $p < p'$, consider the distribution $F(\cdot | p)$ that splits p into $\{0, p'\}$, as explained in Subsection 4.1, so that h^* is no steeper than w). That is, h^* satisfies all four constraints from Section 4.2.1.

Recall that h^m is defined to be the smallest function satisfying the four requirements. This function is well-defined, because if h, h' are two functions satisfying these requirements, the lower envelope $h'' = \min\{h, h'\}$ does as well, and if (h_n) , $n \in \mathbb{N}$, is a converging sequence of functions satisfying them, so does $\lim_{n \rightarrow \infty} h_n$.

We now show that h^m cannot be improved upon. By monotonicity of the operator T , it follows that, starting from $h_0 := w$ and iterating, the resulting sequence $h_1 = T(h_0 - w) + w$, $h_2 = T(h_1 - w) + w$, etc. must converge to h^m .

To show that h^m cannot be improved upon, it suffices to consider arbitrary two-point distributions splitting p into $p_1 < p < p_2$.¹⁹ If all three beliefs belong to an interval in which $h^m < \bar{h}$, the result follows from the concavity of h^m on such intervals. If $p_1 = 0$, the result follows from the fact that h^* is no steeper than $p \mapsto w(p)/p$. If $p_1 > 0$ is such that $h^m(p_1) = \bar{h}(p_1)$, such a splitting is impossible, as $V_0(p_1) = 0$, and so the type-0 Agent would not pay $X > 0$, and hence the type-1 Agent could not be indifferent. Hence, we are left with the case in which $p_1 > 0$, $h^m(p_1) < \bar{h}(p_1)$, and $h^m(\tilde{p}) = \bar{h}(\tilde{p})$ for some $\tilde{p} \in [p_1, p_2]$, which can be further reduced to the case $h^m(p_2) = \bar{h}(p_2)$. The side bet X must equal $V_1(p_1) - V_1(p_2)$, and because $V_0(p_2) = 0$, we have $V_1(p_2) = (p_2 - w(p_2))/p_2$. We must have

$$V_0(p_1) = \frac{p_1 - w(p_1) - p_1 V_1(p_1)}{1 - p_1} \geq X = V_1(p_1) - V_1(p_2).$$

¹⁹Note that, with arbitrarily many periods, we can always decompose more complicated distributions into a sequence of two-point distributions. But the linearity of the optimization problem actually implies that two-point distributions are optimal.

This implies that $h_1(p_1) \leq 1 - (1 - p_1) \frac{w(p_2)}{p_2}$, or, rearranging, and using the formula for $V_1(p_2)$,

$$\frac{w(p_2)}{p_2} \leq \frac{1 - h(p_1)}{1 - p_1}.$$

Note, however, that, because h is no steeper than $w(p)/p$,

$$h(p_1) \geq h(p_2) - \int_{p_1}^{p_2} w^a(p) dp,$$

(recall that $w^a(p) := w(p)/p$) and hence, replacing $h(p_1)$ and rearranging,

$$w^a(p_2) \leq \frac{1}{p_2 - p_1} \int_{p_1}^{p_2} w^a(p) dp,$$

a contradiction, given star-shapedness (if w is weakly star-shaped on the entire interval $[p_1, p_2]$, the bet is feasible, but worthless).