

Selling Information*

Johannes Hörner[†] and Andrzej Skrzypacz[‡]

June 8, 2012

Abstract

We characterize optimal selling protocols/equilibria of a game in which an Agent first puts hidden effort to acquire information and then transacts with a Firm that uses this information to take a decision. We determine the equilibrium payoffs that maximize incentives to acquire information. Our analysis is similar to finding *ex ante* optimal self-enforcing contracts since information sharing, outcomes and transfers cannot be contracted upon. We show when and how selling and transmitting information gradually helps. We also show how mixing/side bets increases the Agent's incentives.

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 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 this paper, we study a dynamic game of selling information in which players cannot use external enforcement of contracts. Motivated by a moral hazard problem of acquiring information, we describe equilibria that maximize *ex ante* efficiency of a decision problem in which an Agent needs to acquire information that a Firm can later use to make a decision.

The selling information game is divided into rounds of communication. Within a round, the players first can transfer payments and then the Agent can send 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 has a binary information about the optimal Firm decision) 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 so that the type-1 Agent can always pass it but the type-0 Agent can pass it only with some known probability (so it is not a cheap talk). Easy tests have a high probability of being passed by the type-0. Divisibility of information means that there is a rich set of tests with varying difficulties.

In this game we construct tight bounds on the limits of the difference between type-1 and type-0 Agent payoffs as the number of possible communication rounds grows to infinity (we show an example where maximizing this difference is necessary for optimal incentive provision to acquire information). We characterize three such bounds: when we consider only pure-strategy equilibria (in which type-1 always passes the test), when we allow for mixed-strategy equilibria (when type-1 may be mixing between passing and failing a test) and finally if we allow for noisy tests that even the type-1 may not be able to pass (in the absence of noisy tests, the same outcome can be achieved with the help of a trusted intermediary, for instance).

Since we assume that the agents cannot commit to payments or information disclosure, our equilibria can be viewed as the best self-enforcing contracts that the players would like to coordinate on *ex ante*. Alternatively, these equilibria describe the maximum payoffs achievable in any equilibrium without external enforcement (as a function of the communication protocol) and hence allow us to divide the value of explicit contracts into the coordination part and the enforcement part.

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 (see Arrow, 1959). 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 difference in payoffs. That is, it is usually better for incentives if the Agent 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 leads the Firm to be indifferent between both decisions. 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 very rich set of tests and arbitrarily many rounds of communication, the benefit of splitting does not depend on either assumption.

Second, we show how mixed strategies can help improve performance of the contract. In the best pure-strategy equilibrium the type-0 Agent collects (in expectation) 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 type-1 Agent and the Firm may have different (endogenous) risk attitudes (more precisely, if the sum of their continuation payoffs is not concave).

Mixed strategies can be further improved upon if the players have access to tests that both types can fail with positive probability, or alternatively, by assuming that players have access to a trusted intermediary that can “noise up” the tests (i.e., the intermediary’s role is to allow the Agent to commit to a mixed strategy). Such tests allow us to exploit also the difference in beliefs between the Firm and the two types of the Agent, regarding the very own evolution of the Firm’s belief about the Agent’s type. This form of communication makes it possible for type-1 of the Agent to use side-bets (in which the Agent pays the Firm upon a failure of the test) to extract the *entire* expected surplus of the information.

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 party 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.¹

Most of the paper (Sections 2 and 3) analyzes the information sales problem for the specific payoff structure that is inherited from the motivating game of information acquisition. However, there is nothing particular about this motivating example. In Section 4, we generalize our results to arbitrary specifications of how the Firm’s payoff varies with its belief about the Agent’s information. 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

¹We 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.

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, i.e. before her type is known. In our case, this is trivial and does not deliver differential payoffs to the Agent's types. Therefore, such an equilibrium does not provide the Agent with incentives to engage in inventive activities in the first place (which determine the probability with which an Agent is informed). To do so requires identifying the best equilibrium from the point of view of a particular type of the Agent.

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.

But it is important to stress that by information, we mean here information that is relevant for commonly known choice, such as an investment opportunity. The object of this information

is not unknown *per se*.

Our focus (proving one owns information) and instrument (tests that imperfectly discriminate between an Agent that holds information or not) 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). Eső and Szentes (2007) take a mechanism design approach to this problem, while Gentzkow and Kamenica (2009) apply ideas similar to Aumann and Maschler (1995) to study 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 shall motivate this example by considering the following decision problem.

2.1 The Decision Problem

Consider the problem of an Agent who must decide whether to acquire information or not. This information will then be sold to a Firm in a second stage. This second stage is the ultimate focus of our analysis. Here we describe this first stage in which information is acquired.

There are two states of Nature, $s_N \in \{L, H\}$, with prior $\mathbb{P}[s_N = H] = \rho \in (0, 1)$. The Agent privately chooses an effort level $e \in [0, \bar{e}]$ at cost $c(e)$. The differentiable function c is strictly increasing and convex, with $c(0) = c'(0) = 0$, and $c'(\bar{e}) > \rho$.

Given e , the Agent privately observes a signal, the *private state* $\omega \in \Omega := \{0, 1\}$, with

$$\mathbb{P}[\omega = 0 | s_N = L] = 1, \mathbb{P}[\omega = 1 | s_N = H] = e.$$

Hence, conditional on private state 1, the probability of $s_N = H$ is 1, while

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

Given effort e , the unconditional probability of the Agent being in the private state 1 is $p_0 := \rho \cdot e$.

Information is useful because an investment decision must be taken, whose return depends on the state of nature. Not investing yields a safe (i.e., state-independent) payoff normalized to 0. Investing yields a payoff 1 when $s_N = H$ and $-\hat{\gamma} < 0$ when $s_N = L$. Hence, conditional on the Agent's private state, this implies that investing yields an expected return of

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

if $\omega = 0$, and 1 if $\omega = 1$. We assume that $\gamma > 0$ (for all feasible e), i.e., investing is optimal if and only if the private state is $\omega = 1$. Hence, the expected surplus is $\rho \cdot e - c(e)$, so that the

first-best effort solves

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

However, it is not the Agent, but the Firm who takes the investment decision and reaps its benefits. The Firm observes neither the exerted effort level nor the resulting private state. Based on its expectation e^* about the Agent's effort level, the Firm forms a posterior belief p_0 that $\omega = 1$. This belief affects how the surplus will be split in the second stage. Anticipating revenues $V_1(p_0)$ and $V_0(p_0)$ from the second stage, if the private state is $\omega = 1$ or 0, respectively, the Agent's effort e^* must maximize

$$\rho e V_1(p_0) + (1 - \rho e) V_0(p_0) - c(e),$$

and so, assuming that $V_1(p_0) > V_0(p_0)$, e^* solves

$$c'(e^*) = \rho(V_1(p_0) - V_0(p_0)). \quad (1)$$

Unless $V_1(p_0) = 1$ and $V_0(p_0) = 0$, equilibrium effort is below first-best.² This is a standard hold-up problem, although investment is unobservable here. Because c is convex, social welfare is then maximized when the difference $V_1(p_0) - V_0(p_0)$ is highest.

This gives us the motivation and the objective function for the game played in the second stage that determines the split of the surplus, to which we now turn.³

2.2 The Game of “Selling Information”

Some basic ingredients of this game are inherited from the decision problem. As this game can be understood independently from the decision problem, we repeat them here, so that the exposition be self-contained.

There are two risk-neutral players: an Agent (she) and a Firm (it). There are two states of the world, $\omega \in \Omega := \{0, 1\}$. The Agent is privately informed of the state of the world at the

²Shares of surplus V_0, V_1 will always be in $[0, 1]$.

³It is straightforward to extend this decision problem to the case of more general outside options, as described in Section 4.

beginning of the game, but the Firm is not. The Firm's prior belief that the state is 1 is p_0 , which is common knowledge. The fact that the Agent is perfectly informed is a normalization.⁴

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

Because the Agent knows the state, call her the type-1 Agent if $\omega = 1$, and the type-0 Agent otherwise.

Note that, absent any information revelation, the Firm's optimal action is to invest if and only if its belief p that the state is 1 satisfies

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

and obtain thereby a payoff of

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

where $x^+ := \max\{0, x\}$. While our analysis will cover both the case in which the prior belief p_0 is below or above p^* , we shall often focus on the more interesting case in which p_0 is smaller than p^* , unless stated otherwise. The payoff $w(p)$ is the Firm's *outside option*, and we shall generalize the analysis to outside options with rather arbitrary specifications in Section ??.

In each of the K rounds before the action is taken, the Firm and Agent can make a monetary transfer, and the Agent can reveal some information if she wishes to. More precisely, the strategy has two parts. In rounds $k = 1, \dots, K$, as a function of the history of transfers and information disclosures up to that point, the Agent and the Firm can simultaneously make a non-negative transfer t_k^A and t_k^F , respectively, to the other party.⁵ Second, once these transfers are made and

⁴Here, a state of the world is what we called a private state in the decision problem. With that interpretation of the information that is available to the Agent, the fact that she is perfectly informed is somewhat tautological.

⁵The 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. Indeed, as we shall see, such payments are irrelevant when only pure strategies are considered. But they play a critical role once more general strategies are considered.

observed, the Agent may disclose some verifiable information.⁶

Information disclosure/gradual persuasion is modeled as follows. The Agent chooses a number $m \in [0, 1]$. This choice is observed by the Firm. Number $1 - m$ represents the difficulty of the test that the Agent picks: The type-1 Agent can always pass the test (though she can choose to fail it), while the type-0 Agent can only pass it with probability m . The realizations of tests are independent across periods (and values of m), conditional on the state.

Note that the Agent can always choose an uninformative test if she wishes to, by picking $m = 1$. This is interpreted as not revealing any information. If $m < 1$ and the Agent passes the test, we say that information gets disclosed.

Note that, given any belief $p \in (0, 1)$ that the Firm might assign to state 1 and for any $p' \geq p$, there exists a test that leads the Firm to update its belief to p' , if the Agent passes it. Indeed, if the Agent picks the value

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

independently of her type, and does not flunk it on purpose, it follows from Bayes' rule that the posterior belief assigned to $\omega = 1$ is equal to

$$\frac{p}{p + (1 - p)m} = p'.^7$$

If the Agent fails the test, then the Firm correctly updates its belief to zero.⁸

The set of possible tests that we assume is rich, and implies that information is perfectly divisible.⁹ This allows us to conduct the analysis entirely in terms of beliefs, and to make abstraction from issues relative to the type of information that is being disclosed, leaving open some fascinating questions (for instance, in which order should information be released?). Tests

⁶It is easy to see that nothing hinges on this timing. Payments could be made sequentially rather than simultaneously, and could occur after rather than before information disclosure.

⁷A modeling issue arises for $p = 0$. What if the Firm, after some history of transfers and disclosures, assigns probability 0 to $\omega = 1$, but the Agent then passes a test with $m = 0$? But our purpose is to identify the best equilibrium, not to characterize the set of all equilibria, and so this issue is irrelevant: the equilibria we shall describe remain equilibria if it is required that players cannot switch away from probability 1 beliefs, and remain the best equilibria if this requirement (not imposed by perfect Bayesian equilibrium) is dropped.

⁸That is, unless the type-1 Agent is expected to flunk it on purpose with positive probability.

⁹Yet our result that it is better to “split information” by using two easier tests instead of a difficult one also holds when the set of tests is coarse.

only serve the purpose of modeling how beliefs can evolve gradually, and could be replaced with any other formalism achieving the same end. But richer sets of tests could be conceived of and will be considered in the analysis: for instance, we might wish to consider tests that even the type-1 Agent could fail with some positive probability, so that the Firm's belief that the Agent is of type 1 can go down just as gradually as it can go up.

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

2.2.1 Histories and Payoffs

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 result of that test (which is either positive, 1, or negative, 0). The set of all such histories is denoted H_k (set $H_0 := \emptyset$). Given some final history h_K (this does not include the Firm's final action to invest or not), the Agent's realized payoff is simply 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 state 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 .

2.2.2 Strategies and Equilibrium

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 $t_k^F := \tau_k^F(h_k)$ 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 $\{\tau_k^A, \mu_k^A, \rho_k^A\}_{k=0}^{K-1}$, where (i) $\tau_k^A : \Omega \times H_k \rightarrow \mathbb{R}_+$ is a probability transition specifying the transfer $t_k^A := \tau_k^A(h_k)$ 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 information that is released in round k (i.e., the value of m), as a function of the state, the history up to the current round, and

This game admits a plethora of equilibria. Our focus is to identify the equilibrium that maximizes the spread $V_1 - V_0$. Given the decision problem of Subsection ??, the motivation is two-fold. First, if the Firm and the Agent could coordinate *ex ante* (i.e. before the decision problem) and make side-payments, then clearly it would be in their interest to choose an equilibrium that maximizes the *ex ante* payoffs, as a form of a self-enforcing contract (or relational contract). Second, we are interested in the upper bound on efficiency that can be achieved without property rights, through such self-enforcing contracts, to better understand the agency costs, and how they depend on the coordination failures vs. on the constraints from the way information is sold and acquired (that is, in the spirit of mechanism design, we separate the question of what is the most any equilibrium can achieve from the question of how to coordinate on that equilibrium).

To recap, we are interested in the limit of the difference in the Agents' payoffs as the number of rounds becomes arbitrarily large.¹² To do so, we shall relax the problem by assuming that players have access to a public randomization device at the beginning of every round (a draw from a uniform distribution), as this will facilitate one argument. The resulting equilibria that we consider (whether we consider pure or mixed strategies, or allow a mediator) turn out not to take advantage of this device, so that the findings hold for the model without it.

2.3 Preliminary Observations

If the probability of state 1 is p , given the history h_k , then the expected surplus (assuming that the Firm takes an optimal eventual decision) is $p \cdot 1 + (1 - p) \cdot 0 = p$. This means that continuation payoffs must satisfy

$$pV_1(h_k) + (1 - p)V_0(h_k) + W(h_k) \leq p. \quad (2)$$

From any history onward, the Agent can secure a payoff of zero, independently of her type:

$$V_1(h_k) \geq 0, V_0(h_k) \geq 0. \quad (3)$$

is straightforward and omitted.

¹²The equilibrium we shall obtain is also an equilibrium of the infinite-horizon, undiscounted game, but taking limits allows us to uniquely pin down the limiting strategy profile.

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

$$w(p) = (p - \gamma(1 - p))^+ . \quad (4)$$

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

It is easy to see that (??) (with $p = p_0$), (??) and (??) define the set of feasible and individually rational (continuation) payoffs. Note that the type-1 Agent cannot receive more than $1 - w(p)/p$, the entire *expected* surplus. While this is the maximum she can hope for, this is still short of the actual surplus, *given* state 1, $1 - w(p)$; hence, even appropriating the entire expected surplus does not solve the moral hazard in the decision problem altogether.

We conclude this section with a series of observations about the selling information game. Fix K and p_0 throughout.

- *There exists an equilibrium (the “worst” equilibrium) that minmaxes both players simultaneously:* Making no transfers (expecting none) and releasing no information is an equilibrium, with payoffs

$$W = w(p_0), V_1 = V_0 = 0.^{13}$$

Although there are many ways for the Agent to signal her information through transfers or deviations in terms of the test difficulty that she picks, and therefore, many out-of-equilibrium beliefs to “worry” about, such beliefs play no role: observable deviations by the Firm do not affect its beliefs (this is the “no signalling what you don’t know” ingredient of perfect Bayesian equilibrium), and observable deviations by the Agent can be deterred through the threat of reverting to this worst equilibrium, independently of how this affects the Firm’s belief.¹⁴

An equilibrium is *efficient* if the constraint (??) is binding, that is, if the type-1 Agent discloses all her information eventually, on the equilibrium path.

- *If an equilibrium gives (V_0, V_1) to the Agent, there is an efficient equilibrium that does so:*

¹³For brevity, we often write V_ω for $V_\omega(h_k)$, but because the maximum equilibrium payoff only depends on the Firm’s belief (and the number of periods left), we also sometimes write $V_\omega(p)$, where p is this belief. Finally, we also write V_ω for the resulting function of the belief. Hopefully, no confusion will arise.

¹⁴This also implies that the equilibrium payoffs that we shall determine can easily be obtained as well for alternative orders of moves within a period, such as disclosure before transfer, etc.

Indeed, the Agent can always disclose the state in the last period *on the equilibrium path*. This cannot weaken the incentives for the players to carry out the planned transfers (because it can only increase the payoff from following the specified equilibrium actions), but it guarantees that the correct action is taken.¹⁵

Some efficient equilibrium payoffs giving all the surplus to one of the parties are easy to describe.

- *There is an equilibrium in which the Firm receives $W_0 = p_0$:* no transfers are ever made, and the type-1 Agent reveals the state in the last period, so that the posterior belief is 1 with probability p_0 , and 0 otherwise.

- *There is an equilibrium in which the Agent's expected payoff is $p_0 V_1 + (1 - p_0) V_0 = p_0 - w(p_0)$:* the Firm pays this amount in the first period, and the Agent reveals the state. If the Firm fails to pay, or the Agent deviates, play reverts to the worst equilibrium.

This shows that attaining the maximum *expected* payoff of the informed player is trivial in our game, unlike in many games with incomplete information (see Aumann and Maschler, 1995). Note also that, since the type-1 Agent can always mimic the type-0 Agent, her payoff must be at least as high as the type-0's payoff. This implies that the maximal equilibrium payoff for the type-0's Agent is the one that maximizes the Agent's *ex ante* payoff, as described above.

However, all these equilibria are terrible for providing incentives to acquire information: in all of them the two types of the Agent earn the same payoff and hence there is no return to the effort. What is non-trivial is to identify an equilibrium that maximizes $V_1 - V_0$, the difference in payoffs of the two types. We now argue that maximizing V_1 among all equilibria is “equivalent” to maximizing $V_1 + W$, the sum of the Firm's and type-1 Agent's payoff, as well as to maximizing $V_1 - V_0$:

- *An equilibrium that maximizes V_1 also maximizes $V_1 + W$ over all equilibria:* Given some

¹⁵This is clear if only at most the Agent randomizes on path, as her payoff from releasing additional information at the end does not affect her incentives. If the Firm is supposed to randomize on path, there exists an equivalent equilibrium that takes advantage of the public randomization device in which, conditional on the device's outcome, its action is pure, and so its incentives from following the equilibrium action are reinforced by this additional disclosure. This is the only point in the analysis in which the relaxation (to a game with a public randomization device) is used.

equilibrium yielding payoffs (V_0, V_1, W) , note that

$$V_1 + W \leq V_1 + w(p_0).$$

Otherwise, by simply starting from the equilibrium that yields V_1 to the type-1 Agent and W to the Firm, and by increasing the initial transfer that the Firm is asked to make by an amount $W - w(p_0)$, we would obtain another equilibrium in which the type-1 Agent gets a payoff strictly above V_1 –a contradiction. Given that $w(p_0)$ is fixed, the conclusion follows.

Therefore, the equilibrium that maximizes the type-1 Agent's payoff cannot leave any surplus to the Firm.

- *An equilibrium that maximizes V_1 also maximizes $V_1 - V_0$:* Efficient equilibria, to which attention can be restricted to, satisfy

$$p_0 V_1 + (1 - p_0) V_0 + W = p_0, \tag{5}$$

so that

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

Thus, given the prior belief p_0 , maximizing the payoff difference $V_1 - V_0$ is equivalent to maximizing the sum $V_1 + W$, but as we have already remarked, this is in turn implied by maximizing V_1 only. Therefore, we can simplify further and focus on maximizing V_1 , the type-1 Agent's payoff.

- *The set of equilibrium payoffs 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, the highest equilibrium payoff for the type-1 Agent has a well-defined limit given p_0 that we shall seek to identify. The functions V_1^p, V_1^m, V_1^{int} are the (pointwise) limit payoffs (as p_0 varies) in pure, mixed unobservable, and mixed observable strategies, that we shall consider in turn.




Figure 1: A feasible action

3 Equilibrium Analysis

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?

3.1 Pure Strategies

We start by considering pure strategies by the Agent. A pure strategy calls for the Agent to disclose a specific piece of information at each round, i.e. for both types to choose a test with a particular equilibrium-path m and for type-1 to pass it for sure and type-0 to pass it with probability m .

This implies that, from the Firm's point of view, and ignoring the uninteresting case in which the Agent is supposed to reveal nothing ($m = 1$), its posterior will take one of two values: either it will jump from p_0 up to some p' , if the piece of information is revealed. Or it will jump down to zero. This is illustrated in Figure ???. The two arrows indicate the two possible posterior beliefs. Note that, as a stochastic process, and 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, however, from the Agent's point of view. Given her knowledge of the state, she assigns different probabilities to these posterior beliefs than the Firm. If she is the type-1 Agent, she knows for sure that the belief will not decrease over time. If she is the type-0 Agent, the expectation of the posterior belief is below p_0 (the process is then a supermartingale).

More generally, an equilibrium outcome specifies a martingale splitting, summarized by the

sequence of Firm's beliefs that the state is 1, conditional on all pieces of information having been exhibited up that round (all test difficulties chosen and the results of the tests). On the equilibrium path, as long as the Agent passes the tests, Firm's beliefs follow a non-decreasing sequence $\{p_0, \dots, p_{K+1}\}$ which starts at the Firm's prior belief, p_0 , and ends up at $p_{K+1} = 1$ (assuming, without loss, that the equilibrium is efficient). If a piece of information fails to be disclosed (i.e. the Agent fails a test), the posterior immediately drops to zero.

Of course, an equilibrium must also specify transfers, as well as how players behave off the equilibrium path. The most effective punishment for deviations (whether in terms of information disclosure or payment) is reversion to the worst equilibrium, and this is assumed throughout.

It turns out that type-1 Agent payoff decreases if the Firm is given any payoff in excess of its outside option in this or future periods. It is obvious for the first round, since the payoffs are transferred one-to-one between the Firm and the Agent. On the one hand, the Agent could demand more in earlier rounds by promising 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 to the type-1 Agent of promising this surplus. The reason is that the Firm assigns a lower probability than the type-1 Agent to the posterior increasing (and promising surplus after the posterior drops to zero is not incentive compatible).

Therefore, if the Firm's belief in the next round is either p_{k+1} , or 0, given the current belief p_k , then the Firm is willing to pay

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

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

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

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.

Unless the Agent can reveal the information slowly, the answer is negative: If $K = 1$, the highest equilibrium payoff to the type-1 Agent is equal to $p_0 - w(p_0)$. With one round of communication, the payoff of the Agent can come only from the payment in the first (and only) round. Therefore, the payoffs of the two types of Agents have to be the same in all equilibria (for $K = 1$). To identify the best for the type-1 Agent, recall that we can focus on efficient equilibria, in which the posterior is either 0 or $p_1 = 1$. Because beliefs must follow a martingale from the Firm's point of view, it must be that the probability that the posterior is p_1 is p_0/p_1 , because

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

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

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


Note that, when $p_0 \leq p^*$, the highest payoff in one round that the type-1 Agent can get in equilibrium is simply the prior p_0 . Note also that this payoff is increasing in $p_0 \leq p^*$.

This immediately suggests one way to improve on the payoff with as little as two rounds. In the first step, the Agent discloses for free the piece of information leading to a posterior belief of p^* (or 0, if she fails to do so). In the second round, the equilibrium of the one-round game is played, given the belief p^* . This second and only payment yields

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

The right panel of Figure ?? illustrates. 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 clearly larger than the payment with only one round, given by the length of the vertical segment at p_0 .

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 will be given at the end of this subsection.

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




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

p increases.

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

What does the maximum payoff converge to as the number of rounds increase? Here is a heuristic derivation of the solution. Note that, for $p \geq p^*$, the payment that the type-1 Agent can extract from the Firm if the following posterior belief is $p' \in \{0, p + dp\}$ is (observing that, from




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

the martingale property of the Firm's beliefs, the test must be passed with probability $\frac{p}{p+dp}$,

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

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

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 will follow as immediate corollary from the next one. Note that this payoff is

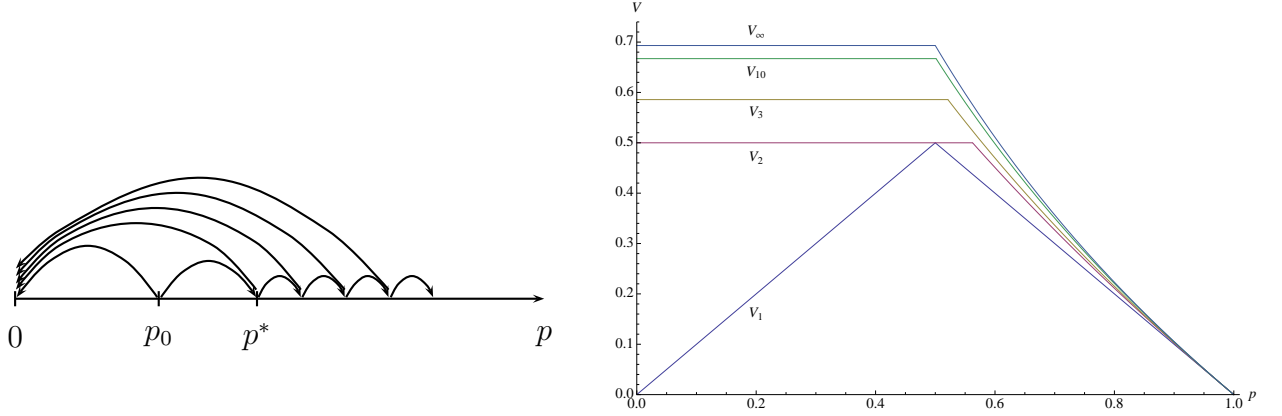


Figure 5: Revealing information in many steps (left); Payoff as a function of K (right).

independent of p_0 (for $p_0 < p^*$). Indeed, the first chunk of information, leading to a posterior belief of p^* , 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 what it is 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 ??.

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 maximal equilibrium payoff of the type-1 Agent with K rounds, given the Firm's prior belief p_0 , 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 ?? 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 produce the relevant information: hence she loses less, and might even gain (for instance, she might not have been able to produce the first piece of evidence that is given away at $p < p^*$). As a result, and quite generally, 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 in the analysis of mixed strategies that we do next.

An implication of this analysis is that, with pure strategies, there is no role for payments going from the Agent to the Firm. We believe, but have not shown, that the converse also holds, and that, without payments from the Agent to the Firm, one cannot improve on the equilibrium in pure strategies.

3.2 Mixed Strategies

We now consider mixed strategies by the Agent. Specifically, consider the following scenario. The type-1 Agent passes the test with positive but non-unitary probability; that is, she flunks on purpose some of the time. The type-0 Agent passes the test whenever she is able to. This

requires (i) the type-1 Agent to be indifferent between the two resulting continuations, and (ii) the type-0 Agent to (weakly) prefer not flunking the test.

In this case, failure to exhibit information does not lead to a posterior of zero. Indeed, the type-1 Agent might conceivably mix in such a way that exhibiting information leads to a lower posterior (though this won't occur in the analysis).

Whether one views mixed strategies as plausible in their own right or not, such dynamics of beliefs would also result from pure strategies with an appropriately extended set of actions: if the Agent can commit to run a test which is noisy (e.g., applying her expertise to a particular task, or letting the Firm experiment with, or make measurements of, her invention), the posterior belief will not necessarily drop to zero after a failure (operating systems *do* crash occasionally). In fact, such tests endow the Agent with even more commitment than mixed strategies as considered here, as they do not require the Agent to be indifferent over the resulting outcomes. The importance of such commitment will be evaluated in the next subsection.

One might wonder what the type-1 Agent could gain from using mixed strategies. The rationale is actually well-known. The type-1 Agent and the Firm have both differences in preferences over posterior beliefs, and differences in beliefs about the event that these posterior beliefs materialize. The next subsection will show how to take advantage of the heterogeneity in beliefs. Mixed strategies take advantage of the heterogeneity in preferences.

To illustrate this, consider the case in which $\gamma = 1$, so that $p^* = 1/2$, and consider the limiting case $K = \infty$, for simplicity. Using the best equilibrium (for the type-1 Agent) as a benchmark, the preferences of the Firm are piecewise affine in p ($w(p) = (2p - 1)^+$). Meanwhile, the type-1 Agent has a payoff function that is convex in p (over $[1/2, 1]$), given by $-\ln p$. We shall use side bets to take advantage of this differential attitude towards the resolution of uncertainty. Of course, if the type-1 Agent gains from such side bets, and the Firm does not lose from them (as its payoff is already down to its outside option), it must be that the type-0 Agent loses. Her payoff function is given by $V_0(p) = 1 + (p \ln p)/(1 - p)$. See the left panel of Figure ??.

Side bets, however, require payments to go back and forth between the Firm and the Agent.¹⁶

¹⁶As mentioned, we do not know what payoffs can be obtained if mixed strategies are allowed, but payments from the Agent to the Firm are not. We cannot rule out that some of the back payments of the bets could take the form of deductions on later payments by the Firm, but is unclear how far such deductions could substitute

If the Firm pays more than the “fundamental” expected value of the information disclosed, in anticipation of the returns of such side bets, it had better be that the Agent has incentives to honor such payments if necessary. If the posterior belief dropped to zero, even the threat of reversion to the worst equilibrium could not discipline the Agent into paying back. Therefore, the stakes of such bets are limited on two accounts: the type-0 Agent should be willing to make the requisite payments if the case occurs, and the type-1 Agent must be indifferent between the two continuation equilibria. Note that, if the type-1 Agent is indifferent between the two continuations, the type-0 Agent prefers the one starting with the higher posterior belief, given their relative preferences over starting beliefs, so that the type-0 Agent will disclose the requisite information, whenever she is able to (as we argued at the end of the previous subsection).

The left panel of Figure ?? illustrates how the mixing works, starting from a given belief $p > 1/2$. If information is disclosed, the Firm becomes more optimistic, with a corresponding posterior of $p + \Delta$, for some $\Delta > 0$. If it does not, the Firm becomes more pessimistic, with a posterior of $p - \Delta > 1/2$: the type-1 Agent randomizes in the right proportion for this posterior to arise, given that the type-0 Agent will disclose the information whenever she is able to. 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 Agent is expected to pay the Firm an amount $X > 0$ if the event $p' = p - \Delta$ realizes. We must set X so that type-1 Agent is willing to randomize. 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 $X = \ln(p + \Delta) - \ln(p - \Delta)$. 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 discloses the information whenever she is able to. We must also pick Δ sufficiently small to ensure that $V_0(p - \Delta) \geq X$, so that the type-0 Agent will not renege on the payment. Because, by definition of X , $X = 0$ when $\Delta = 0$, it is possible to find a small enough $\Delta > 0$ for this to hold.

Because both posterior beliefs are equally likely, the Firm is willing to pay $X/2$ upfront in

for explicit payments by the Agent.




Figure 6: Construction of the side bet (left); Maximum limit payoff $V_1^m + w$, $\gamma = 1$ (right).

exchange for this contingent future payment. This gives a total payoff of

$$\hat{V}_1(p) := \frac{\ln(p + \Delta) - \ln(p - \Delta)}{2} + \ln(p - \Delta) = -\frac{\ln(p + \Delta) + \ln(p - \Delta)}{2} > -\ln p$$

to the type-1 Agent, where the second term is the continuation payoff from the next round onward (which is equal across posterior beliefs, by construction), and the strict inequality follows from Jensen's inequality.

We have just improved on our upper bound based on pure strategies. The key here was the convexity of the type-1 Agent, relative to the Firm's payoff function (i.e. the convexity of $w(p) + V_1(p)$). What is the limit of using such mixing/side bets to improve V_1 ?

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 $h(p) := V_1^m(p) + w(p)$. There are two possibilities that could prevent an extra round with side bets to improve upon a given equilibrium payoff. Either $V_0^m(p) = 0$ at some p , so that by feasibility and individual rationality $h(p) = 1 - (1 - p)w(p)/p$ and it cannot increase any more. Or, $h(p)$ is locally concave, preventing further improvements through side bets. As we add rounds with the side-bets the difference in curvatures of $V_1(p)$ and $w(p)$ goes down. Does it vanish before we reduce $V_0(p)$ down to zero?

Other than these two bounds on h (that it has to be either equal to the upper bound or locally concave) we additionally know that h cannot be steeper than $w(p)/p$: indeed, starting from p_0 , we can always use a pure-strategy with posterior beliefs in $\{0, p_1\}$, so that

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

or

$$\frac{h(p_1) - h(p_0)}{p_1 - p_0} \leq \frac{w(p_1)}{p_1}. \quad (6)$$

Finally, h must exceed w , as the type-1 Agent's payoff is non-negative. What is the lowest function that satisfies these four requirements? In our example, it gives us:

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

See the right panel of Figure ?? . In Appendix, we prove that this is the limit maximum equilibrium payoff of the type-1 Agent with mixed strategies.¹⁷ That is, full extraction occurs for high enough ($p \geq p^m$, in which case $V_0(p) = 0$) but not for low beliefs. Still, this is a marked improvement upon pure strategies.

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 $p_0 < p^*$,*

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

3.3 Intermediary/Noisy Tests

Mixed strategies only allowed us to take advantage of the differences between the Firm and the type-1 Agent as long as their preferences had different curvatures. This constrains how much

¹⁷Roughly, 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.

surplus can be channeled from the type-0 Agent to the type-1 Agent. While the type-1 Agent and the Firm had different beliefs regarding the occurrence of future events –in particular, the Firm’s next belief– such differences could not be leveraged, because the type-1 Agent had to be indifferent over both continuations, so that it did not matter for the type-1 Agent how likely each posterior belief was.

If the Agent could commit to such a mixed action, this constraint no longer applies. As usual, this is a matter of the definition of what is observable: in the case of random, but informative tests, all parties could observe that the Agent is running such a test, and the Agent would no longer be in control of its outcome. The previous subsection considered unobservable mixed actions. Alternatively, if the Agent delegates the decision to disclose the information to a disinterested third party –what we call an intermediary– the Agent can no longer control whether the information gets disclosed or not.

In this section, we allow such an intermediary. Formally, we drop the requirement that the Agent be indifferent over the actions in the support of her mixed strategy. As we argue, this allows the type-1 Agent to further improve on her payoff. The key still lies in the design of side bets. This time around, those bets take advantage of the difference in beliefs.

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

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




Figure 7: The Role of Side Bets

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

Consider then the following scheme when there are arbitrarily many rounds. In the second round, it is understood that the Agent will make a payment of $V_0(\frac{1}{6})$ to the Firm if and only if the realized message is low (in particular, there is no payment by the Agent to the Firm if the realized message is high). Aside from this one-time, conditional payment from the Agent to the Firm, all payments by the Firm to the Agent, and all information disclosures from the Agent to the Firm occur from the second round onward as in the pure-strategy equilibrium without an intermediary (which is obviously possible even with an intermediary), given the realized message.

In the initial round, before the message is sent, the Firm must pay the difference between its expected continuation payoff and its current outside option, 0. (If it fails to do so, we switch to the worst equilibrium, as usual). How much is the Firm willing to pay? Note that, if there was no payment from the Agent to the Firm conditional on a low message, it is not willing to pay anything, since its outside option after either message is still 0. Nevertheless, because it expects

to receive $V_0\left(\frac{1}{6}\right)$ in an event whose probability is $1/2$ from its point of view, it is willing to pay up to $V_0\left(\frac{1}{6}\right)/2$ upfront in this scheme. How much is this scheme worth to the type-1 Agent? Her expected payoff is:

$$\frac{1}{2}V_0\left(\frac{1}{6}\right) + \frac{1}{4}\left[V_1\left(\frac{1}{6}\right) - V_0\left(\frac{1}{6}\right)\right] + \frac{3}{4}V_1\left(\frac{1}{2}\right) = V_1\left(\frac{1}{3}\right) + \frac{1}{4}V_0\left(\frac{1}{6}\right) > V_1\left(\frac{1}{3}\right).$$

To see this, note that she gets $V_0\left(\frac{1}{6}\right)/2$ up-front, $V_1\left(\frac{1}{6}\right) - V_0\left(\frac{1}{6}\right)$ in the event that the message is low (an event to which she assigns probability $1/4$) and $V_1\left(\frac{1}{2}\right)$ in the event that the message is high. Since $V_1(p)$ is constant for $p \in [0, \frac{1}{2}]$, the equality follows. As a result, with this scheme, her payoff with a prior $1/3$ is strictly larger than $V_1\left(\frac{1}{3}\right)$, her maximal payoff without an intermediary.

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

Observe that such a scheme is not possible without an intermediary, because the type-1 Agent is not indifferent over realized messages. She strictly prefers the high message to obtain, so that such a scheme cannot be replicated by mixed strategies without an intermediary.¹⁸ Second, note that the payment that the Agent makes if a low message occurs is not informative *per se*. This is because this payment is no larger than V_0 , and the continuation payoffs of the Agent is at least

¹⁸The payoff that is used as continuation payoff here –the maximum type-1 equilibrium payoff in pure strategies– is affine over the beliefs considered; hence, mixed strategies (without an intermediary) cannot improve her payoff, since the type-1 Agent and the Firm have identical risk attitudes. Alternatively, we could have used as continuation payoff the maximum type-1 equilibrium payoff in mixed strategies: because this payoff is also constant over these beliefs, and the type-0 Agent gets a positive payoff, the same construction would work, and improve on this maximum payoff.

as much, independently of her type. Higher payments would not work, because the type-0 Agent would not be willing to make it given the continuation equilibrium, and so the occurrence of a payment or not would convey information about the Agent's type. From the left panel of Figure ??, it is clear that, the closer the expected payoff pV_1 of the type-1 Agent is to the total surplus p , the smaller is the resulting V_0 , and so, the smaller the scope for such a scheme becomes. But as long as V_0 remains strictly positive, such schemes remain possible.

We concluded the previous subsection by noting the two constraints preventing full surplus extraction. Here, only $V_0(p) \geq 0$ remains. Hence, it should be no surprise that such bets allow full surplus extraction. The maximum equilibrium payoff of the type-1 Agent tends to, as $K \rightarrow \infty$,

$$V_1^{int}(p) := 1 - \frac{w(p)}{p},$$

giving us the following corollary:

Lemma 4 *As $K \rightarrow \infty$, the maximum equilibrium payoff to the type-1 Agent with an intermediary tends to, for $p_0 < p^*$,*

$$V_1^{int}(p_0) = 1.$$

4 General Outside Options

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

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. These assumptions on w are maintained throughout the next three subsections.

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 we are taking w

as a primitive here, we do not assume such a property here. We consider the three cases of pure-strategy, mixed-strategy, and of an intermediary in turn.

4.1 Pure Strategies

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 *star-shaped*, i.e., if and only if the average, $w(p)/p$, is a strictly increasing function of p .¹⁹ 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.²⁰ 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 $w(p) = (p - (1 - p)\gamma)^+$.²¹

¹⁹This 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 $\alpha \geq 2$.

²⁰In case $w(p)$ is not differentiable, then $w'(p)$ is the right-derivative, which is well-defined in case w is star-shaped.

²¹In our main example, w is (globally) *weakly* star-shaped: that is, the function $p \mapsto w(p)/p$ is only weakly

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 payoff in that round is given by

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

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

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

Hence, splitting is better if and only if

$$\frac{1}{\bar{p} - \underline{p}} \int_{\underline{p}}^{\bar{p}} \frac{w(p)}{p} dp \leq \frac{w(\bar{p})}{\bar{p}}, \quad (7)$$

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

Equation (??) 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? Note that, as mentioned, the payoff from a jump is $\underline{p}w(\bar{p})/\bar{p} - w(\underline{p})$, while the marginal benefit from 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 ?? . 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

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.




Figure 8: Splitting information with an arbitrary outside option

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 ??), the following result should not be too unexpected.

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

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

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

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

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

In our main example, $\text{sha } w(p) = w(p)$ 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.

4.2 Mixed Strategies

The description of the maximum payoff is somewhat more complicated in this case, and we restrict attention in this section to the case in which w is weakly star-shaped.

As discussed in the main example, the (limiting) maximum payoff function V_1 must obey several constraints. Stated equivalently in terms of the function $h = V_1 + w$, it must be the case that:

1. The function h is no steeper than $p \mapsto w(p)/p$, as explained in Subsection ??;²²
2. The function h is bounded above, because V_0 is non-negative:

$$h(p) \leq \bar{h}(p) := 1 - \frac{1-p}{p}w(p);$$

²²Formally, (??) holds for all p_0, p_1 .

3. On any interval in which $h < \bar{h}$, h must be weakly concave;
4. The function h weakly exceeds w , because V_1 is non-negative.

Note that, from the definition of \bar{h} , it is no steeper than $p \mapsto w(p)/p$, because w is star-shaped. Hence, whenever $h = \bar{h}$, h also satisfies the first requirement. In particular, the function \bar{h} satisfies all four requirements, but it is generally not the only function that does (cf. our main example).

The only constraint that might not be obvious is the third, briefly mentioned in Subsection ???. Let us illustrate its necessity via a simple example. Suppose that h is not concave, i.e. there exists $p_1 < p < p_2$ such that

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

Assume, in addition, that $V_0(p_1) > V_1(p_1) - V_1(p_2)$ (this is not implied by $h(p) \leq \bar{h}$, but see below). We construct a bet that strictly improves on $V_1(p)$ with one more period. Suppose that the agent pays $V_1(p_1) - V_1(p_2)$ to the principal if and only if the posterior drops to p_1 , and that play reverts then (or if the posterior belief turns out to be p_2) according to the equilibrium that achieves V_1 . Note that 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, given its prior p ,

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

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 $V_1(p_2)$, so that, adding up,

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

Note that the constraint $V_0(p_1) > V_1(p_1) - V_1(p_2)$ is always satisfied if p_1, p_2 are close enough to p and $V_0(p_1) > 0$, and so h must be locally concave at any p at which $V_0(p) > 0$.²³

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:*

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

4.3 Intermediary

Finally, we come back to the case in which the Agent can commit to mixed actions, perhaps because such mixed actions are observable. The maximum payoff has its simplest expression in this case, and the result does not require to assume that w is star-shaped.

As in the main example, the next theorem is established by considering (local) bets that take advantage of the difference in beliefs between the type-1 Agent and the Firm, and that can be constructed as long as $V_0 > 0$. In this way, the type-1 Agent can extract all the surplus, from the type-0 Agent as well as of 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:*

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

5 Final Remarks

We described ways for self-enforcing contracts based on gradual persuasion/communication and possibly mixed strategies and side bets to help resolve the moral hazard/holdup problem

²³This hinges on continuity of V_1 and V_0 ; V_1 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 follows then from the continuity of V_1 .

in one-shot interaction to acquire information. Clearly, in real-life applications these kinds of contracts can be aided by repeated interactions and reputation building.

Although it may not appear that way at first glance, we claim that the model is quite robust, as we discuss the following extensions.

1. It is an important assumption in our model that we are selling information rather than a service or a physical good, since that allows us to assume that the transaction (i.e. taking the test) does not have any physical cost to the Agent. If taking tests was costly, to preserve results we would need to assume that taking a test is contractible and that flunking the test on purpose is as costly for the type-1 as passing it. If tests were not contractible the standard hold-up problem would apply: in the last round of transactions the Agent would always renege on taking the test and hence the Firm would never pay for it, and the equilibrium would unravel. Che and S kovics (2004) suggests the following solution to this known problem: if we relax the equilibrium concept to be an epsilon-equilibrium and assume that easier tests are proportionally less expensive (for example, if we interpret harder tests as taking many easier tests at once), then splitting tests would allow to resolve that aspect of hold-up and hence gradualism would have an additional and independent benefit for sustaining good equilibria (this rationale resolves the problem of the Agent holding up the Firm, while our results are about resolving the opposite hold-up).
2. In our model it does not really matter that it is the Agent who chooses the difficulty of the test. We would obtain the same results in case the Firm was choosing the difficulty. What is important, however, is that either taking the test itself is not contractible or that the Agent has the option of flunking the test on purpose. The reason is that otherwise the Firm would deviate to paying nothing in all but the last round and then offer an epsilon contingent on the Agent taking the hardest test. It would not be a profitable deviation in our equilibria since there the Agent would respond by taking the money and flunking the test for sure.
3. Suppose that there is discounting with every round of communication, and the Firm can decide to take its investment decision before the K rounds are over. Then it is no longer true

that adding another round of communication will strictly increase the Agent's maximum equilibrium payoff. This is both because the Agent faces a trade-off between collecting more money overall and collecting it earlier, and because the Firm will ultimately prefer to take its outside option rather than wait for another period, once the possible benefits from waiting become too small. Hence, in the best equilibrium, the number of rounds in which communication actually takes place is bounded, so that the exact number of rounds available will be of no importance, provided that there are sufficiently many of them. 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. While this version with discounting does not lend itself to closed-form formulas, it is easy to see that, as the discount factor approaches one, the payoff to the type-1 Agent must tend to the payoff in the undiscounted game. Furthermore, in our leading example, numerical simulations show that for the pure-strategy case this convergence occurs at a geometric rate.

4. Suppose that the Agent cares to some extent that the Firm takes the correct action (say, *ceteris paribus*, her payoff increases by some small $\varepsilon > 0$). Then, in the one-shot game, it is dominant for the Agent to reveal all her information, and so the Firm will not make any payment. This logic clearly extends to the finite horizon game, no matter how long the horizon is. On the other hand, this unraveling argument does not extend to the infinite-horizon game (say, with little but positive discounting), and it is possible to construct equilibria in our leading example in which the Agent is paid for a gradual release of information. Of course, the value of ε does put bounds on how extreme the Firm's posterior belief can become before the Agent discloses all information. Nevertheless, our results are robust, inasmuch as the maximal equilibrium payoff to the type-1 Agent will be continuous at $\varepsilon = 0$ (if we allow for discounting).
5. As we mentioned in the introduction, our gradual tests are related to the literature on zero-knowledge proofs. The main difference between our paper and most of that literature is that we assume that the Firm knows which actions are relevant and with every piece of information its outside option changes, while in these other models although the Firm

becomes more and more certain that the Agent knows that the state is 1, the Firm does not know which action is optimal in that state. For example, our model fits a situation where the Agent may have information about a particular investment while in these other models the Agent has information about some investment opportunity. This is also similar to a model where an inventor shows elements of its invention to the Firm, but unless the Firm learns all elements, it cannot “steal” the idea. Mapping this situation to our model would mean that $w(p)$ is constant for all $p < 1$ and increases discontinuously at 1. That would suggest an immediate, familiar solution: the Agent should reveal almost all details other than the “last key,” increasing the Firm’s posterior close to 1, and then sell just that remaining piece. In our model that does not work since the information is valuable to the Firm *per se*. It is possible that the two models could be much more similar if one assumed that the Firm had done some research as well and may already know how to make some of the elements of the invention. If so, then an inventor would always risk that by showing additional elements to the Firm, she would make herself obsolete. However, since such a model requires some private information on the side of the Firm, it would not be equivalent to our model and the analysis of that situation remains an open question.

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.

- [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

A.1 Proof of Lemma ?? and Lemma ??

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

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.²⁴ Further, the type-1 Agent’s maximal payoff with k rounds to go is equal to

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

²⁴In this proof, when we say that the equilibrium involves setting the posterior belief p_1 , we mean that, from the type-1 Agent’s point of view, the posterior belief will be p_1 , while from the point of view of the Firm, the posterior belief will be a random variable p' with possible values $\{0, p_1\}$.

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

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
t

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 5 *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 6 *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. □

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 Lemma ?? and Theorem ??

We start with the theorem, which implies the lemma.

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 ,

$$\int_{[0,1]} p' dF_k(p'|p) = p, \text{ or } \int_{[0,1]} (p' - p) dF_k(p'|p) = 0. \quad (8)$$

To put it differently, $F_k(\cdot|p)$ is a mean-preserving spread of the Firm's prior belief p .²⁶

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'

²⁶The 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.

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

$$\int_0^1 (w(p') + X') dF_k(p'|p) - w(p) + \int_0^1 V_{1,k-1}(p', X') \frac{p'}{p} dF_k(p'|p), \quad (9)$$

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 (??) and the feasibility of X').

A.3.1 The Optimal Transfers

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

Lemma 7 *Fix the prior belief p and the number of remaining rounds k . The best equilibrium payoff of the type-1 Agent, as defined by (??), 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 understand the function $V_{1,k}(p, X)$. Note that, as observed earlier, we can always assume that the equilibrium is efficient: take any equilibrium, and assume that, in the last round, on the equilibrium path, the type-1 Agent discloses the state. This modification can only relax any incentive (or individual rationality) constraint. This means that payoffs must satisfy (??), 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 (??), 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. □

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

The key intuition here is that the type-1 Agent assigns a higher probability to the event that the posterior belief will be $p' > p$ than does the Firm and conversely, a lower probability to the event that $p' < p$, because she knows that the state is 1. Therefore, the type-1 Agent wants to

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

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 (??), we obtain that

$$V_{1,k}(p) = \sup_{F_k(\cdot|p)} \int_0^1 v_{k-1}(p'; p) dF_k(p'|p) - w(p), \quad (10)$$

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 (??), 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 (??) 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 ??: Recall that the function to be maximized is

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

or re-arranging,

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

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

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

or re-arranging,

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

This gives

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

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

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

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

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

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

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 ?? and Theorem ??

We adapt the arguments from the proof of Theorem ??. 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

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

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 ???. 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 ??), so that h^* is no steeper than w . That is, h^* satisfies all four constraints from Section ??.

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 $p_1 < p < p_2$.²⁷ 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

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

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, since h is no steeper than $w(p)/p$,

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

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).

²⁷Note 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.