

Selling Information

Johannes Hörner

Yale University

Andrzej Skrzypacz

Stanford University

A firm considers hiring an agent who may be competent for a potential project or not. The agent can prove her competence but faces a holdup problem. We propose a model of persuasion and show how gradualism helps mitigate the holdup problem. We show when it is optimal to give away part of the information at the beginning of the bargaining and sell the remainder in dribs and drabs. The agent can appropriate only part of the value of information. Introducing a third party allows her to extract the maximum surplus.

I. Introduction

A firm contemplates hiring an expert. For example, it considers hiring a consulting company to lead a potential reorganization of its division; an

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 Graduate School of Economics (JOCS), the Collegio Carlo Alberto, Turin, Essex University, the European University Institute, Harvard–Massachusetts Institute of Technology, University of Illinois at Urbana–Champaign, Northwestern University, Oxford University, Simon Fraser University, Stanford University, University of British Columbia, University of Western Ontario, University of California, San Diego, Yale, X-HEC Paris, Paris Sciences Economiques, Toulouse School of Economics, the European Summer Symposium in Economic Theory at Gerzensee 2012, the Society for Economic Dynamics 2009 meeting, the Stony Brook 2010 game theory conference, and the 2011 Asian meeting of the Econometric Society for useful comments and suggestions.

Electronically published October 31, 2016
[*Journal of Political Economy*, 2016, vol. 124, no. 6]
© 2016 by The University of Chicago. All rights reserved. 0022-3808/2016/12406-0001\$10.00

advisor to help with an acquisition; or an inventor to help build a new product feature, or a product itself, in the case of an agency engaged in a procurement process. The agent has private information whether she is competent (type 1) or incompetent (type 0). Hiring an incompetent agent leads to a negative expected net present value (ENPV) from the project. Hiring a competent agent yields a positive ENPV. While the firm could hire the expert/agent by offering her a payment equal to her outside option, the firm's problem is that it is uncertain whether hiring the agent and going forward with the project is the best course of action. The firm could make its decision on the basis of the prior belief, but taking into account the agent's information would result in efficiency gains.

How could the agent persuade the firm that she is competent and at the same time be rewarded for competence? How should the agent sell that information optimally so that the firm would make an informed hiring decision and the agent would be the most motivated to acquire competence?

To answer these questions, we propose a game-theoretic model of gradual persuasion between the agent and the firm. Persuasion can be thought of as preliminary projects or temporary employment spells (which either side can revoke at any time). During the trial period the agent is paid and gradually reveals information. In particular, the agent performs tasks (which we call "tests") that are informative about her type. In every period, the firm observes the task chosen by the agent, and its outcomes, and then decides whether to continue the trial period, fire the agent, or hire her and start the main project. Neither the firm nor the agent can commit to future payments, hiring decisions, test taking (choice of tasks), or information sharing. There is neither commitment nor contractibility. Most importantly, the firm cannot commit to pay the agent contingent on the information she reveals. In practice, there are obvious ways in which the two sides can commit, more or less formally. The agent might be concerned about her reputation for competence, for instance. The legal structure might provide some protection to intellectual property rights (as patents do).

Motivated by incentives to acquire competence, we characterize equilibria that maximize the difference in expected payoffs of the agent of type 1 and type 0 (as we show, this is equivalent to solving for equilibria that maximize the payoffs of the competent agent).

With one-time information transmission (whether information disclosure is full or partial), both types of the agent would get the same payment. Maximizing incentives to acquire competence calls for multistage communication. The game we consider is therefore dynamic: there is a multistage information selling stage (which we call the game of information sale) followed by the final decision of the firm, which depends on the information the firm has acquired. Within a round of communica-

tion, the two parties make voluntary monetary transfers, and then the agent can disclose some information to the firm.

We assume that information is *verifiable* and *divisible*. In particular, the information transmission is modeled as tests to verify the agent's type. Verifiability of information means that each test has a known difficulty: the competent agent can always pass it (in the baseline model), but the incompetent agent passes it with a probability commensurate to its difficulty. Easier tests have a higher probability of being passed by an incompetent agent. Divisibility of information means that there is a rich set of available tests of varying difficulties.

Our main result is that, very generally, "splitting information over time" increases the competent agent's payoff. That is, the competent agent's payoff is higher in equilibria in which she takes two tests (and is paid for them) in a sequence than if she takes both of them at once (which is equivalent to taking one harder test). The intuition is that when we split the tests, the firm still pays the same amount on average; but the incompetent agent may fail the first test, so that type gets a smaller fraction of the expected payment.

That effect explains the structure of the best equilibrium in our first model: first, an initial chunk of information is given away for free that leads the firm to the utmost confusion regarding whether hiring the agent is the correct decision. If the agent passes this first test, she is then hired on a "temporary basis"; and during the trial, as long as the agent performs, she sells information in dribs and drabs and gets paid a little for each bit (and a failure of a test leads to a termination). These features seem consistent with the practice of preliminary reports/trial periods before firms make large financial commitments to projects with the help of external experts. Our finding that selling information gradually is beneficial to reward competence of the seller should come as no surprise to anyone who was ever involved in consulting. The free first consultation is also reminiscent of standard business practice.

We derive a tight bound on the competent agent's equilibrium payoff as the number of possible communication rounds grows to infinity. Although the closed-form expression for the limit payoff relies on the divisibility of information and the arbitrary number of rounds that we allow, there is no discontinuity: the benefit of splitting information does not depend on either assumption.

While the acquisition of competence is our leading interpretation for the model, we also develop a more general version that could be useful for other applications. For example, the firm may be buying the agent's advice (not needing the agent's help to execute the project) and the agent could be privately informed about a state of the world, which in turn affects the ENPV of the project. The firm's action set does not need to be binary either (e.g., the firm may be choosing a size of its investment

in the project). To model such situations, we allow for more general payoffs if the firm decides to make decisions without further information from the agent. In particular, these payoffs capture the possibility that revealing information the agent has about the state makes it easier for the firm to make good decisions without further agent's input. We show that the benefits of "splitting information" apply to this more general model and characterize the best equilibrium payoffs for type 1. We prove that selling information in small bits is profitable as long as the firm's payoff function (mapping the current posterior belief to the firm's ENPV if it was to make a decision without any further information from the agent) is star shaped, that is, as long as its average value is increasing in the belief.

Splitting information might help the agent in extracting more surplus from the relationship, but it does not suffice to extract the entire surplus. In the last section of the paper we show how enlarging the class of "test technologies" to include some that involve noisy communication between the agent and the firm can help. In particular, we show that, even with our extreme assumptions of noncontractibility and noncommitment, with rich enough tests, the competent agent can extract the entire expected value quite generally.

Related literature.—The paper is related to the literature on holdup, 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 can depend on the future pieces that are disclosed (or failed to be). An agent that eventually reveals himself to be incompetent by the same token destroys the value of the consideration that he had established. This property of beliefs—that they can go up or down—is an essential ingredient for our results. 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.

The formal maximization problem and, in particular, the structural constraints on information revelation are reminiscent of the literature on long cheap talk. See, in particular, Forges (1990) and Aumann and Hart (2003) and, more generally, Aumann and Maschler (1995). As is the case here, the problem is how to "split" a martingale optimally over time. That is, the firm's belief is a martingale, and the optimal strategy specifies its distribution over time. Ely, Frankel, and Kamenica (2015) is another analysis of optimal martingale splitting, although information has instrumental value there.

There are important differences between our paper and the motivation of these papers, however. First and foremost, in contrast to that lit-

erature, payoff-relevant actions are taken before information disclosure is over, since the firm pays the agent as information gets revealed over time, so communication and payments are concurrent here. In fact, with a mediator, the agent also makes payments to the firm during the communication phase. As in Matthews and Postlewaite (1995) or 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. Pure 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 (2013) consider the problem of information exchange between two informed parties in a repeated game without transfers and establish a folk theorem. In all these papers, the focus is on identifying the best equilibrium from the agent's perspective in the *ex ante* sense, before her type is known. In our case, this is trivial and does not deliver differential payoffs to the agents' types (i.e., a higher payoff to the competent type).

The martingale property is distinctive of information, and this is a key difference between our setup 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; Marx and Matthews 2000; Compte and Jehiel 2004). 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, in contrast to contribution games.) The opportunity cost of giving information away is a function of the equilibrium to be played. So, in contrast to public goods game, the marginal (opportunity) cost of information is endogenous. Relative to sales of private goods, the marginal value of information cannot be ascertained without considering the information as a whole, very much as for public goods.

Our focus (proving one's knowledge) and instrument (tests that imperfectly discriminate for it) are reminiscent of the literature on zero-knowledge proofs, which also stresses the benefits of repeating such tests. This literature, which 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, in contrast to 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.

Our leading motivating example of an agent that tries to convince a firm to hire her is closely related to that of the standard model of signaling, as in Spence (1973). There are three differences. First, in terms of technology, in job market signaling, all messages can be sent by all types, but at different costs. Here, instead, messages are free, but the message space is type dependent. Second, in terms of market structure, in standard job market signaling, it is usually assumed that firms compete for the worker, who reaps the entire surplus, so the issue of how much surplus signaling allows her to appropriate does not really come up. Third, in terms of objective function, our goal is to understand what equilibrium maximizes the payoff difference between the competent and incompetent types, a question that is not usually central to the analysis of signaling. These are not major differences, however: type-dependent messages can be viewed as actions that are prohibitively costly to some types;¹ perfect competition among firms does not seem like an essential ingredient of Spence's analysis, and likewise for the fact that we ignore here any outside option for the agent. The major difference, in our view, is that we are interested in how persuasion should be dynamically structured, to provide maximum rewards for (and so incentives to acquire) competence. With the exception of Forges (1990), who analyzes a very different game, we are not aware of an analysis of longer communication or signaling in a job market-related setting.

Less related are some papers in industrial organization. Our paper is complementary to Anton and Yao (1994, 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, 1990). Esó and Szentes (2007) take a mechanism design approach to this problem, while Kamenica and Gentzkow (2011) apply ideas similar to those in 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.

¹ This interpretation of our tests does not quite work because the incompetent agent does not know whether she will be lucky or not in passing a given test.

II. Hiring Competent Experts and Motivating Competence

We start with a simple model in which we explicitly derive how the firm's payoff of hiring the agent as a function of the firm's beliefs about her competence arises from a decision problem.

A. Setup

There is one agent (she) and a firm (it). The agent is of one of two possible types: $\omega \in \Omega := \{0, 1\}$; she is either competent (1) or not (0). We also refer to these as type 1 and type 0. The agent's type is private information. The firm's prior belief that $\omega = 1$ is $p_0 \in (0, 1)$. In Section II.G, we discuss how this belief might come about.

The game is divided into K rounds of communication (we focus on the limit as K grows large), followed by an action stage. In the action stage the firm must choose either action I (hire the agent and "invest" in the project) or action N (not hire the agent and "not invest" in the project). Not investing yields a payoff of 0 independently of the agent's type. Investing yields an expected net payoff of 1 if $\omega = 1$ and $-\gamma < 0$ if $\omega = 0$ (net of the costs of hiring the agent). That is, investing is optimal if the agent is competent, as such an agent has the skill, know-how, or information to make the investment thrive. However, if the agent is incompetent, it is safer to abstain from investing.

In each of the K rounds of communication timing is as follows. First, the firm and the agent choose a monetary transfer to the other player, t_k^A and t_k^F , respectively.² Second, after these simultaneous transfers, the agent chooses whether to reveal some information by undergoing a test.³

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

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

² The reader might wonder why we allow the agent to pay the firm. After all, it is the agent who owns the unique valuable good, information. Such payments will turn out to be irrelevant given the testing technology in this section but will play a role in the second part of the paper with more complex communication.

³ Nothing hinges on this timing. Payments could be made sequentially rather than simultaneously and occur after rather than before the test is taken.

⁴ Because the level of difficulty is determined in equilibrium, it does not matter that the agent chooses it rather than the firm.

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

In words, by disclosing information, the agent affects the firm's belief that she is competent. Persuasion can be a gradual process: after the agent discloses some information, the firm's posterior belief p' can be arbitrary, provided that the prior belief p is not degenerate. But the firm uses Bayesian updating. Viewed as a stochastic process whose realization depends on the disclosed information, the sequence of posterior beliefs is a martingale from the firm's point of view.⁵

B. Histories, Strategies, and Payoffs

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

$$h_k = \{(t_k^A, t_k^F, m_k, r_k)\}_{k=0}^{k-1},$$

where

$$(t_k^A, t_k^F, m_k, r_k) \in \mathbb{R}_+^2 \times [0, 1] \times \{0, 1\}.$$

Here, m_k is the difficulty of the test chosen by the agent in stage k and r_k is the outcome of that test (which is either positive, 1, or negative, 0). The set of all such histories is denoted H_k (set $H_0 := \emptyset$).

A (behavior) strategy σ^F for the firm is a collection $(\{\tau_k^F\}_{k=0}^{K-1}, \alpha^F)$, where (i) τ_k^F is a probability transition $\tau_k^F: H_k \rightarrow \mathbb{R}_+$, specifying a transfer $t_k^F := \tau^F(h_k)$ as a function of the (public) history so far, as well as (ii) an action (a probability transition as well), $\alpha^F: H_K \rightarrow \{I, N\}$ after the K th round. A (behavior) strategy σ^A for the agent is a collection $(\{\tau_k^A, \mu_k^A\}_{k=0}^{K-1})$, where (i) $\tau_k^A: \Omega \times H_k \rightarrow \mathbb{R}_+$ is a probability transition specifying the transfer $t_k^A := \tau^A(h_k)$ in round k given the history so far and given the information she has, and (ii) $\mu_k^A: \Omega \times H_k \times \mathbb{R}_+^2 \rightarrow [0, 1]$ is a probability transition specifying the difficulty of the test (i.e., the value of m), as a function of the

⁵ More abstractly, as in the literature on repeated games with incomplete information, we can think of an agent's strategy as a choice of a martingale—the firm's beliefs—given the consistency requirements imposed by Bayes's rule. For concreteness, we model this as the outcome of a series of tests whose difficulty can be varied. Alternatively, we may think of information as being divisible and the agent choosing how much information to disclose at each round; the incompetent agent might be a charlatan who might be lucky or skilled enough to produce some persuasive evidence. For now, we do not allow the competent agent to flunk the test on purpose, nor do we consider tests so difficult that even the competent agent might fail them. This means that the sample path of the belief martingale is either nondecreasing or absorbed at zero. We discuss richer communication possibilities in Sec. IV.

agent’s type, the history up to the current round, and the transfers that were made in the round. All choices are possibly randomized, but in this and the next section we restrict attention to pure-strategy equilibria.

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

In terms of payoffs, we assume there is neither discounting nor any other type of frictions during the K rounds (e.g., taking the tests is free). We discuss frictions in the next section.

In the absence of any additional information revelation, the firm’s optimal action is to invest if and only if its belief p that her type is 1 satisfies

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

Hence, its payoff from the optimal action is given by

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

where $x^+ := \max\{0, x\}$. While our analysis covers the cases in which the prior belief p_0 is below or above p^* , we have in mind the more interesting case in which p_0 is smaller than p^* . The payoff $w(p)$ is the firm’s *outside option*. Since we assumed that the firm makes a binary investment decision, the specific outside option reduces to a call option. We consider a richer class of outside option specifications in the next section.

To focus attention on the communication stage of the game, we assume for now that if the firm decides to invest, it hires the agent at a salary equal to her outside option (we consider the perhaps more realistic case in which the agent strictly prefers to be hired in Sec. II.F). As a result, the agent does not care about the firm’s investment decision or about its own competence per se. All she cares about is getting paid as much as possible over the K rounds of communication.

The firm cares about the payoff from the investment decision, net of any payments to the agent during the communication stage (recall that the investment payoffs already include the cost of hiring the agent).

Given some final history h_K (which does not include the firm’s final action to invest or not), the type ω agent’s realized payoff is the sum of all net transfers over all rounds, independently of her type:

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

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

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

If instead the firm decides to invest, it receives

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

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

A prior belief p_0 and a strategy profile $\sigma := (\sigma^F, \sigma^A)$ define a distribution over $\Omega \times H_K \times \{I, N\}$, and we let $V(\sigma)$, $W(\sigma)$, or simply V , W , denote the expected payoffs of the agent and the firm, respectively, with respect to this distribution. When the strategy profile is understood, we also write $V(h_k)$, $W(h_k)$ for the players' continuation payoffs, given history h_k . We further write V_0 , V_1 for the payoff to the agent when we condition on the type $\omega = 0, 1$.

The solution concept is perfect Bayesian equilibrium, as defined in Fudenberg and Tirole (1991, definition 8.2).⁶ We assume that players have access to a public randomization device at the beginning of each round (a draw from a uniform distribution), as this facilitates an argument in a proof. The best equilibrium that we identify (in this and later sections) turns out not to take advantage of this device, so that results do not depend on it.

C. Preliminary Remarks

This game admits a plethora of equilibria, but our focus is on identifying the best equilibrium for the competent agent. It is not difficult to motivate our interest in this equilibrium. After all, rewarding agents for their expertise is socially desirable if acquiring it is costly. Clearly, there are many ways to model acquisition of competence; the online appendix provides a particular example.

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

⁶ Fudenberg and Tirole define perfect Bayesian equilibria for finite multistage games with observed actions only. Here instead, both the type space and the action sets are infinite. The natural generalization of their definition is straightforward and is omitted.

fices to make sure that all players receive at least their minmax payoff (0 and $w(p)$) at any time.

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

$$pV_1(h_k) + (1 - p)V_0(h_k) + W(h_k) \leq p. \tag{1}$$

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

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

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

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

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

D. The Best Equilibrium for the Competent Agent

We now turn to the focus of the analysis: what equilibrium maximizes the payoff of the competent agent, and how much of the surplus can she ap-

appropriate? Note that this maximum payoff is nondecreasing in K , the number of rounds: players can always choose not to make transfers or disclose any information in the first round. Hence, for any p_0 , the highest equilibrium payoff for the type 1 agent has a well-defined limit as $K \rightarrow \infty$ that we seek to identify.

In this and the next section we consider equilibria in which the competent agent always passes any test she takes. She is not allowed to “flunk” the test on purpose, a possibility that we will allow in Section IV: there we enrich the description of the game to allow the agent to choose whether to pass the test after she chooses the difficulty m . In that richer game the analysis in this section is equivalent to restricting attention to pure strategies (a restriction that we recall in formal statements).

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

Instead of describing the information part of an equilibrium outcome by the tests taken so far $\{m_k\}$ and their results, we can equivalently describe it by *martingale splitting*, that is, the sequence of the firm’s beliefs that the type is 1, conditional on all tests so far. As long as the agent passes the tests, the firm’s equilibrium beliefs follow a nondecreasing sequence $\{p_0, \dots, p_{K+1}\}$ starting at the firm’s prior belief, p_0 , and ending at $p_{K+1} = 1$ (assuming, without loss, that the equilibrium is efficient). If the agent fails a test, the posterior drops to zero.

An equilibrium must also specify payments. It turns out that the type 1 agent’s payoff decreases if the firm is ever granted any payoff in excess of its outside option. On the one hand, the agent could demand more in earlier rounds by promising to leave some surplus to the firm in later rounds. On the other hand, the willingness to pay of the firm for this fu-




FIG. 1.—A feasible action

ture surplus is lower than the cost of such a promise to the type 1 agent. The reason is that the firm assigns a lower probability than the type 1 agent to the posterior increasing (and payments once the posterior drops to zero are not individually rational). Finally, it is not hard to see that there is no point here in having the agent make any payments. To sum up, if the firm’s belief in the next round is either p_{k+1} or 0, given the current belief p_k , then the equilibrium specifies that the firm pays her willingness to pay

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

where p' is the (random) belief in the next round, with possible values zero and p_{k+1} , and $\mathbb{E}_f[\cdot]$ is the expectation operator for the firm.

This leaves us with the determination of the sequence of posterior beliefs.⁷

E. A Geometric Analysis

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

Consider first the case $K = 1$. In this case, the highest equilibrium payoff to the type 1 agent is indeed equal to $p_0 - w(p_0)$. Suppose that a successful test takes the posterior to $p_1 \geq p_0$. Using the martingale property, it must be that the probability that the posterior is p_1 is p_0/p_1 , because

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

Hence, the firm is willing to pay

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

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

Note that, when $p_0 \leq p^*$, $w(p_0) = 0$, and the highest payoff in one round that the type 1 agent can achieve is the prior p_0 , which is increasing in $p_0 \leq p^*$. This suggests one way to improve on the payoff with two rounds. In the first round, the agent takes a test whose success leads to a posterior of p^* for free. Indeed, the firm is not willing to pay for a test

⁷ In addition, an equilibrium must also specify how players behave off the equilibrium path. As discussed, the most effective punishment for deviations is reversion to the babbling equilibrium, and this is assumed throughout unless mentioned otherwise.

that does not affect its outside option. In the second round, the equilibrium of the one-round game is played, given the belief p^* . This second and only payment yields

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

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

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

Is the splitting that we described optimal with two periods to go? As it turns out, it is so if and only if $p_0 < (p^*)^2$. But there are many other ways of splitting information with two periods to go that improve on the one-round equilibrium and, among them, splits that also improve over the one-period equilibrium when $p_0 > p^*$. The optimal strategy is given at the end of this subsection.

Can we do better with more rounds? Consider figure 3. As shown on the left panel, information is revealed in three stages. First, the belief is split into 0 and p^* . Second, at p^* (assuming this belief is reached), it is split in 0 and p' . Finally, at p' , it is split in 0 and 1. The right panel shows how to determine the type 1 agent's payoff. The two solid (vertical) segments represent the maximum payments that can be demanded at the second and third stages. (No payment is made in the first, as the splitting does not affect the firm's outside option.)

To understand these payments, note that in each round the agent is paid the difference between the expectation of the value of the outside




FIG. 2.—Revealing information in two steps




FIG. 3.—Revealing information in three steps: evolution (*left*) and payoff (*right*)

option tomorrow and the current outside option. The expectation is weighted by the probabilities of each posterior belief: geometrically, it is the value, at the current belief, of the line that connects the outside option at the two possible posterior beliefs tomorrow (so 0 and p' for the first positive payment and 0 and 1 for the second).

Thus, the added lengths of the solid vertical segments measure the type 1 agent's payoff. Compare with our two-stage equilibrium, in which all information is disclosed once the belief reaches p^* . The payment for our three-stage equilibrium is measured the same way: it is the length of the vertical segment connecting the outside option at p' to the point on the line connecting the outside options at 0 and 1. So it is the sum of the left vertical segment and the dotted segment above it. As is clear, the right solid segment exceeds the left dotted segment: the payoff with three stages must be larger, as the chords from the origin to the point $(p, w(p))$ become steeper as p increases. With three steps as depicted, the type 1 agent reaps two payments that add to more than the one payment with two steps.

Intuitively, the firm is willing to pay more with three steps because it is less likely to have to make all these payments if it is dealing with the type 0 agent. By the time the posterior belief reaches p' , by definition, there is a chance that the type 0 agent has been found out; hence, the second payment is a conditional payment, and because the conditioning event is suggestive that the agent is indeed competent, the firm is willing to pay more in total when some of the payments are conditional.

We could go on: information splitting is beneficial. Figure 4 illustrates the total payoff that results from a splitting that involves many small steps (which is the sum of all vertical segments).

Does it follow that the competent agent extracts the maximum value of information as $K \rightarrow \infty$? Unfortunately, no: see the right panel of fig-




FIG. 4.—Revealing information in many steps (*left*); forgone profit at each step (*right*)

ure 4. As the firm’s belief goes from $p - dp$ to p , the firm must pay (using the martingale property, the test must be passed with probability $(p - dp)/p$)

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

yet its outside option increases by $w(p) - w(p - dp)$. The type 1 agent gives up the difference $w(p)dp/p$ in the process. This forgone payoff need not be large when the step size dp is small; but then again, the smaller the step size, the larger the number of steps involved. Note that this forgone payoff does not benefit the firm, which is always charged its full willingness to pay. The type 0 agent reaps this payoff. As a result, her payoff does not vanish, even as $K \rightarrow \infty$.

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

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

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

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

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

LEMMA 1. Consider the binary-hiring model: $w(p) = [p - (1 - p)\gamma]^+$. As $K \rightarrow \infty$, the maximum payoff to the type 1 agent in pure strategies tends to, for $p_0 < p^*$,

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

If $p_0 \geq p^*$, $V_1^p(p_0) = -\gamma \ln p_0$.⁸

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

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

LEMMA 2. The maximum equilibrium payoff of the type 1 agent in a game with K rounds is recursively given by

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

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

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

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 $p_0 < (p^*)^{K/(K-1)}$ 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 $(p^*)^{(K-1)/K}$, $(p^*)^{(K-2)/K}$, ..., 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^* ,

⁸ The superscript p refers to the restriction to pure strategies.




FIG. 5.—Revealing information in many steps (*left*); payoff as a function of K (*right*). Color version available as an online enhancement.

given the prior belief $p_0 < p^*$, followed by a succession of very small disclosures, leading the firm's belief gradually up all the way to one. The right panel of figure 5 shows how the payoff varies with K .

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

F. Agent Prefers to Be Hired

So far, we have assumed that the firm can hire the agent at her outside option, so that the agent is indifferent whether the firm hires her or not. A perhaps more realistic assumption is that the agent's compensation is strictly higher than her outside option, so that she strictly prefers being hired. How does that affect our analysis?

Assume that the agent's surplus from employment is smaller than the expected net losses the firm would incur if it hired an incompetent expert, so that not investing remains the efficient action if the agent is incompetent (otherwise, investing would be optimal in both states and communicating competence would not be important).

In this case the equilibrium that maximizes V_1 (or $V_1 - V_0$) has the same equilibrium path as described above. The only difference is in the off-path behavior supporting it.

First, in our construction above, deviations are punished by a babbling equilibrium. That means that if $p \geq p^*$ and the agent deviates, the firm still hires her. When the agent strictly prefers being hired, this no longer suffices. Once p is sufficiently high, the incompetent type would prefer not to take any more tests (even if she gets compensated for them) to avoid the risk of losing employment. In particular, she would not take the last test in round K that is fully revealing. To sustain our equilibrium outcome we can use the following continuation equilibrium: if the agent ever fails to take the test she is expected to take, the firm's belief about her competence drops to 0, and the babbling equilibrium is then played. This clearly provides incentives for the agent to take the prescribed tests.

Second, one may worry that the firm would not make any payments to the agent knowing that she wants to be employed and hence has strict incentives in the last period to reveal enough information to get employed. This creates no difficulty: our equilibrium calls for the agent to take for free a test in the first period so that the firm's beliefs increase to p^* . After that, it is always a best response for the firm to hire her. So if the firm deviates to not paying for future tests, the deviation to a babbling equilibrium (in which the agent reveals no more information and is still hired at the end) remains incentive compatible.

Positive employment surplus has two additional effects on the equilibrium that are worth pointing out. First, it increases $V_1 - V_0$, because only the competent agent captures the employment surplus. Since our equilibrium is separating, leaving employment surplus to the competent agent strengthens incentives for acquisition of competence. Second, one may be worried that in our original construction the agent is indifferent between revealing and not revealing the last piece of information in round K , because some considerations left out of the model might break this indifference, leading to unraveling. With a positive employment surplus the agent has strict incentives to take the last, fully revealing test, because otherwise the firm would think that she must be "hiding something" and not hire her. At the same time, it does not lead to unraveling because the strict incentives hold only on the equilibrium path. If the firm deviates by not paying for some of the tests, the agent would still be hired, so it would be rational not to take the last test (the competent agent would be indifferent and the incompetent agent strictly prefers to follow the punishment). For further discussion of robustness, see Section III.

G. *Free Entry*

The prior belief p_0 plays an important role in the analysis, so it is worth discussing how it might come about. One natural way of endogenizing it is to explicitly model a prior stage in which agents must decide on whether

to invest in competence, so that the benefit of doing so, which depends on p_0 , must equal its cost. As a result, the fraction of agents doing the necessary investment gives rise to the value of p_0 that makes precisely this fraction willing to do so.

More directly, we may assume that there is a unit mass of agents (on the short side of the market), whose choice is to enter either as incompetent for free or as competent at a cost of $c \in (0, 1)$. There are plenty of equilibria, including one in which only incompetent agents enter and $p_0 = 0$.⁹ The best equilibrium from a social point of view is the one maximizing p_0 ; indeed, each competent agent generates a surplus of 1 at a cost of c . But note that it is not an equilibrium for every agent to enter as a competent one, as the resulting belief $p_0 = 1$ would strip them from any additional rewards from further persuading the firm of their competence. In the best equilibrium, both types of agents must enter, and so the constraint $V_1(p_0) - V_0(p_0) = c$ must hold. Maximizing p_0 thus means selecting the equilibrium that maximizes the difference $V_1 - V_0$, as this is the one for which it is possible to pick the highest p_0 satisfying the constraint. This is precisely the equilibrium that we have characterized.

Hence, the best equilibrium yields a prior belief p_0 that solves

$$\begin{aligned} V_1(p_0) - V_0(p_0) &= c, \\ p_0 V_1(p_0) + (1 - p_0) V_0(p_0) + w(p_0) &= p_0. \end{aligned}$$

If $c \in (-\gamma[1 + (1 + \gamma) \ln p^*], -\gamma \ln p^*)$, the unique solution has $p_0 = -[(c + \gamma \ln p^*)/(1 - c)] \in (0, p^*)$.¹⁰ If, on the other hand, $c \leq -\gamma[1 + (1 + \gamma) \ln p^*]$, the unique solution lies in $(p^*, 1)$. As one would expect, a lower cost leads to a higher prior belief that an agent is competent, as it is then cheaper for him to enter.

III. General Outside Options

Assuming that the outside option is given by a call option, as in our main example, provides a simple illustration of the benefits of splitting, as well as closed-form expressions. However, the analysis can be generalized.

Such a generalization has two benefits. First, it clarifies what drives the benefits of splitting information. Second, it encompasses a broader class of applications. Plainly, there is no reason to confine ourselves to binary decisions by the firm. For instance, the firm might choose between hiring the agent if it deems it profitable, remaining in an arm's-length relationship with the agent, or stopping the relationship altogether.¹¹

⁹ A small cost of entering as incompetent would eliminate this equilibrium.

¹⁰ If $c \geq -\gamma \ln p^*$, there is no equilibrium with competent entrants.

¹¹ We thank a referee for suggesting such a ternary decision, as well as useful variations.

Hence, it makes sense to assume that the firm’s outside option depends on its belief in the agent’s competence in arbitrary ways. Convexity seems to be a natural property to impose, but as we show, an even weaker condition suffices to generalize our results.

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 nonnegative continuous function $w(p)$, and normalize $w(0) = 0, w(1) = 1$.¹² We further assume that $w(p) \leq p$, for all $p \in [0, 1]$, for otherwise full information disclosure is not socially desirable. This payoff can be thought of as the reduced form of some decision problem that the firm faces, as in our baseline model. In that case, w must be convex, but since it is a primitive here, we do not assume so.

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

The answer, as it turns out, depends on the shape of the outside option. It is in the interest of the type 1 agent to split information as finely as possible for any prior belief p_0 if and only if the function w is (strictly) *star shaped*, that is, if and only if the average, $w(p)/p$, is a strictly increasing function of p .¹³ 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.¹⁴ 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)

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

¹² Without loss we may assume w to be nondecreasing.

¹³ This condition already appears in the economics literature in the study of risk (see Landsberger and Meilijson 1990). It is weaker than convexity.

¹⁴ In case $w(p)$ is not differentiable, then $w'(p)$ is the right derivative, which is well defined in case w is star shaped.

which generalizes the formula that we have seen for the special case $w(p) = [p - (1 - p)\gamma]^+$.¹⁵ 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 jumps from \underline{p} to \bar{p} ; the payment from the firm to the agent in that round is given by

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

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

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

Hence, splitting is better if and only if

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

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

Equation (2) also explains why splitting information finely is not a good idea if the average outside option is strictly decreasing over some range $[\underline{p}, \bar{p}]$, as the inequality is reversed in that case. What determines the jump? As mentioned, the payoff from a jump is $\frac{\bar{p}}{\underline{p}} w(\bar{p}) - w(\underline{p})$, 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

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

See figure 6. The left panel illustrates how having two rounds improves on one round. Starting with a prior belief p_0 , the highest equilibrium payoff the type 1 agent can receive in one round is given by the dotted black segment. If instead information is disclosed in two steps, with an intermediate belief p_1 , the type 1 agent's payoff becomes the sum of the two solid (vertical) segments, which is strictly more, since $w(p)/p$ is strictly increasing. The right panel illustrates the jump in beliefs that oc-

¹⁵ In our main example, w is (globally) weakly star shaped; i.e., the function $p \mapsto w(p)/p$ is only weakly increasing. The formula for the maximum payoff in the limit $K \rightarrow \infty$ is the same whether there is a jump in the first period or not. But for any finite K , splitting information disclosures over the range $[p_0, p^*]$ is suboptimal, as it is a "wasted period," whose cost vanishes only in the limit.




FIG. 6.—Splitting information with an arbitrary outside option. Color version available as an online enhancement.

curs 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 nonnegative 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 the right panel of fig. 6), the following result should not be too unexpected.

THEOREM 1. The maximum equilibrium payoff to the type 1 agent in pure strategies as $K \rightarrow \infty$ tends to

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

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 that one applies it to the largest weakly star-shaped function that is smaller than w . In words, the maximum payoff to the type 1 agent is the area between the marginal and the average outside option of the firm, after “regularizing” this outside option by considering the largest weakly star-shaped function below it.

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

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

where $\text{cl } A$ is the closure of set A . 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.

Frictions.—The only friction assumed so far has been the finiteness of the horizon. Clearly, this is a simplification. In practice, delaying the action has a cost in terms of discounting. Taking tests can entail intrinsic costs as well. Finally, the agent is unlikely to be wholly unconcerned by the firm’s action. She might be able to make money out of her information elsewhere; in this case, she would balk at giving away for free the last bit. Or, to the contrary, she might have a slight preference for the firm taking the right action, all else being equal; she would then not resist giving away this last bit for free if this was the only way to prevent the firm from making a mistake. The firm might then be tempted to forgo the payments in hopes that this occurs.

Considering each friction one by one, it is not hard to see that our results are robust to small perturbations. First, suppose that every additional round of communication is discounted, with some common discount factor $\delta \in (0, 1)$. Plainly then, there is no benefit in having arbitrarily many rounds. The reason is that the agent faces a trade-off between collecting more money overall and collecting it earlier, and the firm ultimately prefers taking its outside option rather than waiting for another period, once the benefits from waiting become small. Hence, in the best equilibrium, the number of rounds in which communication actually takes place is bounded. However, as long as the players are not too impatient, the best equilibrium still involves a gradual release of information. It is easy to see that, as discounting vanishes, the payoff to the competent agent must tend to her payoff in the undiscounted game. In the Appendix, we prove the following result.

LEMMA 3. Suppose that w is star shaped and that players discount rounds at rate $\delta \geq 1$. As $\delta \rightarrow 1$, the maximum equilibrium payoff of the type 1 agent in pure strategies tends to the undiscounted limit $V_1^p(p_0)$.

In our main example, it is easy to show that this convergence occurs at a rate that is geometric in $1 - \delta$.

Suppose now that the horizon is infinite (with low discounting), and introduce a cost to the agent taking a test or transmitting/certifying information. (This is equivalent to each bit of information having an op-

portunity cost.) Assume that conveying information is socially efficient; more precisely, assume that the ratio of the cost of taking the posterior belief of the firm from p to $p' > p$ (denoted $c(p, p') \geq 0$) to the benefit $p'w(p')/p - w(p)$ to the firm is bounded below one, uniformly in p, p' .

We can then scale down the information as time passes so as to make sure that the continuation payoff of the two parties incentivizes them to make payments and to take the test.¹⁶

In the case of an intrinsic preference of the agent for the firm taking either the right or a given action (worth, say, a given $\$v(\omega, a) \geq 0$, $a = I, N$, to the agent, where it is assumed that incentives are weakly aligned: $v(1, I) \geq v(1, N)$, $v(0, N) \geq v(0, I)$), the procedure must be modified so that the agent releases a last chunk of information whose value to the firm still exceeds $\delta \max_{\omega, a} v(\omega, a)$; if the due payment takes place, the agent releases this last chunk; if not, the agent postpones releasing this information, expecting the firm to make up for such a careless slip by making the payment in the next round (if and only if the agent did not release this information). See the Appendix for a formal proof of the following result.

LEMMA 4. Suppose that w is differentiable and star shaped and that players discount rounds at rate $\delta \geq 1$. Suppose either costs $c(p, p')$ or intrinsic preferences $v(\omega, a)$ as described. As $\delta \rightarrow 1$ and either $\max_{p < p'} c(p, p') \rightarrow 0$ or $\max_{\omega, a} v(\omega, a) \rightarrow 0$, the maximum equilibrium payoff of the type 1 agent in pure strategies tends to the undiscounted limit (without cost of intrinsic preferences) $V_1^p(p_0)$.

We stress that this robustness applies to small perturbations only. Unraveling certainly applies to our model for certain kinds of perturbations, as it does in related models of contribution games with irreversibility (e.g., Admati and Perry 1991; Marx and Matthews 2000). For instance (and this is certainly not the only possibility), if the agent derives a strictly positive gain from the firm making the right (hiring) decision and there is a finite horizon, the firm can certainly wait until the last period and get all the information for free.

IV. Noisy Information Transmission

So far, we have assumed that the competent agent always passes the test, which implies that the firm's posterior belief is either nondecreasing or absorbed at zero.

¹⁶ As usual, one need not think of the information sale phase as lasting literally forever: the low discount factor can be thought of as a probability of terminating this phase, and it can be generated by the players themselves, using, for instance, a jointly controlled lottery; in that case, the duration of this phase is (almost surely) finite.

There are two reasons why even the competent agent may fail. First, she may be able to choose to flunk the test (it turns out that such an option may improve on the equilibria considered so far). In practice, it is hard to see what prevents an agent from failing intentionally a given test: software can be crippled or slowed down, prototypes can be damaged or impaired, and imprecise or even incorrect answers can be given. To model this possibility, we add a third dimension to the agent's strategy; namely, in every round, after a test has been privately performed, the agent has the choice, in case of a success, to report a failure. As further notation is not needed, we refer the interested reader to the working paper (Hörner and Skrzypacz 2012) for a formal definition. Because the model considered in Sec. II.D corresponds to the special case in which the competent agent always passes the test—the only interesting pure strategy in the extended model—we refer to this version as the noisy model. Formally, this is the same model as before, but mixed strategies are considered and, as we will see, make a difference.

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

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

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

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

A. *The Value of Lower Posteriors: An Illustration*

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

Suppose that the firm and the agent agree to the following (self-enforcing) scheme. If the test fails, the posterior belief falls to $p - \Delta$, for some $\Delta > 0$. If the test succeeds, the posterior belief jumps to $p + \Delta$. Pick Δ such that $p^* < p - \Delta < p + \Delta < 1$. Such posterior beliefs are achieved by mixing by the type 1 agent (or by a mediator on her behalf), given that the type 0 agent will disclose that the outcome of the test is a success whenever she is lucky. Because the possible posterior beliefs are symmetric around p , the two events (that information gets disclosed or not) must be equally likely from the firm's point of view.

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

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

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

passes the test whenever she is lucky. The left panel of figure 7 illustrates how the mixing works, starting from a given belief $p > p^*$.

Given that the firm pays $X/2$ up-front and that, by construction, the continuation payoff of the type 1 agent is the same whether the posterior belief goes up or down (namely, $\ln(p - \Delta)$), her expected payoff is

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

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

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

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

We generalize the previous scheme to this case: the agent pays $V_1^m(p_1) - V_1^m(p_2)$ to the principal if and only if the posterior drops to p_1 , and play reverts then (or if the posterior belief turns out to be p_2) to the equilibrium that achieves V_1^m . The type 1 agent is indifferent between both pos-




FIG. 7.—Construction of the scheme (left); maximum limit payoff $V_1^m + w, \gamma = 1$ (right). Color version available as an online enhancement.

terior beliefs and so is willing to randomize. Given her assessment of the likelihood of each of these events, the firm is willing to pay up-front

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

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

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

Note that the participation constraint for the incompetent agent, $V_0^m(p_1) > V_1^m(p_1) - V_1^m(p_2)$, is always satisfied if p_1, p_2 are close enough to p and $V_0^m(p_1) > 0$, and so h must be locally concave at any p at which $V_0^m(p) > 0$.¹⁷

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

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

B. Maximum Payoff with Noise and Mediation

1. Noise

First, consider the case of noise—the extension of the baseline model to mixed strategies. Two constraints have been derived on the limiting value of h , the sum of the payoffs of the firm and the competent agent. First,

¹⁷ This hinges on continuity of V_1^m and V_0^m ; V_1^m is continuous because it is always possible to use the same disclosure strategy starting at p_2 as the continuation strategy given p_1 would specify from the first posterior belief above p_2 onward; the first payment must be adjusted, but the continuity in payoffs as $p_1 \rightarrow p_2$ then follows from the continuity of w . Continuity of V_0^m follows from the continuity of V_1^m .

it must be less than $\bar{h} = 1 - (1 - p)w(p)/p$, as implied by feasibility given that the type 0 agent's payoff is nonnegative. Second, on any interval on which $h < \bar{h}$, the function h must be locally concave. There is a third constraint on h that is rather obvious: h must exceed w , the outside option of the firm, as the type 1 agent's payoff is nonnegative.

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

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

or

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

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

DEFINITION 1. Given the outside option w , we define $h^m : [0, 1] \rightarrow \mathbb{R}_+$ as the smallest function such that,

1. for all p , $h^m \leq \bar{h}(p) := 1 - (1 - p)w(p)/p$;
2. if $h^m < \bar{h}$ on $[p_0, p_1]$, then h is locally concave on $[p_0, p_1]$;
3. for all p , $h^m \geq w(p)$;
4. for all $p_1 > p_0$,

$$\frac{h^m(p_1) - h^m(p_0)}{p_1 - p_0} \leq \frac{w(p_1)}{p_1}.$$

As it turns out, equilibrium imposes no additional restriction on h , as we show in the Appendix.¹⁸ What is the smallest function that satisfies these four requirements?¹⁹ In our main example, some algebra gives that

¹⁸ Roughly, any function satisfying these properties cannot be improved on 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.

¹⁹ One might wonder why the smallest function h satisfying the requirements is the appropriate one; the reason is that, starting from the highest equilibrium payoff with one round and applying the two schemes that we have described, we recursively obtain higher values for h as the number of rounds increases, but we cannot "overtake" the smallest function that satisfies the four requirements.

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

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

LEMMA 5. Consider the main example: $w(p) = [p - (1 - p)\gamma]^+$. As $K \rightarrow \infty$, the maximum payoff to the type 1 agent in the model with noise tends to, for $p_0 < p^*$,

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

If $p_0 > p^*$, this limit is

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

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

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

THEOREM 2. Assume that w is weakly star shaped. As $K \rightarrow \infty$, the maximum payoff to the type 1 agent in the model with noise tends to

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

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

2. Mediation

It turns out that a similar reasoning can be used to characterize the maximal V_1 in case a mediator can send noisy messages based on the test results (or the agent has access to noisy tests). The only difference is that the scheme that involves payments by the agent in case the posterior drops is no longer constrained by the indifference of the competent agent, which imposed that h was locally concave whenever h fell short of the upper bound \bar{h} . So we are left with the other three restrictions on the func-

tion h . It turns out, as before, that the solution is given by the smallest function satisfying these requirements. The maximum payoff has a particularly simple expression, and the result does not require w to be star shaped.

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

THEOREM 3. As $K \rightarrow \infty$, the maximum payoff to the type 1 agent with an intermediary tends to

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

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

V. Final Remarks

This paper describes self-enforcing contracts based on gradual persuasion to facilitate sale of information. Clearly, in real-life applications, the mechanism that we describe is limited by the extent to which information is divisible or tests are available. On the other hand, it can be facilitated by repeated interactions and reputation building.

As mentioned in the introduction, our mechanism is reminiscent of zero-knowledge proofs. But gradualism is a technological constraint in this literature. There is no counterpart to the firm's outside option, and the only objective is to convince the other party that the agent holds the information. It is as if $w(p) = 0$ for $p < 1$ and $w(1) = 1$, in which case it is optimal to reveal all details but the "last key," increasing the firm's posterior close to one, and then to sell just that remaining piece. Gradualism arises in our mechanism precisely because the firm's outside option depends on its belief, as is plausible in most economic applications. In fact, often the buyer has private information as well, and an inventor always risks making herself obsolete by revealing additional information to the firm. Considering such a model, in which both parties hold private information, is left for future research.

Appendix

Proofs

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

A. Proof of Lemma 2 and Lemma 1

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

Our induction hypothesis is that, with $k \geq 1$ periods to go and a prior belief $p = p_0$, the best equilibrium involves setting the next (nonzero) posterior belief, p_1 , equal to $p_1 = p^{(k-1)/k}$ if $p^{(k-1)/k} \geq p^*$ (i.e., if $p \geq (p^*)^{k/(k-1)}$ for $k \geq 2$), and equal to p^* otherwise.²⁰ Further, the type 1 agent’s maximal payoff with k rounds to go is equal to

$$V_{1,k}(p) = \begin{cases} k\gamma(1 - p^{1/k}) - [p - \gamma(1 - p)]^- & \text{if } p \geq (p^*)^{k/(k-1)} \\ V_{1,k-1}(p^*) & \text{if } p < (p^*)^{k/(k-1)}. \end{cases}$$

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 $p^{(K-1)/K}$, the relevant specification given that $p^{0/1} = 1 \geq p^*$. 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

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

²⁰ In 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 $\{0, p_1\}$.

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

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

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 $(p/p_1)w(p_1) + V_{1,k}(p_1)$. Note first that, given the induction hypothesis, all values $p_1 \in [p, (p^*)^{k/(k-1)}]$ 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 $p_1 \in [(p^*)^{k/(k-1)}, p^*]$ (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 $p_1 = p^{k/(k+1)}$, 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 $p \geq (p^*)^{(k+1)/k}$.

Therefore, the unique candidates for p_1 are $\{p^*, \max\{p^*, p^{k/(k+1)}\}, 1\}$. Observe that setting the posterior belief p_1 equal to $\max\{p^*, p^{k/(k+1)}\}$ 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).

B. Proof of Theorem 1

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

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

Given a nonnegative function f on $[0, 1]$, let $\text{sha } f$ denote the largest weakly star-shaped function that is smaller than f . This function is well defined because (i) if f_1, f_2 are two weakly star-shaped functions lower than f , the pointwise max-

imum g (i.e., $g(p) := \max\{f_1(p), f_2(p)\}$) is star shaped as well;²¹ and (ii) the limit of a convergent sequence of star-shaped functions is star shaped. Bruckner and Ostrow (1962, theorem 2) also show that a nonnegative star-shaped function (with $w(0) = 0$) must be nondecreasing.

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

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

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

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

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

Let us start by showing that this payoff can be achieved asymptotically (i.e., as $K \rightarrow \infty$). Let J_w denote the complement of I_w , which is a union of disjoint open intervals. Let $\{(p_n^-, p_n^+)\}_{n \in \mathbb{N}}$ denote an enumeration of its endpoints. Finally, let $\check{p}_0 := \min\{p \in I_w, p \geq p_0\}$. Note that, for all n , by continuity of w (using that $\hat{w}(p_n^+)/p_n^+ = \hat{w}(p_n^-)/p_n^-$ 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 $K \in \mathbb{N}$, with $p_0^K = \hat{p}_0$, $p_K^K = 1$, 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 $\hat{p}_0 > p_0$, the type 1 agent gives away the information for free that leads to a posterior \hat{p}_0 ; afterward, 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)

$$\begin{aligned} \sum_{k=0}^{K-1} p_k^K [w^a(p_{k+1}^K) - w^a(p_k^K)] &\rightarrow 1 - \hat{w}(\check{p}_0) - \int_{I_w \cap [p_0, 1]} \hat{w}^a(p) dp \\ &= 1 - \hat{w}(\hat{p}_0) - \int_{\hat{p}_0}^1 \hat{w}^a(p) dp. \end{aligned}$$

²¹ Given $p_1 < p_2$, let $g(p_1) = f_i(p_1)$, $g(p_2) = f_j(p_2)$. Then $g^a(p_2) = f_j^a(p_2) \geq f_i^a(p_2) \geq f_i^a(p_1) = g^a(p_1)$.

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

$$p_0[w^a(p_1) - w^a(p_0)] + p_1[w^a(p_2) - w^a(p_1)] + \dots + p_{K-1} \cdot [w^a(1) - w^a(p_{K-1})]$$

$$= 1 - w(p_0) - (1 - p_{K-1})w^a(1) - \dots - (p_1 - p_0)w^a(p_1),$$

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 6. 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 7. 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 , that is,

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

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

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.

C. Proof of Lemma 3

When there is discounting, the payment made when the current belief is p and the next belief is $p' > p$ (or 0) is given by

$$t(p) = \frac{p}{p'} w(p') - w(p).$$

Here, for simplicity, we assume that the function w is strongly star shaped on the relevant interval of beliefs $[p_0, 1]$; that is, there exists $\nu > 0$ such that, for all $p \in [p_0, 1]$, $pw'(p) - w(p) > \nu$. If as in our main example the function is star shaped but not strongly star shaped, given the prior (it is strongly star shaped if $p_0 > p^*$), an elementary adaptation is required (e.g., in the main example, if the prior p is lower than $p < p^*$, give for free a piece of information such that the posterior p_1 is strictly above but arbitrarily close to p^* , and then follow the same argument).

Fix the prior $p_0 > 0$. We consider a strategy in which p/p' is kept constant in all but a final period of disclosure. That is, assume that $p' = (1 + k)p$, for some $k > 0$ (and all $p < 1/(1 + k)$). Once $(1 + k)p > 1$ given p , information is disclosed in one shot (for the maximum price).

We set $k = \sqrt{1 - \delta}$. The assumption that w is strongly star shaped ensures that there exists $\bar{\delta} < 1$, for all $\delta \in (\bar{\delta}, 1)$, and all $p \in [p_0, 1/(1 + \sqrt{1 - \delta})]$,

$$\delta \frac{p'}{p} w(p') - w(p) = \frac{\delta}{1 + \sqrt{1 - \delta}} w((1 + \sqrt{1 - \delta})p) - w(p) > 0.$$

We write $p^{(0)} = p_0$, $p^{(n)} = (1 + k)^n p_0$, and compute the value function V from this policy:

$$\begin{aligned} V(p_0) &= \frac{\delta}{1 + k} w(p^{(1)}) - w(p) + \delta V(p^{(1)}) \\ &= \frac{\delta w(p^{(1)})}{1 + k} - w(p) + \delta \left[\frac{\delta w(p^{(2)})}{1 + k} - w(p^{(1)}) \right] \\ &\quad + \dots + \delta^T \left[\frac{\delta w(p^{(T+1)})}{1 + k} - w(p^{(T)}) \right], \end{aligned}$$

where T is the number of steps until the belief $1/(1 + \sqrt{1 - \delta})$ is exceeded, that is, $T = \max\{t : (1 + k)^t p \leq 1/(1 + \sqrt{1 - \delta})\}$. Alternatively,

$$T = \left\lfloor \frac{-\ln[(1 + \sqrt{1 - \delta})p]}{\ln(1 + k)} \right\rfloor.$$

Rearranging,

$$\begin{aligned} V(p_0) &= \frac{\delta^{T+1}}{1 + k} w(p^{(T+1)}) - w(p) - \delta \frac{k}{1 + k} w(p^{(1)}) \\ &\quad - \dots - \delta^T \frac{k}{1 + k} w(p^{(T)}) \\ &\geq \frac{\delta^{T+1}}{1 + k} w(p^{(T+1)}) - w(p) - \left[\frac{k}{1 + k} w(p^{(1)}) \right. \\ &\quad \left. + \dots + \frac{k}{1 + k} w(p^{(T)}) \right]. \end{aligned}$$

Note that, for all $k < T$,

$$\frac{p^{(k+1)} - p^{(k)}}{p^{(k+1)}} = \frac{k}{1 + k},$$

so

$$\begin{aligned} V(p_0) &\geq \frac{\delta^{T+1}}{1 + k} w(p^{(T+1)}) - w(p) - \left\{ [p^{(1)} - p^{(0)}] \frac{w(p^{(1)})}{p^{(1)}} \right. \\ &\quad \left. + \dots + [p^{(T)} - p^{(T-1)}] \frac{w(p^{(T)})}{p^{(T)}} \right\} \\ &\geq \frac{\delta^{T+1}}{1 + k} w(p^{(T+1)}) - w(p) - \int_{p_0}^1 \frac{w(p)}{p} dp, \end{aligned}$$

because w is star shaped. To prove that this lower bound converges to $w(1) - w(p_0) - \int_{p_0}^1 [w(p)/p] dp$, as desired, it now suffices to show that $\delta^{T+1}/(1 + \sqrt{1 - \delta}) \rightarrow 1$. This follows from $\delta^{1/\ln(1+\sqrt{1-\delta})} \rightarrow 1$. (Take logs to get $\ln \delta / \ln(1 + \sqrt{1 - \delta})$; use l'Hospital's rule to conclude that this has the same limit as $2[(1 - \delta) - \sqrt{1 - \delta}]/\delta \rightarrow 0$.)

D. Proof of Lemma 4

Part 1: Let us first consider the case of positive costs $c(p, p')$. Modify the construction of the equilibrium used for lemma 3 as follows. Instead of having a final period once $(1 + k)p > 1$, assume that from that round onward, the agent discloses information on that path in each round in a way that leads the firm's belief to follow a strictly increasing sequence $(p_n)_{n \in \mathbb{N}}$, with $p_n < 1$ for all n , and $(p_{n+2} - p_{n+1})/(p_{n+1} - p_n) > 1 - \varepsilon$ for some $\varepsilon > 0$ (take, for instance, $p_n = 1 - h_1/(n^2 + h_2)$ for some large constants $h_1, h_2 > 0$). If the agent ever deviates or the firm fails to pay $\delta(p_n/p_{n+1})w(p_{n+1}) - w(p_n)$ in the n th round along this sequence, equilibrium reverts to babbling, with no payments ever made again. Because $(p_{n+2} - p_{n+1})/(p_{n+1} - p_n) > 1 - \varepsilon$,

$$\frac{w(p_{n+2})/p_{n+2} - w(p_{n+1})/p_{n+1}}{w(p_{n+1})/p_{n+1} - w(p_n)/p_n} > 1 - h_3\varepsilon$$

for some constant h_3 (this is where differentiability of w is invoked), so that tomorrow's payment exceeds the cost $c(p_n, p_{n+1})$ today for ε small enough. Pick δ close enough to one and ε small enough for this to remain the case when discounting is taken into account. It follows that the agent has no incentive to deviate.

Part 2: Let us now consider intrinsic preferences $v(\omega, a)$. Again, we modify the construction of lemma 3 by specifying that the last round is reached in the round immediately before the first time the payment falls short of $\delta \max_{\omega,a} v(\omega, a)$. Taking $v(\omega, a)$ to zero, this occurs at T , for δ sufficiently close to one. As described in the text, equilibrium strategies call for the entire information to be disclosed as soon as this payment is made. If the agent discloses more information than prescribed at earlier rounds, babbling is played: no payment is ever made again. Be-

cause future payments always exceed the benefit from disclosure, it is optimal for the agent (as well as for the firm) to follow this equilibrium strategy.

E. Proof of Theorem 3

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

$$\int_{[0,1]} p' dF_k(p' | p) = p \quad \text{or} \quad \int_{[0,1]} (p' - p) dF_k(p' | p) = 0. \tag{A1}$$

To put it differently, $F_k(\cdot | p)$ is a mean-preserving spread of the firm’s prior belief p .²²

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,

$$\bar{v}_k^F := \int_0^1 [w(p') + X(p')] dF_k(p' | p) - w(p),$$

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 follows from the definition of the event p' , and the second follows from Bayes’s 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

²² The notation $[0, 1]$ for the domain of integration emphasizes the possibility of an atom at zero. 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.

$$\int_0^1 [w(p') + X'] dF_k(p' | p) - w(p) + \int_0^1 V_{1,k-1}(p', X') \frac{p'}{p} dF_k(p' | p), \quad (A2)$$

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 [A1]) and the feasibility of X' .

1. The Optimal Transfers

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

LEMMA 8. Fix the prior belief p and the number of remaining rounds k . The best equilibrium payoff of the type 1 agent, as defined by (A2), 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

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

The type 1 agent's continuation payoff is then given as

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

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

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

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

$$p[V_{1,k}(p) - X^*] + X^* + w(p) = p$$

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 (A2), the contribution to the type 1 agent's payoff from this posterior is equal to

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

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. QED

The intuition behind this lemma is that to maximize V_1 , because the type 1 agent assigns a smaller probability to the posterior decreasing than the type 0 agent, it is best to promise as high a rent as possible to the firm if the posterior belief is lower than the prior belief and as low as possible if it is higher. The function X^* describes this upper bound. As in the example, this bound turns out to be the entire continuation payoff of the type 0 agent in the best equilibrium for the type 1 agent with $k - 1$ periods to go. We can express this bound in terms of the firm's belief and the type 1 agent's continuation payoff, given that the equilib-

rium is efficient. Of course, it is possible to give even higher rents to the firm, provided that the equilibrium that is played in the continuation game gives the type 0 agent a higher payoff than the equilibrium that is best for the type 1 agent. The proof of this lemma establishes that what is gained in the initial period by considering higher rents is more than offset by what must be relinquished in the continuation game, in order to generate a high enough payoff to the type 0 agent.

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

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 (A2), we obtain that

$$V_{1,k}(p) = \sup_{F_k(\cdot | p)} \int_0^1 v_{k-1}(p'; p) dF_k(p' | p) - w(p), \tag{A3}$$

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 (A1); 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 (A3) is $V_1(p) = [p - w(p)]/p$. Recall that this value is the upper bound on $V_{1,k}(p)$ that we derived earlier, so it is the highest payoff that we could have hoped for. We may now finally prove the theorem.

Proof of theorem 3. Recall that the function to be maximized is

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

or rearranging,

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

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

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

or rearranging,

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

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 one-half to $p - \varepsilon$ and one-half 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

$$x(p) + (1-p)w(p) \leq \frac{1}{2} \cdot \frac{1-p}{1-p+\varepsilon} [x(p-\varepsilon) + (1-p+\varepsilon)w(p-\varepsilon)] \\ + \frac{1}{2} [x(p+\varepsilon) + (1-p-\varepsilon)w(p+\varepsilon)] + \varepsilon w(p+\varepsilon) \\ = \frac{1}{2} [x(p-\varepsilon) + (1-p+\varepsilon)w(p-\varepsilon)] \\ + \frac{1}{2} [x(p+\varepsilon) + (1-p-\varepsilon)w(p+\varepsilon)] \\ + \varepsilon \left[w(p+\varepsilon) - w(p-\varepsilon) - \frac{x(p-\varepsilon)}{1-p+\varepsilon} \right].$$

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 almost everywhere on I . Rearranging 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 almost everywhere differentiability of z gives $\int_I x(p)/(1-p) \leq 0$. 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 zero. Given the definition of x , the claim regarding the limiting value of $V_{1,k}$ follows. QED

F. Proof of Lemma 5 and Theorem 2

We adapt the arguments from the proof of theorem 3. Recall that w is assumed to be weakly star shaped (in particular, nondecreasing). Consider a mixed-strategy

equilibrium. In terms of beliefs, such an equilibrium can be summarized by a distribution $F_{k+1}(\cdot | p)$ that is used by the agent (on the equilibrium path) with $k + 1$ rounds left, given belief p , and the continuation payoffs $W_k(\cdot)$ and $V_k(\cdot)$. As before, we may assume that the equilibrium is efficient, and so we can assume that, given that the firm obtains a net payoff of X_k (i.e., given that $W_k = w(p) + X_k$), the type 1 agent receives $V_{1,k}(p, X_k)$, the highest payoff to this type given that the firm receives at least a net payoff of X_k . Since $V_{1,k}$ maximizes the sum of the firm's and type 1 agent's payoff, it holds that, for all k, p and $X \geq 0$,

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

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

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

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

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

that is, such that the belief of the firm follows a martingale. To emphasize the importance of the posterior $p' = 0$, we alternatively write this constraint as

$$\int_0^1 (p' - p) dF_{k+1}(p' | p) = pF_{k+1}(0 | p),$$

where

$$\int_0^1 dF_{k+1}(p' | p) := 1 - F_{k+1}(0 | p).$$

If the type 1 agent randomizes, she must be indifferent between all elements in the support of its mixed action, that is, for all $p' > 0$ in the support of $F_{k+1}(\cdot | p)$, $V_{1,k}(p', X') = \underline{V}_k$, for some \underline{V}_k independent of p' . Assume (as will be verified) that in all relevant arguments, p' and $X \geq 0$ are such that it holds that

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

Recall that this is always possible if X is small enough (cf. lemma 8). 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 zero 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 Sec. IV.A), so that h^* is no steeper than w . That is, h^* satisfies all four constraints from Section IV.B.1.

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

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

To show that h^m cannot be improved on, it suffices to consider arbitrary two-point distributions splitting p into $p_1 < p < p_2$.²³ 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 $h_1(p_1) \leq 1 - (1 - p_1)[w(p_2)/p_2]$, or, rearranging and using the formula for $V_1(p_2)$,

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

²³ Note 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.

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

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

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

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

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

References

- Admati, A. R., and M. Perry. 1991. "Joint Projects without Commitment." *Rev. Econ. Studies* 58:259–76.
- Admati, A. R., and P. Pfleiderer. 1988. "Selling and Trading on Information in Financial Markets." *A.E.R. Papers and Proc.* 78:96–103.
- . 1990. "Direct and Indirect Sale of Information." *Econometrica* 58:901–28.
- Anton, J., and D. Yao. 1994. "Expropriation and Inventions." *A.E.R.* 84:190–209.
- . 2002. "The Sale of Ideas: Strategic Disclosure, Property Rights, and Contracting." *Rev. Econ. Studies* 69:513–31.
- Aumann, R. J., and S. Hart. 2003. "Long Cheap Talk." *Econometrica* 71:1619–60.
- Aumann, R. J., and M. B. Maschler. 1995. *Repeated Games with Incomplete Information*. Cambridge, MA: MIT Press.
- Bruckner, A. M., and E. Ostrow. 1962. "Some Function Classes Related to the Class of Convex Functions." *Pacific J. Math.* 12:1203–15.
- Che, Y.-K., and J. Sákovičs. 2004. "A Dynamic Theory of Holdup." *Econometrica* 72:1063–1103.
- Compte, O., and P. Jehiel. 2004. "Gradualism in Bargaining and Contribution Games." *Rev. Econ. Studies* 71:975–1000.
- Ely, J., A. Frankel, and E. Kamenica. 2015. "Suspense and Surprise." *J.P.E.* 123: 215–60.
- Eső, P., and B. Szentes. 2007. "The Price of Advice." *Rand J. Econ.* 38:863–1179.
- Forges, F. 1990. "Equilibria with Communication in a Job Market Example." *Q.J.E.* 105:375–98.
- Forges, F., and F. Koessler. 2008. "Long Persuasion Games." *J. Econ. Theory* 143: 1–35.
- Fudenberg, D., and J. Tirole. 1991. *Game Theory*, Cambridge, MA: MIT Press.
- Goldwasser, S., S. Micali, and C. Rackoff. 1985. "The Knowledge Complexity of Interactive Proof-Systems." In *Proceedings of the 17th ACM Symposium on the Theory of Computing*, 291–304. Providence, RI: Assoc. Computing Machinery.
- Gul, F. 2001. "Unobservable Investment and the Hold-Up Problem." *Econometrica* 69:343–76.
- Hörner, J., and A. Skrzypacz. 2012. "Selling Information." Discussion Paper no. 1743RR, Cowles Found., Yale Univ.
- Kamenica, E., and M. Gentzkow. 2011. "Bayesian Persuasion." *A.E.R.* 101:2590–2615.
- Landsberger, M., and I. Meilijson. 1990. "Lotteries, Insurance and Star-Shaped Utility Functions." *J. Econ. Theory* 52:1–17.
- Marx, L. M., and S. A. Matthews. 2000. "Dynamic Voluntary Contribution to a Public Project." *Rev. Econ. Studies* 67:327–58.

- Matthews, S. A., and A. Postlewaite. 1995. "On Modeling Cheap Talk in Bayesian Games." In *The Economics of Informational Decentralization: Complexity, Efficiency, and Stability; Essays in Honor of Stanley Reiter*, edited by J. O. Ledyard, 347–66. Boston: Kluwer.
- Rosenberg, D., E. Solan, and N. Vieille. 2013. "Strategic Information Exchange." *Games and Econ. Behavior* 82:444–67.
- Spence, M. 1973. "Job Market Signaling." *Q.J.E.* 87:355–74.