SIGNALING COVERTLY ACQUIRED INFORMATION [1]

MEHMET EKMEKCI AND NENAD KOS

Date: November 4, 2019

We study how information acquisition is affected by signaling. A sender decides whether to learn his type at a cost prior to taking the signaling action. A receiver responds after observing the signaling action but not whether the sender acquired information. In the benchmark model with observable information acquisition decision the sender never acquires information. However, if information acquisition is covert, the sender does acquire information in the unique outcome surviving a form of never weak best response refinement when the information is cheap. While more than one outcome can survive the refinement, sufficient conditions on the sender's payoff are provided that guarantee uniqueness.

KEYWORDS: Signaling, information acquisition, refinements

JEL Classification Numbers: D82, G34.

Boston College, Economics Department. Email: ekmekci@bc.edu.

Department of Economics and IGIER, Bocconi University, Via Roentgen 1, 20136 Milan, Italy Email: nenad.kos@unibocconi.it.

[1] We would like to thank Dilip Abreu, Alp Atakan, Rahul Deb, George Mailath, Andrea Prat, Larry Samuelson and seminar audiences at Boston College, Boston University, Rice University, NYU, Royal Holloway, Aalto University, Norwegian Business School - Oslo, EUI, LMU, CSEF-IGIER Symposium 2018, ESSET Gerzensee 2019, Queen Mary Economic Theory Workshop.

Akerlof (1970) demonstrated how asymmetry of information can wreak havoc in markets. The buyers' inability to differentiate between various qualities of products can unravel to the point where nothing but the worst items are traded. Several solutions have been proposed to remedy the problem, a particularly elegant one in the form of signaling due to Spence (1973). Signaling provides an opportunity for the agents to credibly reveal their information. However, the inherent freedom to choose out of equilibrium beliefs results in a wealth of equilibria, and low predictive power of the model. In response, much work and attention has been directed towards taming of the out of equilibrium beliefs and restricting the set of equilibria through refinements. Most commonly used refinements, as intuitive criterion, D1 , D2 (Cho and Kreps (1987)), divinity (Banks and Sobel (1987)) etc., have originated from strategic stability (henceforth stability-based refinements), a concept introduced by Kohlberg and Mertens (1986). A common implication of these refinements is the selection of the most efficient separating equilibrium. Each type undertakes a distinct observable action thereby fully revealing the sender's private information. This leads us the main question of this paper: if the sender's information is revealed through his action, what good is that information to him and why should he acquire information to start with?

The model we study is the following. A sender can covertly learn his type at a cost c. After learning it, if he chose to, he takes a signaling action, which is less costly for a higher type of the sender. Majority of the paper focuses on the case of two types; we refer to the more cost-efficient type as the high type. The receiver observes the sender's action, but not whether he acquired information and replies with an action of his own. The sender would like the receiver to take as high an action as possible and is risk neutral with respect to the receiver's action. The receiver, on the other hand, wants to take an action as close as possible to the state.

Before exploring the environment with covert information acquisition, we examine the benchmark model where the information acquisition decision is observable. Were one to try to sustain an equilibrium where information is acquired, the receiver would on the equilibrium path on average take the action equal to the expected value of the type. On the other hand, if the sender deviated to not acquiring information he would have nothing to signal and the receiver would take the action equal to the average type (assuming that the sender cannot signal what he does not know). Since in both cases the sender gets the same benefit from the receiver yet does not have to signal nor pay for information when no information is acquired, he has a profitable deviation. Thus, in the unique equilibrium outcome under costly information, the sender does not acquire information nor signals. This result reinforces the question, where does the information come from in signaling models. Needless to say, in some environments it is easy to argue that the sender acquired information for reasons extraneous to the model. However, for a tool as commonly used as signaling, one ought to understand whether every and each time it is used one has to find an explanation for information through motives completely outside of the model.

The set of equilibria in the game with unobservable information acquisition is much larger. There are equilibria without information acquisition and no signaling, equilibria with information acquisition, and even equilibria without information acquisition but strictly positive amount of signaling. When information is free, these equilibria coexists with information acquisition followed by the well-known equilibria arising in Spence model: pooling, separating, and semi-separating. However, the set of equilibria shrinks significantly as soon as information becomes costly. Namely, in any equilibrium where information is acquired with positive probability the two types must separate themselves strongly, that is, each type must strictly prefer their own equilibrium action(s). Were that not the case, the sender could deviate to not acquiring information and choosing the action that gives both types the equilibrium payoff, but forgoing the cost of information. Some immediate implications are that under costly information: (i) in equilibrium there can be no pooling after information acquisition; (ii) when information is acquired with probability one the low type does not signal ($s_L = 0$), and (iii) the Riley outcome—the most efficient separating equilibrium in the Spence model—cannot be supported as an equilibrium after information acquisition when information is costly. In the most efficient equilibrium in which the sender acquires information the sender is indifferent between his equilibrium play and deviating to not acquiring information followed by pretending to be the high type. Interestingly, in such an equilibrium the high type must burn more surplus than in the Riley outcome, and increasingly so in the cost of information. Namely, if the sender deviates to not acquiring information he suffers due to signaling regardless of the state of the world, whereas he suffers only as the high type after acquiring information. The demonstrated multiplicity of equilibrium outcomes, and the associated low predictive power of the model, all but necessitate a refinement.

Stability based refinements are defined and operate on signaling games—games where a privately informed sender takes an action that is followed by the receiver's action—which the game analyzed in this paper is not. What is more, these refinements do not readily extend to the game analyzed here.[1] In light of the above, the natural candidate for the refinement is the strategic stability itself. However, two obstacles stand in the way of applying strategic stability directly. First, it is only defined for finite games, while the extension to more general environments has not been established. Second, checking for all the equilibria of every perturbation of the game (in the sense of stability) is rather laborious, to say the least. As a middle ground, we introduce a type of never weak best response (NWBR) refinement, where we check whether a certain equilibrium outcome survives the refinement by constructing the set of all equilibria with the given outcome and iteratively deleting strategies that are never weak best response to any strategy in the set of equilibria of the game obtained after deletion. Precise definition of the procedure is provided in the main text.

We start refining equilibria in a stylized model where the sender's cost of signaling takes the quadratic form. We show that there exists a threshold cost of information $c^* > 0$ such that for any $c < c^*$ only the most efficient separating outcome with information acquisition survives the refinement, while for $c > c^*$ only the outcome in which the sender does not acquire information and undertakes the least costly action survives. At the threshold c^* a vast array of equilibria survives the refinement. Most importantly, that information is acquired when cheap stands in stark contrast with the result under observable information acquisition where information is never acquired.

We further explore the more general single-crossing environment where the sender's cost of signaling is given by a function $g(s, \theta)$, which is increasing in the signal, decreasing in the type and has a negative cross derivative. We identify additional conditions that guarantee generic uniqueness of equilibrium outcomes surviving the refinement, a log-supermodularity condition on the marginal cost of signaling $g_s(s, \theta)$, and show that single-crossing is enough to guarantee uniqueness when information is cheap. In particular, at low cost of information only the most efficient outcome with information acquisition survives. As the cost of information grows, the high type must burn more and more surplus. Interestingly, under some single-crossing cost of signaling functions g for intermediate region of cost of information the only equilibrium outcome surviving the refinement is the one where the sender does not acquire information, yet undertakes a costly action. The sender signals that she is uninformed rather

[1] The above mentioned refinements are based on comparisons of sets of beliefs for which each type could profitably deviate to a given action. Such type-by-type comparisons do not suffice here. Consider, an equilibrium in which the sender does not acquire information and undertakes the least costly action. The high type does not exist in such an equilibrium. To bring him into existence, so to say, the sender would first have to deviate to acquiring information. Yet to assess whether he indeed wants to do so, one would need to know what both the low and the high type will do after the deviation.

than the informed low type.

Related Literature Our paper builds on the seminal work of Spence (1973) and the subsequent literature on refinements, see Kohlberg and Mertens (1986), Cho and Kreps (1987), Banks and Sobel (1987), Cho and Sobel (1990); most of which was discussed above. Majority of work on signaling is focused on signaling games—games in which a privately informed sender takes an action which is observed by the receiver who in response takes an action of his own—which our game is not. A comprehensive review of the literature on signaling games is beyond the scope of this paper; the interested reader is recommended to consult Riley (2001) and Sobel (2009).

Grassi and Ma (2016) study referrals where two experts compete for clients. Each expert is randomly matched to a client and may acquire information about whether the client is a good fit or not, and may refer the client to the other expert for a fee. The model exhibits a plethora of Perfect Bayesian equilibria among which the authors focus on the ones with "passive beliefs", requiring that out of equilibrium beliefs on the sender's information acquisition decision and private signal do not depend on the pricing decision.

The question of information acquisition on the side of receiver, which likewise leads outside of the scope of signaling games, has received much more attention. Under various degree of generality (and in different applications) it has been studied in Banks (1992), Bester and Ritzberger (2001), Stahl and Strausz (2017), Bester et al. (2019). A common question is how the sender's incentive to signal trades off with the receiver's incentives to acquire information.[2]

In and Wright (2017) study games where the sender (who has no private information to start with) takes two actions in a row and only the second one is observed by the receiver. The second action serves a signal of the first. The authors introduce a refinement that is easily applied and delivers unique equilibrium in many environments. Their results rely on the property that there are no nature's moves between the two sender's actions, thus they cannot be applied to the environment studied in this paper.

[2]Similar ideas arise in the work where the receiver observes a signal a signal of given precision about the sender for free; see Feltovich et al. (2002), Alós-Ferrer and Prat (2012), Daley and Green (2014).

Setting

We study a game between two players: a sender and a receiver. The sender first decides whether to acquire information, then takes an unproductive but costly action. The receiver observes the sender's action, but not whether he acquired information. The environment is a canonical signaling game preceded by an information acquisition stage in which the sender can procure information.

More precisely, there are two states of the world $\Theta = \{\theta_L, \theta_H\}$, a low and a high state, respectively. The prior probability of state θ_H is $\lambda \in (0, 1)$. In the first stage the sender decides whether to learn the state θ at some cost $c \geq 0$ or not, in the second stage he chooses an action $s \in S = [0, \infty)$.[3] In the final stage, the receiver chooses an action $r \in R = [0, \infty)$ after having observed s but not whether the sender acquired information. We refer the set $T := \theta_u \cup \Theta$ as the set of types of the sender, where t denotes a generic element of the set. Type θ_u is the sender's type when he does not acquire information.

The sender's payoff in state $\theta \in \Theta$ is $r - g(s, \theta)$, where g satisfies the following properties: g is C^2, $g_s > 0$ and $g_{s\theta} < 0$ for every $s > 0$. Moreover, $g(0, \theta) = 0$ and $\lim_{s\to\infty} g(s, \theta) = \infty$ for each $\theta \in \Theta$. The sender wants the receiver to take as high an action as possible. Signaling is, however, costly, and the marginal cost of signaling is decreasing in the state. The receiver, on the other hand, maximizes $-(\theta - r)^2$. It is easy to see that the receiver takes an action that is equal to the expected value of θ given his posterior.

Note that $g(s, \theta) = \int_0^s g_s(x, \theta)dx$. Because $g_{s\theta} < 0$ for every $s > 0$, $g_s(x, \theta) > g_s(x, \theta')$ for every $\theta' > \theta$, $x > 0$. Hence, $g(s, \theta') < g(s, \theta)$ if $\theta' > \theta$ and $s > 0$.

Benchmark

A natural benchmark for our model is the environment in which the receiver can observe whether the sender acquired information. The solution concept we apply is Perfect Bayesian equilibrium (PBE) as defined in Fudenberg and Tirole (1991), with the interpretation of the no signaling of what you do not know to imply that the receiver's posterior is equal to the prior when the sender does not acquire information, on or off the equilibrium path.

[3] We explore the case in which the sender receives an imperfect signal about the state of the world by paying a cost in an extension later.

Towards characterizing the set of PBE, consider an outcome where the sender does not acquire information and takes a strictly costly action $s^* > 0$. After observing the sender not acquiring information and choosing s^*, the receiver would optimally respond with $r = E[\theta]$. However, if the receiver encountered the sender who did not acquire information and chose $s < s^*$, he would also hold the belief equal to the prior and respond with $E[\theta]$. The sender would, therefore, find it profitable to deviate towards not acquiring information and $s = 0$. That is to say, no outcome without information acquisition and $s^* > 0$ can be supported as a PBE.

The outcome where the sender does not acquire information and chooses $s^* = 0$, on the other hand, can be sustained as a PBE in several ways. For example with passive beliefs: no matter what off equilibrium behavior the receiver encounters his posterior remains equal to the prior.

As for equilibria with information acquisition. One can conjecture information acquisition to be followed by pooling, separation or even semi-separation of the two types. All these continuations have a feature in common, namely, that the beliefs on the equilibrium path are a martingale and, therefore, the receiver's expected action is $E[\theta]$. However, by deviating to not acquiring information and choosing $s = 0$, the sender could likewise induce the receiver to play $E[0]$, yet forgo the cost of information and possible costs of signaling. In consequence, as long as $c > 0$ the only equilibrium outcome that can be supported in a PBE is the one where the sender does not acquire information and consequently refrains from signaling. A second equilibrium outcome arises when $c = 0$: the one where the sender acquires information and both types pool on $s^* = 0$.[4]

Proposition 1 *For $c > 0$ the only equilibrium outcome that can be supported in a PBE is the one where no information is acquired, the sender chooses $s = 0$ and the receiver $r = E[\theta]$. For $c = 0$, a second equilibrium outcome can be supported where information is acquired, but both types pool on $s^* = 0$.*

The above result sets a somewhat daunting tone, if information acquisition is observable and the information costly, it will not be acquired; even if costless, it will not be conveyed. Sender's private information can, thus, not be justified through observable information acquisition preceding a signaling game. The rest of the paper focuses on the environment where

[4]These equilibria can be refined away with the refinements we discuss later in the paper. A reader should notice that the continuation game after information acquisition corresponds to the standard signaling model where the refinements like the intuitive criterion and D1 yields the separating outcome as the only reasonable.

the receiver does not observe whether the sender acquired information.

Equilibria

The opportunity to acquire information opens venues for behavior not present in signaling games. The sender has an opportunity to save on cost of information and merely pretend that he is informed—to bluff, so to say. The question thus becomes when one can sustain information acquisition as an equilibrium. The solution concept we adopt is Nash equilibrium in which the receiver's strategy is not weakly dominated; for short *an equilibrium*.[5]

We start by examining equilibria in which the sender refrains from acquiring information.

Lemma 1 *Let $\bar{s}_u$ be such that the uninformed sender is indifferent between $(0, \theta_L)$ and $(\bar{s}_u, E[\theta])$. The following statements are true:*

- *For every $s^* \leq \bar{s}_u$, there exists a $c_{s^*} \geq 0$ such that not acquiring information followed by s^* can be supported as an equilibrium if and only if $c \geq c_{s^*}$. Moreover, there exists a $s_u^L \in (0, \bar{s}_u)$ such that $c_{s^*} = 0$ for all $s^* \in [0, s_u^L]$.*
- *Actions $s^* > \bar{s}_u$ cannot be supported in equilibrium where the sender does not acquire information.*

Proof: The action $\bar{s}_u$ as defined in the supposition of the result exists due to the assumption that $g(s, \theta)$ is continuous in s and goes to infinity with s. Now, not acquiring information followed by $s^* > \bar{s}_u$ cannot be supported as an equilibrium. The worst response the sender can expect after not acquiring information and $s = 0$ is $r = \theta_L$, which results in a strictly higher payoff.

As for $s^* \leq \bar{s}_u$, the easiest way to support not acquiring information followed by such an s^* is having the receiver respond to any non equilibrium signaling action s by $r = \theta_L$. The sender, then, does not have an incentive to deviate after having not acquired information. The question remains whether he can find it profitable to deviate to acquiring information.

Let s_u^L be such that the low type is indifferent between $(0, \theta_L)$ and $(s_u^L, E[\theta])$. Due to the

[5]In particular, the receiver's actions after seeing an off path sender's action s is a best reply to some set of beliefs over Θ. Differently, it precludes the receiver from taking actions above θ_H or below θ_L. This condition is the usual admissibility condition imposed on the receiver.

single-crossing assumption $s_u^L < \bar{s}_u$. If the sender were to deviate from not acquiring information and $s^* \leq \bar{s}_u$ (with the reply $r = \theta_L$ for out of equilibrium signaling actions s'), to acquiring information, the high type would choose s^* due to the single-crossing assumption. The low type, on the other hand, would choose 0 if $s^* \leq s_u^L$ and s^* if $s^* \in (s_u^L, \bar{s}_u]$. Consequently, for any $s^* \leq s_u^L$ the sender would after deviating to information acquisition choose the same signaling action he is to choose when he does not acquire information regardless of his type, but incur a cost. He, therefore, has no incentive to deviate for any $c \geq 0$.

If one were to support no information acquisition followed by $s^* \in (s_u^L, \bar{s}_u]$ as an equilibrium outcome, after a deviation to information acquisition the low type sender would strictly prefer $s = 0$—see figure 1b—thereby creating value for information acquisition. The threshold c_{s^*} is defined by the indifference condition:

$$E[\theta] - Eg(s^*, \theta) = \lambda \left(E[\theta] - g(s^*, \theta_H) \right) + (1 - \lambda)\theta_L - c_{s^*},$$

where the left hand-side is the equilibrium payoff from not acquiring information and choosing s^* and the right hand-side from acquiring information, low type choosing $s = 0$ and the high type s^*. □

The most straightforward way to support equilibria is to have the receiver respond to non-equilibrium actions s with $r = \theta_L$. The range of equilibrium signaling actions, s^*, can be split in three regions. For low s^*, equilibria where the sender does not acquire information followed by s^* can be supported for every level of cost c; see figure 1a. In this case, s^* is optimal for both types even after the deviation to acquiring information. The cost of information acquisition would, therefore, be wasted. For the intermediate values of s^*, after the deviation to acquiring information the high type sender would choose s^*, but the low type would strictly prefer $s = 0$; see 1b. The latter necessitate that the cost of acquiring information is large enough in order to prevent the sender from deviating. Finally, large s can never be supported as an equilibrium with no information acquisition. The sender would rather deviate to $s = 0$ without a need to acquire information.

Equilibria with information acquisition introduce several new challenges. When the cost of information acquisition is nil one can sustain equilibria with information acquisition followed by a large set of outcomes familiar from signaling games: pooling, separating and semi-separating. However, when the cost of acquiring information rises above 0 the threat of deviations to not acquiring information alters the set of equilibria.

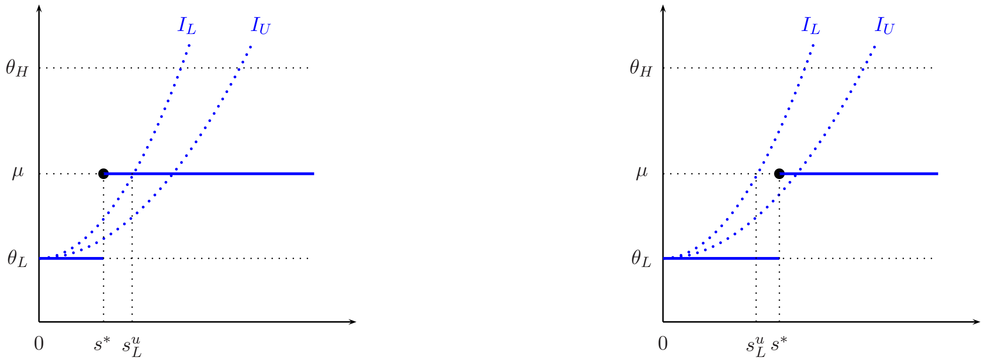

(a) After deviating to information acquisition both types choose s^*.

(b) After a deviation the low type would choose $s = 0$.

Figure 1: Equilibria with signaling but no information.

Lemma 2 *For $c > 0$, in any equilibrium where the sender acquires information with positive probability, each type $\theta \in \Theta$ of the sender must strictly prefer their equilibrium action to an equilibrium action of the other type.*

If after acquiring information one of the two types—say θ_L—was indifferent between his own equilibrium action and some equilibrium action of the high type, the sender could deviate to not acquiring information and take the high type's action, thereby replicating the utility of the high type. Due to the low type's indifference, the sender would also replicate the payoff in the low state. That is to say, the sender could without acquiring information state by state achieve the same utility as if he acquired it but forgo the cost of information.

The above lemma has the following implications. For $c > 0$:

- there are no equilibria where the sender acquires information with positive probability and the low and the high type pool with positive probability;
- information acquisition followed by the Riley outcome (the outcome where the low type chooses s=0, and the high type chooses $s = s_H$ such that the low type is indifferent) cannot be supported as an equilibrium.

The first point states that the two types separate after information is acquired. Needless to say, when the low type can be identified, he has no incentive to signal his type, that is, $s_L = 0$. The high type's equilibrum behavior, on the other hand, must be curtailed in order to prevent the sender from deviating to not acquiring information and pretending to be one of the two types. The relevant constraints are:

$$\text{(1)} \qquad \lambda\left(\theta_H - g(s_H, \theta_H)\right) + (1-\lambda)(\theta_L - g(0, \theta_L)) - c \geq \theta_L - E[g(0, \theta)];$$

which and

$$\text{(2)} \qquad \lambda\left(\theta_H - g(s_H, \theta_H)\right) + (1-\lambda)(\theta_L - g(0, \theta_L)) - c \geq \theta_H - E[g(s_H, \theta)],$$

where the first constraint requires that the sender not find it profitable to deviate to not acquiring information and pretending to be the low type and the second that he cannot profitably deviate to not acquiring information and mimicking the high type. The two constraints reduce to

$$\text{(3)} \qquad \theta_H - g(s_H, \theta_H) - \frac{c}{\lambda} \geq \theta_L - g(0, \theta_H);$$

and

$$\text{(4)} \qquad \theta_L - g(0, \theta_L) - \frac{c}{1-\lambda} \geq \theta_H - g(s_H, \theta_L),$$

respectively. If the sender were to deviate to not acquiring information and pretend to be, say, the high type, his behavior would differ from the prescribed only in the low state. In addition, he would save on the cost of information acquisition. Not to benefit from the said deviation, the low type must prefer his equilibrium action to the high type's sufficiently enough to outweigh the savings on the cost of information. The deviation constraints preventing the sender from deviating to not acquiring information are, therefore, stronger than the deviation constraints preventing the agent from simply misrepresenting the type after having acquired information. In other words, in any equilibrium where the sender acquires information the relevant deviation constraints are the ones toward not acquiring information.

For low costs of information, a wide array of action after information acquisition can be supported in equilibrium for the high type. Of particular interest is the most efficient such equilibrium with information acquisition—the one where the high type burns the least surplus. The lower bound on the high type's signaling action is imposed by the constraint preventing the sender from deviating to not acquiring information and choosing the high type's action; inequality (4). In the most efficient equilibrium with information acquisition,

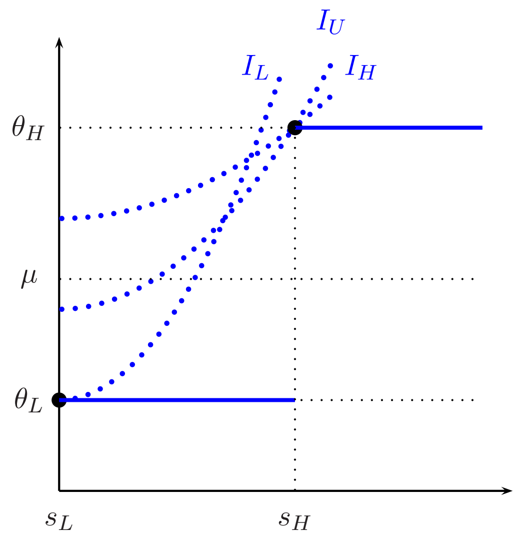

Figure 2: An equilibrium with information acquisition for $c > 0$. The high type burns enough surplus for the low type to strictly prefer his own option.

where the high type's action is denoted s_H^* the said incentive constraint is binding:

$$\theta_L - g(0, \theta_L) - \frac{c}{1-\lambda} = \theta_H - g(s_H^*, \theta_L). \quad (5)$$

The equality uniquely pins down s_H^* for every c; and vice-versa. A cursory glance reveals that s_H^* is increasing in c. At $c = 0$ the Riley outcome obtains, but as c increases, so does s_H^*. The larger the cost of information the more value needs to be created in order to acquire it. True, forcing the high type to burn more surplus reduces the payoff from being informed. However, it reduces it only when the sender is the high type. At, the same time if the sender deviates to being uninformed, he prefers the high type's option. Reducing the high type's payoff, therefore, reduces the deviation payoff more than the equilibrium payoff. Differently, reducing the high type's payoff creates value for learning that the sender is the low type.

There is an upper bound on the cost of information, $\bar{c}$, for which equilibria with information acquisition can be sustained. The incentive constraints (1) and (2) can be rewritten as the upper and the lower bound on the high type's action s_H, respectively. The upper bound is decreasing and the lower increasing in the cost of information. The upper bound on the cost of the information under which information acquisition with certainty can be sustained as equilibrium is, therefore, reached when the two constraints bind simultaneously:

$$\begin{aligned}\theta_L - E[g(0,\theta)] &= \lambda\,(\theta_H - g(\bar{s}_H^*,\theta_H)) + (1-\lambda)(\theta_L - g(0,\theta_L)) - \bar{c} \\ &= \theta_H - E[g(\bar{s}_H^*,\theta)],\end{aligned}$$

where the only value of s_H that can be sustained under $\bar{c}$ is denoted by $\bar{s}_H^*$. Notice that $\bar{s}_H^*$ is such that the uninformed sender is indifferent between $(0,\theta_L)$ and $(\bar{s}_H^*,\theta_H)$; see Figure 3.

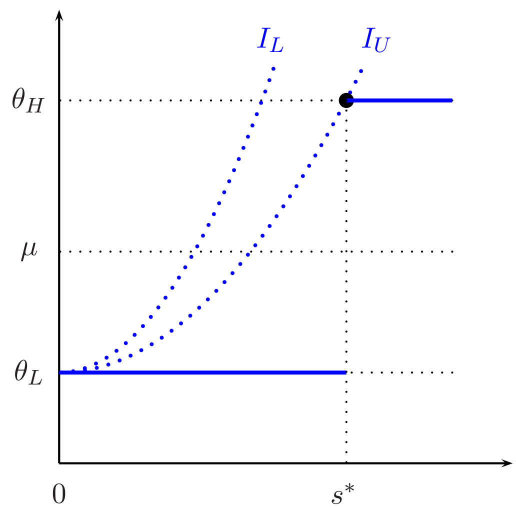

Figure 3: An equilibrium with information acquisition and the highest s_H.

The above result is summarized in the next lemma.

Lemma 3 *Equilibria with information acquisition exist for $c \leq \bar{c}$. In any such equilibrium $s_L = 0$. In the most efficient equilibrium with information acquisition the high type's action, s_H^*, is given by equality (5). Moreover, s_H^* is increasing in c.*

Finally, there is a myriad of equilibria in which the sender randomizes over information acquisition decisions. We omit the characterization of those.

Refinements

The profusion of equilibria gives rise to a variety of behavior: the sender can acquire information, not acquire it, even undertake a strictly costly actions after remaining uninformed.

This leaves the modeler with rather little predictive power. In an attempt to narrow down the players' behavior a wide array of refinements has been developed. Perhaps the most commonly used refinement in signaling games—the intuitive criterion of Cho and Kreps (1987)—reduces the set of equilibria in the model with two types to a single outcome. With more than two types a stronger refinements, for example D1 (Cho and Sobel (1990)), are required for uniqueness. The two mentioned refinements are defined on signaling games, games where a privately informed sender takes an action that is followed by the receiver's action. The game analyzed in this paper is not a signaling game and therefore the before-mentioned refinements do not apply directly. In what follows we argue that such refinements cannot be easily extended to accommodate our environment.[6]

Consider an equilibrium in which the sender does not acquire information and undertakes the least costly action. The above mentioned refinements would start by characterizing the set of beliefs for which a type, say the high type, can profitably deviate to each action. However, the high type in the considered equilibrium does not exist, and neither does for that matter the low. To talk about the high type's deviations, the sender would first need to deviate to acquiring information. But to discern whether the sender has an incentive to do so one would need to determine how both the low and the high type would behave if the sender was informed. Looking at each type's deviations in isolation will not do. Differently, the above mentioned refinements rely on the set of types being fixed, whereas it is endogenous in our game. If the sender does not acquire information there is a single type, if he does there are two, and if he randomizes over information acquisition, there are three types.

A refinement that does apply to our setting is strategic stability of Kohlberg and Mertens (1986). Strictly speaking, strategic stability is defined for finite games, thus, the precise statement would be that it applies to a discretized version of our game. The point is, however, somewhat mute as strategic stability is notoriously difficult to verify—it requires checking every sequence of particular trembles. We, instead, resort to a simplified refinement called never weak best response (NWBR); for a good primer see chapter 11 in Fudenberg and Tirole (1991) and Cho and Kreps (1987).[7] The building block of NWBR is a property of strategically stable sets: if one erases from a game some strategies that are never weak best response

[6] The two mentioned refinements belong to a class of stability-based refinements, thus named due to their connection to strategic stability of Kohlberg and Mertens (1986). An alternative approach would be to explore the refinement proposed in Mailath et al. (1993), which too is defined for signaling games and would therefore have to be generalized for our game. This is beyond the scope of our paper.

[7] The development of equilibrium refinements took a rather interesting path. Introduction of strategic stability (Kohlberg and Mertens (1986)) was followed by development of simpler refinements that can be more readily applied to signaling games while guaranteeing that the outcome(s) they deliver is stable.

to any strategy in a stable set, the newly obtained game has a stable set that is contained in the stable set of the original game one started with; Kohlberg and Mertens (1986). In addition, generically there exists a stable set such that the distribution over outcomes is unique. One can, therefore repeatedly apply NWBR to refine away equilibria. Starting with a set of equilibria with a unique outcome one erases (possibly iteratively) never weak best response strategies. If one arrives at a game where the starting outcome cannot be supported as an equilibrium, then what one started with can not be a stable set. NWBR has commonly been used to validate equilibrium refinements, the standard result being that the refinement does not eliminate anything NWBR would not eliminate itself. In fact, Cho and Sobel (1990) show that D1 is equivalent (in terms of outcomes) to NWBR in monotonic signaling games.

The main reason for adopting the NWBR refinement (defined below) is the following. The model with information acquisition when information is free is the natural counterpart of the standard signaling game. To facilitate the comparison of our findings with the existing results in signaling games we wanted a refinement that is general enough to apply to the game studied here while at same time replicating the results obtained under the stability-based refinements in signaling games.

Our Refinement Let Γ be the normal form representation of the sender-receiver game described in the section Setting. Let A_1 be the set of pure strategies for the sender, and A_2 be the set of pure strategies for the receiver. A mixed strategy for Player i is a probability distribution over Player i's pure strategies, i.e., $\sigma_i \in \Delta A_i$.

A strategy profile $\sigma = (\sigma_1, \sigma_2) \in \Sigma := \Delta A_1 \times \Delta A_2$ is a Nash equilibrium if $u_i(\sigma_1, \sigma_2) = \max_{\sigma_i'} u_i(\sigma_i', \sigma_{-i})$ for each $i \in \{1, 2\}$. A Nash equilibrium is admissible for the receiver if σ_2 is a weakly undominated strategy for the receiver. Each strategy profile leads to an outcome o which is a probability distribution over the terminal nodes of the game.

Fix a sender-receiver game in the normal form Γ', where the set of pure strategies for player i is $A_i' \subseteq A_i$, and fix an outcome o of a Nash equilibrium of Γ' that is admissible for the receiver. Let $\Sigma'(o, \Gamma')$ be the largest set of Nash equilibria of Γ' that are admissible for the receiver, and that lead to the outcome o. Observe that $\Sigma'(o)$ can be the empty set.

Definition 1 *Γ'' is a pruning of (Γ', o) where o is an outcome of some Nash equilibrium of Γ' that is is admissible for the receiver if:*

1. $A_1'' \subset A_1'$ and $A_2'' = A_2'$.

2. *If* $a_1' \in A_1'$*, and if* $a_1' \notin A_1''$*, there exists no* $\sigma' \in \Sigma'(o, \Gamma')$ *such that* a_1' *is a weak best reply to* σ_2'*.*

Pruning of a game with respect to an outcome o erases a strategy of the sender (Player 1) only if that strategy is never a weak best response to any of receiver's (Player 2's) strategies in the set of strategy profiles which are Nash equilibria that are admissible for the receiver. Pruning does not erase any strategy of the receiver. We choose this specification because, the only strategies of the receiver that are never weak best responses in a set of equilibria that lead to a unique outcome are those that do not lead to the outcome. Hence, erasing such strategies would not lead to any change in the power of the NWBR test we define below. Despite not pruning any of receiver's strategies, pruning operates with respect to the set of all equilibria in which the receiver's strategies are admissible.

Definition 2 *An outcome* o *fails the NWBR test if either* i) *it is not an outcome of a Nash equilibrium* σ *of* Γ *that is admissible for the receiver or* ii) *there exists a sequence* $\{\Gamma^n\}_{n=1,2..,k}$ *of games which satisfy:*

1. Γ^1 *is a pruning of* (Γ, o)*.*
2. Γ^n *is a pruning of* (Γ^{n-1}, o) *for every* $n = 2, 3, ..., k$*.*
3. o *is not an outcome of any Nash equilibrium that is admissible for the receiver in* Γ^k*.*

Quadratic Costs

It is instructive to first study the quadratic environment where the seller's cost of signaling takes the form

$$g(s, \theta) = \frac{s^2}{\theta},$$

and then move to more general single-crossing environment.

We split the equilibrium outcomes into three groups—with information acquisition, without information acquisition and with randomization over information acquisition decisions—and study when the outcomes in each group survive the refinement. First, we take a closer look at outcomes without information acquisition, which are further divided into the ones followed by signaling ($s^* > 0$) and the ones without signaling ($s^* = 0$).

Definition 3 *For a fixed set of equilibria we say that an action s is never weak best response for type i, $i \in \{\theta_L, \theta_H\}$, if there is no strategy where the sender acquires information and the type i chooses s that is a weak best response to some receiver's strategy in the set.*

Let $\bar{s}^*$ be such that the low type is indifferent between $(0, \theta_L)$ and $(\bar{s}^*, E[\theta])$; see Figure 4.

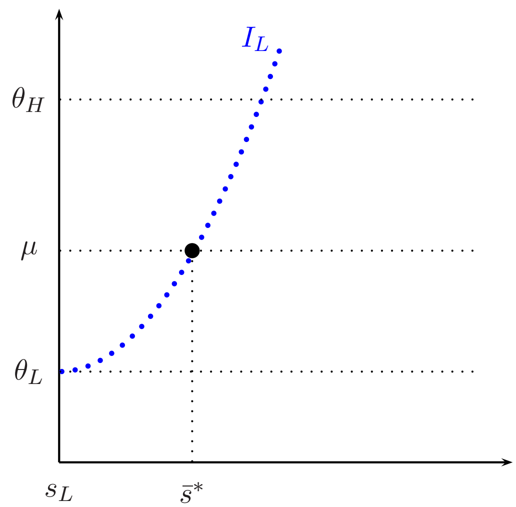

Figure 4: Graphical definition of $\bar{s}^*$.

Lemma 4 *Outcomes with no information acquisition and $s^* > \bar{s}^*$ can be refined away for every cost of information, while each equilibrium outcome with no information acquisition and $s^* \in (0, \bar{s}^*]$ can be refined away for every cost of information bar one; denoted by $b(s^*)$. In particular, at $c = b(s^*)$, the outcome with no information acquisition and $s^* \leq \bar{s}^*$ survives the refinement.*

Proof: See Appendix for the proof. □

Consider an outcome in which the sender does not acquire information and chooses s^*. In any equilibrium with the outcome the receiver responds with $E[\theta]$. Lemma 1 established limits to the amount of signaling s^* that follows no information acquisitions—if too much signaling was required, the uninformed sender is better off being considered to be the low type. The refinement restricts the signaling further. If the sender were to choose an $s^* > \bar{s}^*$ after not acquiring information, the actions between $\bar{s}^*$ and s^* would be never weak best response for the low type. This is due to the receiver's response never being above the uninformed agent's indifference curve through $(s^*, E[\theta])$ and the latter being below the low type's indifference curve through $(0, \theta_L)$ on the interval $(\bar{s}^*, s^*)$. In the game obtained after pruning all the strategies where the low type chooses an action in $(\bar{s}^*, s^*)$, the uninformed sender could profitably deviate to the same interval.

On the other hand, equilibria where the sender does not acquire information and signals moderately, $s^* \leq \bar{s}^*$, cannot be refined away invariably. For each $s^* \leq \bar{s}^*$ such an equilibrium survives the refinement for precisely one value of the cost of information. Fix an equilibrium outcome with no information acquisition and $s^* \leq \bar{s}^*$ and let s_u be such that the uninformed sender is indifferent between $(s^*, E[\theta])$ and (s_u, θ_H):

$$(6) \qquad E[\theta] - E\left[\frac{s^{*2}}{\theta}\right] = \theta_H - E\left[\frac{s_u^2}{\theta}\right];$$

see Figure 5. Actions above s_u are NWBR neither for the low type nor the uninformed sender. One can, therefore, prune away all the strategies where the two play such actions. In the newly obtained game, termed *the refined game*,[8] only the high type can choose actions $s > s_u$, the receiver, therefore, responds to any such action with θ_H; the solid line in the figure. Since the receiver's response in any equilibrium with the prescribed outcome is (weakly) below the uninformed's indifference curve, single-crossing implies that the best option for the high type is s_u (to which he expects the reply θ_H). Furthermore, the sender must be indifferent between his equilibrium play and the deviation to acquiring information followed by s^* as the low type and s_u as the high type. If the deviation yielded a smaller payoff, any strategy with information acquisition would too and the actions slightly below s^* would never be a best response for the low type. Upon erasing the strategies where the low type plays the mentioned signaling actions, one would arrive at a game where the uninformed sender would find it profitable to reduce his signaling action. The indifference condition is:

$$(7) \qquad E[\theta] - E\left[\frac{s^{*2}}{\theta}\right] = \lambda\left(\theta_H - \frac{s_u^2}{\theta_H}\right) + (1-\lambda)\left(E[\theta] - \frac{s^{*2}}{\theta_L}\right) - b(s^*),$$

where the left-hand side is the equilibrium payoff and the right-hand side the payoff from the deviation. Since s_u is pinned down by s^* through equation (6), the indifference can obtain only for one cost of information.

The following result characterizes how the cost at which each of the equilibria without information acquisition survives varies with the amount of signaling s^*.

[8] The refined game plays a prominent role in our analysis, thus warranting the special name. It should also be noted that it is a set of games, one for each s^*, rather than just one.

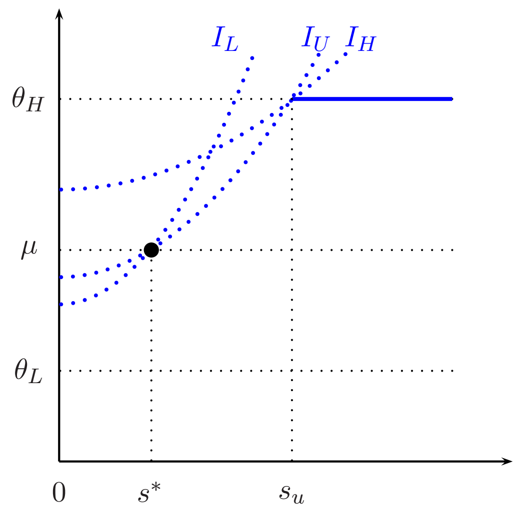

Figure 5: The equilibrium outcome where the sender does not acquire information and chooses s^*. Actions above s_u are NWBR for the low type or the uninformed sender.

LEMMA 5 *There exists a $c_N > 0$ such that*

$$b(s^*) = c_N, \text{ for all } s^* \in [0, \bar{s}^*].$$

Moreover, no equilibrium outcome without information acquisition survives for $c < c_N$ and the only equilibrium outcome without information acquisition that survives the refinement for $c > c_N$ is the one without signaling, i.e. $s^* = 0$.

All the equilibria where the sender does not acquire information and chooses a moderate signaling action s^* survive the refinement at one and the same cost of information c_N.[9] The result relies on a computation that leverages the particular payoff structure studied in this section; more general environments will be examined later in the paper.

Below c_N all the equilibrium outcomes where information is not acquired can be refined away. The only outcome that still requires attention is the outcome without information acquisition or signaling, $s^* = 0$. As above, one can construct s_u and argue that the payoff from the particular deviation towards inquiring information should not be larger than the equilibrium payoff. The indifference as in the case of $s^* > 0$, however, is here not necessary as the sender after not acquiring information cannot deviate to a lower action. Indeed, the

[9]While this makes the survival of equilibrium outcomes with information acquisition and with signaling ($s^* > 0$) non-generic here, these outcomes will play a prominent role in the following section.

equilibrium outcome without information acquisition and $s^* = 0$ survives the refinement for every $c \geq c_N$.

Next we turn attention to equilibria with information acquisition.

LEMMA 6 *There exists a $c_I > 0$ such that no equilibrium outcome with information acquisition survives the refinement for $c > c_I$. For every $c \leq c_I$ the only equilibrium outcome where information is acquired with probability that survives the refinement is the most efficient equilibrium outcome with information acquisition (and separation of types).*

The above result establishes existence of a c_I such that for every $c \leq c_I$ only one equilibrium outcome with information acquisition survives the refinement—the one in which the two types separate themselves, and the sender is indifferent between his equilibrium play and the deviation towards no information acquisition trailed by the high type's equilibrium action—while for $c > c_I$ all the equilibria with information acquisition are refined away. At $c = 0$ the only information acquisition outcome that survives the refinement is the one in which the low type is indifferent between his own and the high type's action (the Riley outcome). As the cost of information increases, the equilibrium outcome with information acquisition that survives the refinement requires the high type to burn more and more surplus. Alternatively, s_H grows with c; see Figure 6. When information becomes costlier the receiver is more likely to question whether the sender truly acquired it. To demonstrate his claim the sender must forgo more and more surplus. When cost becomes too large, the receiver ceases to believe that the sender acquired information all-together. The combined cost of acquiring information and signaling its acquisition would be prohibitive.

The proof of the above result proceeds as follows. First we argue that at $c = 0$ any equilibrium with information acquisition and pooling can be refined away. Moreover, Lemma 2 established that pooling cannot obtain at all in an equilibrium with information acquisition for $c > 0$, even more, that in such equilibria each of the two types strictly prefers their own action. We then show that an equilibrium with information acquisition can be refined away unless the sender is indifferent between the equilibrium play and the deviation towards not acquiring information followed by pretending to be the high type. If the sender strictly preferred his equilibrium play, the actions just below the high type's action would be NWBR for the low type (by Lemma 2) as well as for the uninformed sender. After pruning away those NWBR strategies, the receiver should in any equilibrium respond to an action just below the high type's with $r = \theta_H$, which would provide a profitable deviation for the high

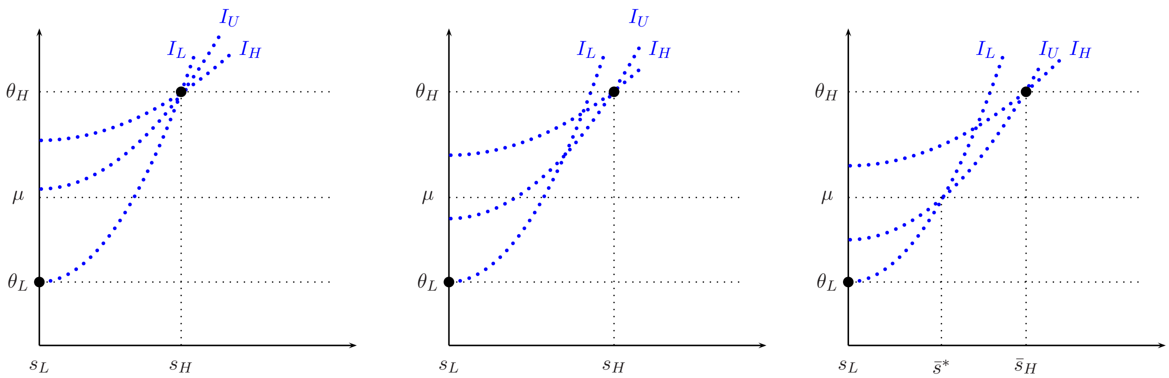

(a) $c = 0$ (b) $c \in (0, c_I)$ (c) $c = c_I$

Figure 6: The most efficient separating equilibria with information acquisition.

type.

As c grows, the high type is forced to undertake more and more signaling, that is, to choose a higher s_H. Equilibria with too high s_H, however, can be refined away. In particular, if one draws the uninformed's indifference curve through the high type's action s_H and the low type's indifference curve through $(0, \theta_L)$, the vertical intercept of the two indifference curves must be at least $E[\theta]$. Should it drop below, the actions just above the intercept are NWBR for the low type; see Figure 7. After pruning the strategies where the low type plays these actions, one obtains a game in which the original outcome cannot be supported as an equilibrium. The final part of the proof painstakingly verifies that the equilibria with information acquisition and indifference cannot be refined away for $c \leq c_I$.

The last class of equilibria to be considered are the equilibria where the sender randomizes over information acquisition decisions.

Lemma 7 *Equilibrium outcomes with randomization over information acquisition decisions survive the refinement at a single cost of information, denoted c_R. All such surviving equilibrium outcomes are separating and with the property that the low type is indifferent between his own and the uninformed sender's equilibrium action, and the uninformed sender between his own and the high type's equilibrium action.*

We have established that only the equilibrium outcome with information acquisition and indifference over information acquisition decisions survives the refinement for each $c < c_I$, that only the equilibrium outcome with no information acquisition and no signaling survives

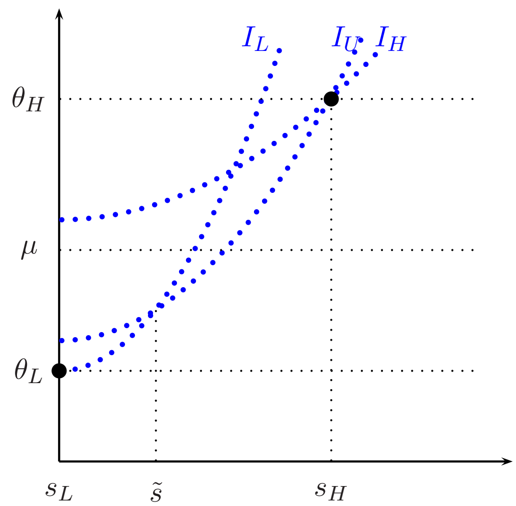

Figure 7: Too much signaling.

the refinement for $c > c_N$ and that the equilibria with randomization over information acquisition decisions survive the equilibrium refinement only for $c = c_R$. The following result relates c_I, c_R and c_N, thereby providing a full characterization of the equilibria that survive the refinement.

Proposition 2 *There exists a $c^* > 0$ such that $c_I = c_N = c_R = c^*$, implying that generically a unique equilibrium outcome survives the refinement. For $c < c^*$, the unique equilibrium outcome that survives the refinement is the most efficient equilibrium outcome with information acquisition (and separation of types), while for $c > c^*$ only the equilibrium with no information acquisition and no signaling survives.*

Proof: The cost of information acquisition under which equilibria without information acquisition can be refined away, c_N, is equal to the constant function b. At the same time, the cost of information c_I is defined so that the low type's indifference curve through $(0, \theta_L)$ and the uninformed sender's indifference curve through (s_H, θ_H) intersect at $(s_{LU}, E[\theta])$, and therefore equal to $b(s_{LU})$. Since b is constant, $c_I = c_N$. In the equilibria with randomization over information acquisition the uninformed sender's indifference curve through s_H intersects with the low type's indifference curve through $(0, \theta_L)$ at a point with the vertical component $E[\theta]$. Thus $c_I = c_N = c_R$. □

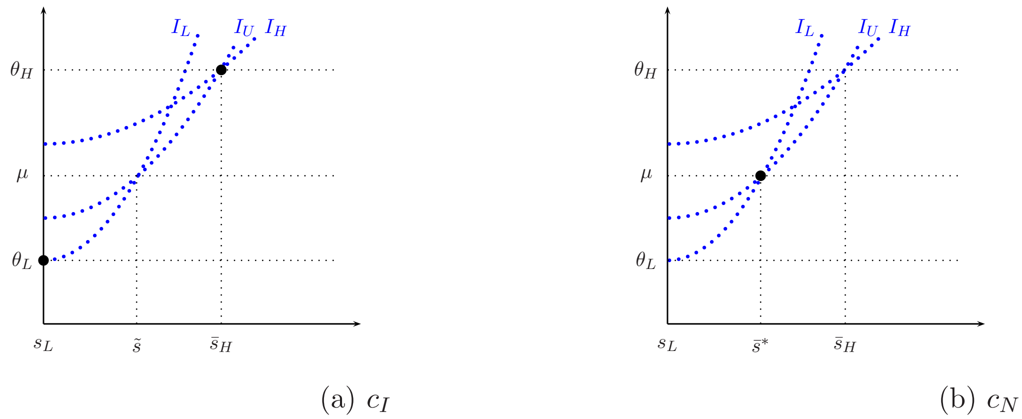

Figure 8: c_I and c_N coincide.

A single equilibrium outcome survives the refinement at every cost of information bar one. When the information is cheap, the sender acquires information, when it is expensive, he does not. This stands in stark contrast to the finding that with observable information acquisition information is never acqured. Covertness of information acquisition (or non-verifiability), thus, provides a rationale for a privately informed sender.

That the lower bound on equilibria without information acquisition, c_N, coincides with the upper bound on equilibria with information acquisition, c_I, is most readily demonstrated graphically; see Figure 8. The cost c_I is determined by the equilibrium in which the low type's indifference curve and the uninformed sender's indifference curve intersect at a point with the vertical component $E[\theta]$. Similarly, the cost of information c_N can be computed through indifference in equilibrium where the sender does not acquire information and chooses the amount of signaling $\bar{s}^*$ such that the low type would be indifferent between $(0, \theta_L)$ and $(\bar{s}^*, E[\theta])$. In the computation of both costs, the high type chooses the same point. The choices of the uninformed sender and the low type are not the same in the two cases, but they do lie on the same indifference curve. The case of c_R is analogous.

General Cost

Thus far we explored a stylized single-crossing environment which streamlined the exposition and simplified some steps in the analysis. In what follows we analyze a more general single-crossing setting as outlined at the beginning of the section Setting.

The first refinement result, Lemma 4, relied on the single-crossing property rather than on the details of the quadratic setting and, therefore, carries over to the more general environment without additional work. Equilibrium outcomes without information acquisition and a moderate amount of signaling, s^*, survive the refinement for precisely one level of cost $b(s^*)$. Mirroring the analysis of Lemma 4, see equation (7), one obtains

$$(8) \qquad b(s^*) = \lambda(\theta_H - g(s_H, \theta_H)) + (1-\lambda)(E[\theta] - g(s^*, \theta_L)) - (E[\theta] - \lambda g(s^*, \theta_H) - (1-\lambda)g(s^*, \theta_L))$$

where $b(s^*)$ represents the cost of information at which the sender is indifferent between not acquiring information followed by s^* and deviating to acquiring information followed by s^* as the low type and s_H as the high type; s_H, in turn, is such that the uninformed sender is indifferent between $(s^*, E[\theta])$ and (s_H, θ_H):

$$(9) \qquad E[\theta] - E\left[g(s^*, \theta)\right] = \theta_H - E\left[g(s_H, \theta)\right].$$

In the quadratic case b did not depend on s^*, resulting in $b(0) = b(\bar{s}^*) = c_N$ and the generic uniqueness of equilibria that survive the refinement. This begs two questions. How does b depend on the primitives of the environment and when is it constant?

Lemma 8 *The sign of $\frac{db}{ds^*}$ is the opposite of the sign of $\frac{d^2 \log(g_s(s,\theta))}{dsd\theta}$. In particular, the cost function b does not depend on s^* when $\frac{d^2 \log(g_s(s,\theta))}{dsd\theta} = 0$.*

The cost function is decreasing in s^* if the marginal cost of sender's action, c_s, is log-supermodular, and increasing if it is log-submodular. It is independent of s^* when $\frac{d^2 \log(g_s(s,\theta))}{dsd\theta} = 0$, for all s.

The change in the indifference cost, b, with respect to the amount of signaling, s^*, can be written as

$$\frac{db}{ds^*} = -\lambda(g_s(s_H, \theta_H) - g_s(s^*, \theta_H)) + \lambda g_s(s^*, \theta_H)\left(1 - \frac{ds_H}{ds^*}\right),$$

where the first term, termed *the marginal cost effect*, captures the increase in the cost if the distance between s^* and s_H remained constant and the second effect, *the shift effect*, accounts for the change in the distance between s^* and s_H as s^* increases. The high type's action after the deviation to information acquisition, s_H, changes with s^* as given by

$$\frac{ds_H}{ds^*} = \frac{\lambda g_s(s^*, \theta_H) + (1-\lambda) g_s(s^*, \theta_L)}{\lambda g_s(s_H, \theta_H) + (1-\lambda) g_s(s_H, \theta_L)}.$$

In the quadratic case g_s is increasing in s and, therefore, the marginal cost effect negative. However, the distance between s^* and s_H shrinks as s^* increases ($\frac{ds_H}{ds^*} < 1$), making the shift effect positive. While in the quadratic case the two effects offset each other, this is not the case in general. The log-supermodularity condition provides the exact dividing line between the two forces.

A more precise characterization of the intermediate case can be provided.

LEMMA 9 *The equality* $\frac{d^2 \log(g_s(s,\theta))}{ds d\theta} = 0$ *holds for all* s *if and only if there exist real functions* f *and* g *such that* $g_s(s, \theta) = f(s) g(\theta)$.

PROOF: See Appendix for the proof. □

REMARK 1 If g_s is log-supermodular, i.e., if $\frac{\partial \ln(g_s(s,\theta))}{\partial s \partial \theta} > 0$, then $b(s)$ is a decreasing function. Because b is a continuous function, the inverse function, $b^{-1}(c)$, is well defined on $[b(\bar{s}^*), b(0)]$ and is onto $[0, \bar{s}^*]$.

The results on equilibrium outcomes with information acquisition also extend to the more general environment. There exist a $c_I > 0$ such that all equilibria with information acquisition can be refined away for $c > c_I$, while for $c < c_I$ only the most efficient equilibrium with information acquisition survives. More precisely, c_I is determined by the requirement that the uninformed sender's indifference curve through the high type's action and the low type's indifference curve through his own action in the most efficient equilibrium outcome with information acquisition intersect at $r = E[\theta]$. The final piece of the puzzle in the analysis under quadratic cost of signaling was the result $c_I = c_N$, which established a threshold above which only an equilibrium without information survives and below it only the efficient equilibria with information acquisition remain. An analogous, yet slightly weaker, result holds here.

REMARK 2 The following equality holds: $c_I = b(\bar{s}^*)$; that is, the highest cost at which an equilibrium with information acquisition survives the refinement, c_I, coincides with the cost at which the equilibrium outcome without information acquisition and the largest amount of signaling, $b(\bar{s}^*)$, survives.

The only remaining equilibria are the ones in which the sender randomizes over information acquisition. Such equilibria survive the refinement only at one cost, denoted c_R. Given the above it is not too surprising that $c_R = c_I = b(\bar{s}^*)$.

The above derived result can be assembled to paint the full picture. If the fundamentals are such that function b is constant, see Lemma 8, one obtains the same result as in the previous section. Lemma 9 further implies that the characterization relied on multiplicative separability of the cost function of signaling, rather than on the cost being quadratic.

The most interesting characterization obtains when b is decreasing.

PROPOSITION 3 *If g_s is log-supermodular, $b(\cdot)$ decreasing, generically a unique equilibrium survives the refinement:*

- *For $c < b(\bar{s}^*)$ it is the most efficient equilibrium outcome with information acquisition.*
- *For $c \in (b(\bar{s}^*), b(0))$, it is the equilibrium outcome in which the sender does not acquire information and chooses $b^{-1}(c)$.*
- *For $c > b(0)$, it is the equilibrium outcome in which no information is acquired and the sender chooses the least costly action, 0.*

When b is decreasing generically a unique equilibrium outcome survives the refinement. When the information is cheap only the most efficient equilibrium outcomes with information acquisition survive. The familiar pattern arises, as the cost of information increases, the high type burns more and more surplus to assure the receiver that he inquired information. This persists up to the cost c_I at which the uninformed's indifference curve through the high type's option and the low type's indifference curve through $(0, \theta_L)$ intersect at a point with the vertical component $E[\theta]$. At the same cost, an equilibrium without information acquisition and $s = \bar{s}^*$ survives too.

The most interesting behavior is exhibited when the cost of information exceeds c_I. The only equilibrium outcome that survives the refinement is the one in which the sender does not acquire information and chooses $s = b^{-1}(c)$. Strikingly, the uninformed sender undertakes costly signaling in order to convince the receiver that he is indeed uninformed rather than an informed sender with a low type. Since b is decreasing, the uninformed sender engages in less and less signaling as the cost of information increases. When the cost of information increases above $b(0)$, in the only equilibrium outcome that survives the refinement, the sender does

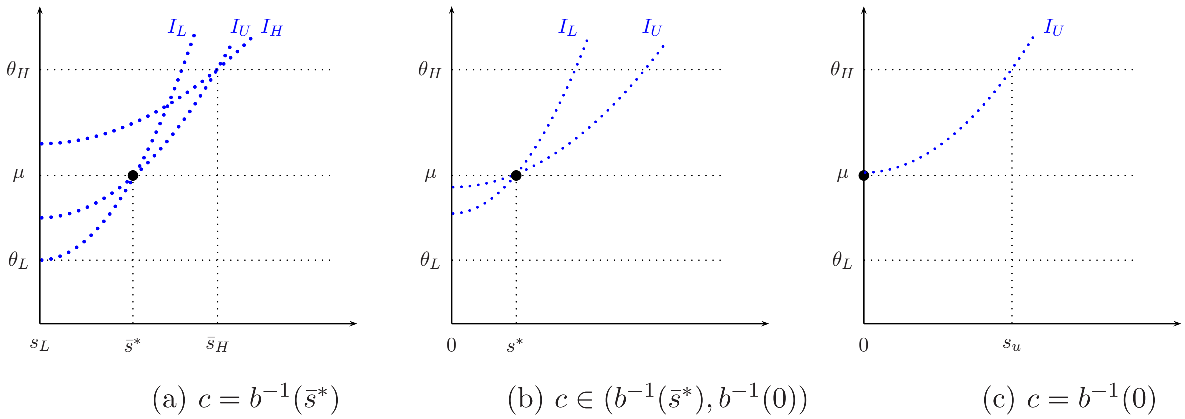

Figure 9: The surviving equilibria for $c \in [b^{-1}(\bar{s}^*), b^{-1}(0)]$ when $b(\cdot)$ is decreasing.

not acquire information and chooses the least costly action.

Proposition 4 *If g_s is log-submodular ($b(\cdot)$ increasing):*

- *A unique equilibrium outcome survives the refinement for $c < b(0)$: the most efficient equilibrium with information acquisition.*
- *A unique equilibrium outcome survives the refinement for $c > b(\bar{s}^*)$: the one in which no information is acquired and the sender chooses the least costly action, 0.*
- *For $c \in [b(0), b(\bar{s}^*)]$ multiple equilibrium outcomes survive NWBR. These equilibrium outcomes include the most efficient equilibrium outcome with information acquisition, and the equilibrium outcome with no information acquisition and $s = b^{-1}(c)$.*

When b is increasing, there is a multiplicity of equilibrium outcomes that survive the refinement in the intermediate region of cost $c \in (b(\bar{s}^*), b(0))$. Equilibrium outcomes that survive are the most efficient equilibrium outcome with information acquisition, no information acquisition equilibrium outcome with $s^* = 0$ and no information acquisition equilibrium outcome with $b^{-1}(c)$. Despite the multiplicity in the intermediate region of cost, for low costs of information ($c < \min\{b(\bar{s}^*), b(0)\}$) a unique equilibrium outcome survives the refinement—the most efficient equilibrium outcome with information acquisition, thus reversing the no information acquisition result obtained under observable information acquisition.

Above we only considered the cases where b is monotonic. More general statements can be made. First, if b is at any point increasing, there will be a multiplicity of equilibria that survive the refinement in the increasing region. Given that b is continuous, this implies that b non-increasing is also necessary for the generic uniqueness of equilibria; within the realm of single-crossing signaling cost functions g. Second, even if b is non-monotonic only the most efficient outcome with information acquisition survives the refinement for low enough c. Since b is continuous and relevant on the interval $[0, \bar{s}^*]$, it attains a minimum. Moreover, it is easy to verify $b(s^*) > 0$ for all $s^* \in [0, \bar{s}^*]$, thus the minimum on the interval is strictly above 0. The above analysis then implies that a unique equilibrium survives the refinement for $c < \min\{b(s^*) : s^* \in [0, \bar{s}^*]\}$.

Proposition 5 *For any $c < \min\{b(s^*) : s^* \in [0, \bar{s}^*]\}$, only the most efficient equilibrium outcome with information acquisition survives the refinement.*

Extensions

This section examines some of the assumptions imposed at the outset and shows how the results extend to more general environments.

Partial Information. The analysis thus far was conducted under the assumption that acquired information reveals the state of the world completely. Alternatively, one could study an environment in which after information acquisition the sender observes one of two signals: ϕ_h, ϕ_l. After signal ϕ_i, $i \in \{l, h\}$, which is observed with ex ante probability p_i, the sender's posterior is π_i. Let λ be the prior belief that the state is θ_H. Naturally:

$$p_h \pi_h + p_l \pi_l = \lambda.$$

In what follows we verify that the results extend for the environment analyzed in section Quadratic Cost; indeed, we verify the results for a somewhat more general environment. In particular, assume that along the usual assumptions imposed on g, it can be written as a product of two functions f and h: $\tilde{g}(s, \theta) = h(s)f(\theta)$ for some functions f and h as covered by Lemma 9.[10]

Abusing notation slightly, we write the cost function as a function of the sender's belief

[10] Observe that if $\tilde{g}(s, \theta) = h(s)f(\theta)$, then $g_s(s, \theta) = h'(s)f(\theta)$.

π that the state is θ_H:

$$\tilde{g}(s,\pi) := \pi f(\theta_H)h(s) + (1-\pi)f(\theta_L)h(s).$$

Taking the cross-partial: $\tilde{g}_{s\pi}(s,\pi) = h'(s)\left(f(\theta_H) - f(\theta_L)\right) < 0$, because $h'(s) > 0$, and $f(\theta_H) < f(\theta_L)$. Therefore, $\tilde{g}$ satisfies the single-crossing assumption. After the signal ϕ_h the sender's indifference curve is flatter than the uninformed agent's, after ϕ_l it is steeper.

The results that only depended on the single-crossing arguments—lemmata 4,6, and 7—extend to the environment with partial information. The main question is whether suitably defined version of $b(\cdot)$ is constant, as in Lemma 5.

Let μ_h and μ_l be the expected value of the state conditional on the sender's signal being high and low, respectively. Define

$$\tilde{b}(s) = p_h\left(\mu_h - \tilde{g}(\pi_h, s_u)\right) + p_l\left(E(\theta) - \tilde{g}(\pi_l, s)\right) - E(\theta) + \tilde{g}(\lambda, s), \tag{10}$$

where s_u solves the equality:

$$E(\theta) - \tilde{g}(\lambda, s) = w_h - \tilde{g}(\lambda, s_u), \tag{11}$$

to be the partial information analogue of $b(\cdot)$.

Lemma 10 *Function $\tilde{b}(\cdot)$ is constant on $[0,\bar{s}]$.*

Proof: See Appendix for the proof. □

More than two types. Suppose that there are n states of the world $\Theta = \{\theta_1, \theta_2, ..., \theta_n\}$, with $\theta_1 < \theta_2 < ... < \theta_n$ and the prior probability of state θ_i being $\lambda_i \in (0,1)$. As initially, we assume that when the sender acquires information, he learns the state perfectly. We argue the robustness of the above obtained results in two ways. First, we show that the equilibrium outcome with no information acquisition and no signaling can be refined away when information is cheap enough. Second, we argue that an outcome with information acquisition cannot be refined away for sufficiently small cost of information.

Lemma 11 *There exists a $\gamma > 0$ such that the equilibrium outcome with no information acquisition and no signaling can be refined away for all $c < \gamma$.*

PROOF: Fix the equilibrium outcome with no information acquisition and no signaling. The receiver after observing $s = 0$ responds with $r = E[\theta]$. Define s_i, $i \in \{1, 2, ..., n\}$ to be the intersection of type θ_i's indifference curve through $(0, E[\theta])$ with the ray θ_n:

$$E[\theta] - c(0, \theta_i) = \theta_n - c(s_i, \theta_i).$$

The assumptions on c imply $s_1 < s_2 < ... < s_n$. In the same fashion we can define s_u as the intersection of the uninformed agent's indifference curve with θ_n. We will consider two cases depending on whether $s_{n-1} \leq s_u$ or $s_{n-1} > s_u$; clearly $s_u < s_n$.

First, suppose that $s_{n-1} \leq s_u$. Then actions above s_u are NWBR for types θ_1 through θ_{n-1}, nor for the uninformed sender. Indeed, a necessary condition for equilibria with no information acquisition and $s = 0$ is that the receiver's response is never above the uninformed sender's indifference curve. At the same time, any type $\theta_1, ..., \theta_{n-1}$'s indifference curve through $(0, E[\theta])$ is not below the uninformed sender's indifference curve. Therefore, any strategy where the sender acquires information and type θ_i, $i \in \{\theta_1, ..., \theta_{n-1}\}$ plays an action $s > s_u$ is dominated by the same strategy with a modification that the same type plays $s = 0$. In a game obtained after erasing any strategy where any type of sender except for θ_n plays $s > s_u$, the receiver should respond to any $s > s_u$ with θ_n. In the newly obtained game, the sender if he were to acquire information optimally chooses $s = 0$ if type $\theta_1, ..., \theta_{n-1}$ and s_u if θ_n. In addition, since type θ_{n-1} is indifferent between $(0, E[\theta])$ and (θ_n, s_{n-1}), the type θ_n strictly prefers the latter point due to single-crossing. Therefore, as long as the cost of acquiring information is smaller than the probability of type θ_n multiplied by the benefit of the high type from choosing (θ_n, s_{n-1}) over $(0, E[\theta])$, the sender can profitably deviate in the pruned game.

If $s_{n-1} > s_u$, then actions above s_{n-1} are NWBR for any type except for possibly θ_n. In the game obtained after erasing any strategy where types θ_1 through θ_{n-1} and the uninformed sender play $s > s_{n-1}$, the receiver should reply to an $s > s_{n-1}$ with $r = \theta_n$. But then the sender has an incentive to deviate to acquiring information followed with $s = 0$ except for s_{n-1} if $\theta = \theta_n$, if the cost of information is low enough. □

The above result establishes that the no information acquisition and no signaling outcome can be refined away for small costs of information. The proof establishes that the result holds for the lowest levels of cost, rather than characterizing all such costs.

Unlike in the two type case, when there are more than two types, information acquisition followed by the Riley outcome can be sustained as an equilibrium, provided information is cheap.

Lemma 12 *There exists a $\gamma_R > 0$ such that information acquisition followed by the Riley outcome can be sustained as an equilibrium for every $c < \gamma_R$.*

Proof: See Appendix for the proof. □

In the two type case, information acquisition followed by the Riley outcome can not be implemented. The sender can deviate to not acquiring information and choosing the high type's option which enables him to replicate the same payoff as with information acquisition state by state without paying for information; the low type is indifferent between own option and the high type's. In the case of more than two types, each type is indifferent among at most two options. Therefore, deviating to not acquiring information enables the sender to replicate the payoff of information acquisition in at most two states. In the remaining states, the sender is strictly better off acquiring information, which is the origin of the value of information.

Lemma 13 *Suppose there are at least three types. There exists a $\hat{c} > 0$ such that information acquisition followed by the Riley outcome cannot be refined away for $c \leq \hat{c}$.*

Proof: See Appendix for the proof. □

Concluding Remarks

We study information acquisition in signaling. In a result reminiscent of Grossman and Stiglitz (1980) the sender never acquires costly information if the decision to acquire it is observable. However, if the decision whether to acquire information is covert and information cheap, the sender does acquire it in the unique equilibrium that survives a form of never weak best response refinement. Interestingly, for low costs of information, as the information becomes costlier the high type sender burns more and more surplus in order to convince the receiver that he indeed acquired it—with costlier information, the receiver requires more convincing. Even as strong a refinement as NWBR does not guarantee uniqueness under single crossing, but additional conditions which guarantee a single outcome survives the refinement are provided. Of note is that in some cases the only outcome surviving the refinement is the one in which the sender does not acquire information yet undertakes a strictly costly action, signaling to the receiver that he is not the informed sender who learned that he is of the low type.

We study an environment where the sender and the receiver have opposing preferences over the receiver's action. One could alternatively consider an environment in which the sender's preferred receiver's action depends on the state of the world and parametrize their disagreement. The conjecture is that as the two player's preferences become more aligned, the sender's incentive to acquire information grows.

While signaling games have been extensively studied, a comprehensive study of signaling in more general games is by and large an uncharted territory with few exceptions; for a recent take see In and Wright (2017). In separate work we intend to study a game in which the sender undertakes an investment in his ability (productivity) with a stochastic outcome, and then undertakes a signaling action. The receiver observes the signaling action, but not the investment.

References

Akerlof, G. A. (1970): "The Market for "Lemons": Quality Uncertainty and the Market Mechanism," *The Quarterly Journal of Economics*, 84, 488–500.

Alós-Ferrer, C. and J. Prat (2012): "Job market signaling and employer learning," *Journal of Economic Theory*, 147, 1787–1817.

Banks, J. S. (1992): "Monopoly Pricingand Regulatory Oversight," *Journal of Economics & Management Strategy*, 1, 203–233.

Banks, J. S. and J. Sobel (1987): "Equilibrium selection in signaling games," *Econometrica: Journal of the Econometric Society*, 647–661.

Bester, H., M. Lang, and J. Li (2019): "Signaling versus auditing," .

Bester, H. and K. Ritzberger (2001): "Strategic pricing, signalling, and costly information acquisition," *International Journal of Industrial Organization*, 19, 1347–1361.

Cho, I.-K. and D. M. Kreps (1987): "Signaling games and stable equilibria," *The Quarterly Journal of Economics*, 102, 179–221.

Cho, I.-K. and J. Sobel (1990): "Strategic stability and uniqueness in signaling games," *Journal of Economic Theory*, 50, 381–413.

Daley, B. and B. Green (2014): "Market signaling with grades," *Journal of Economic Theory*, 151, 114–145.

Feltovich, N., R. Harbaugh, and T. To (2002): "Too cool for school? Signaling and countersignaling," *RAND Journal of Economics*, 33.

Fudenberg, D. and J. Tirole (1991): *Game Theory.*

Grassi, S. and C.-t. A. Ma (2016): "Information acquisition, referral, and organization," *The RAND Journal of Economics*, 47, 935–960.

Grossman, S. J. and J. E. Stiglitz (1980): "On the impossibility of informationally efficient markets," *The American economic review*, 70, 393–408.

In, Y. and J. Wright (2017): "Signaling private choices," *The Review of Economic Studies*, 85, 558–580.

Kohlberg, E. and J. Mertens (1986): "On the strategic stability of equilibria," *Econometrica*, 54, 1003–1037.

Mailath, G. J., M. Okuno-Fujiwara, and A. Postlewaite (1993): "Belief-based refinements in signalling games," *Journal of Economic Theory*, 60, 241–276.

Riley, J. G. (2001): "Silver signals: Twenty-five years of screening and signaling," *Journal of Economic literature*, 39, 432–478.

Sobel, J. (2009): "Signaling Games." *Encyclopedia of Complexity and Systems Science*, 19, 8125–8139.

Spence, M. (1973): "Job Market Signaling," *The Quarterly Journal of Economics*, 87, 355–374.

Stahl, K. and R. Strausz (2017): "Certification and market transparency," *The Review of Economic Studies*, 84, 1842–1868.

A. Appendix

Notation: In the sender-receiver game that we study, a pure strategy of the sender specifies the sender's information acquisition action (I for acquiring information, U for not acquiring information), and a signaling action for each of the three types, θ_H, θ_L and θ_u. However, for purposes of applying the NWBR refinement, specifying the signaling action for a type that does not exist on the path generated by the information acquisition action of a strategy is redundant. Hence, we refer to a pure strategy in which the sender that acquires information, and chooses s if low type, and s' if high type as simply (I, s, s'), and a pure strategy of the sender that does not acquire information, and chooses s as simply (U, s).

Proof of Lemma 4: Consider an equilibrium outcome where the sender does not acquire information and undertakes an $s^* > \bar{s}^*$. We claim that actions between $\bar{s}^*$ and s^* are never week best response for the low type. For the prescribed outcome to be an equilibrium outcome it has to be the case that the receiver's response $r(s)$ is never above the uninformed's indifference curve through $(s^*, E[\theta])$. On the interval $(\bar{s}^*, s^*)$ the low type's indifference curve through $(0, \theta_L)$ is strictly above the uninformed sender's indifference curve through $(s^*, E[\theta])$. This is due to $\bar{s}^*$ being defined so that the low type is indifferent between $(0, \theta_L)$ and $(\bar{s}^*, E[\theta])$. Therefore, any sender's strategy where he acquires information and the low type plays $s \in (\bar{s}^*, s^*)$ with positive probability is dominated by the same strategy where the low type chooses 0 instead of s. Having established that $s \in (\bar{s}^*, s^*)$ is NWBR for the low type, one can prune away all the strategies in which the low type plays these actions. Since in the newly obtained game, actions $s \in (\bar{s}^*, s^*)$ can only be chosen by the uninformed agent or the high type, the receiver best responds with an $r \in [E[\theta], \theta_H]$. But then any action in $(\bar{s}^*, s^*)$ represents a profitable deviation for the uninformed sender.

Next we show that any outcome where the sender does not acquire information and chooses an $s^* \in (0, \bar{s}^*]$ can be refined away for all but one cost of information. Fix an equilibrium outcome with no information acquisition and $s^* \in (0, \bar{s}^*]$. Let s_u be the intersection of the uninformed agent's indifference curve through $(s^*, E[\theta])$ with the ray θ_H:

$$E[\theta] - E\left[\frac{s^{*2}}{\theta}\right] = \theta_H - E\left[\frac{s_u^2}{\theta}\right].$$

Actions $s' > s_u$ are never weak best response for the uninformed sender, as well as for the low type. Any strategy where the sender acquires information and the low type chooses $s \geq s_u$ is dominated by the same strategy with the difference that the low type chooses s^*.

In the newly obtained game, the best option for the high type is (s_u, θ_H); formally, the high type can approach this payoff. This follows from the observation that $(s^*, E[\theta])$ and (s_u, θ_H) are connected by the uninformed's indifference curve and that the high type's indifference curve is flatter.

Next, we argue that if the equilibrium outcome is not to be refined away, the sender must be indifferent between his equilibrium play and deviation towards acquiring information followed by choosing s^* when the low type and s_u (to which the receiver responds with θ_H) when the high type. Suppose not, if the payoff with acquiring information was larger, then it would constitute a profitable deviation in the above-derived game. Strictly speaking $(I, s^*, s_u + \epsilon)$ would be a profitable deviation for an ϵ small enough. On the other hand, if the payoff with information acquisition is strictly lower, there exists an $\epsilon > 0$ such that action in $(s^* - \epsilon, s^*)$ are NWBR for the low type; the high type always optimally chooses s_u to which the receiver responds with θ_H. Indeed, since in every equilibrium with the outcome, the receiver's response cannot result in a pair above the uninformed's indifference curve through $(s^*, E[\theta])$, by moving slightly below s^* when the low type, the sender cannot increase his payoff discontinuously. One can, therefore, prune away all the strategies where the low type plays an action in $(s^* - \epsilon, s^*)$. In thus obtained game the receiver will respond to any $s \in (s^* - \epsilon, s^*)$ with an $r \in [E[\theta], \theta_H]$. But then not acquiring information and choosing an action in $s \in (s^* - \epsilon, s^*)$ is a profitable deviation for the sender.

We established that if an equilibrium outcome without information acquisition and $s^* \in (0, \bar{s}^*)$ is to survive the refinement the sender ought to be indifferent between not acquiring information followed by s^* and acquiring information followed by s^* if low type and s_u if high type:

$$E[\theta] - E\left[\frac{s^{*2}}{\theta}\right] = \lambda\left(\theta_H - \frac{s_u^2}{\theta_H}\right) + (1-\lambda)\left(E[\theta] - \frac{s^*}{\theta_L}\right) - c.$$

The cost that makes the sender indifferent for a given s^*, denoted $b(s^*)$, is

$$b(s^*) = \lambda\left(\theta_H - \frac{s_u^2}{\theta_H}\right) + (1-\lambda)\left(E[\theta] - E\left[\frac{s^{*2}}{\theta}\right]\right).$$

Verification. The final step of the proof is to verify that any outcome where the sender does not acquire information and chooses $s^* \leq \bar{s}^*$, outcome o, survives the refinement when the cost of information is $b(s^*)$. On the way to a contradiction, suppose there is a sequence of pruning with respect to the outcome o such that in the stages $i = 1, ..., k-1$, o is an equilibrium outcome in Γ^i, but not in Γ^k.

We start with some observations: First, the set of equilibrium receiver strategies that lead to outcome o in the games $\{\Gamma^i\}_{i=1,...,k-1}$ are weakly below the indifference curve of the uninformed type's indifference curve that passes through $(s^*, E(\theta))$ (called IC_U), and weakly below the ray θ_H. These are necessary conditions for the optimality of the strategy (U, s^*) for the sender, and an implication of the equilibrium condition requiring that the receiver best responds to some beliefs. Second, we claim that the strategy in which the sender acquires information, chooses s^* as the low and s_u as the high type, for short (I, s^*, s_u), cannot be pruned away in any step. The sender is indifferent between (U, s^*), and (I, s^*, s_u) when $r(s^*) = E(\theta)$, and $r(s_u) = \theta_H$. Therefore, the strategy could be pruned away only if the receiver could not assign the belief to θ_H after observing s_U, or equivalently, if all the strategies where the high type plays s_u have been pruned already, with the possible exception of (I, s_u, s_u). That is, for (I, s^*, s_u) to be NWBR it would have had to been pruned already.

We argue that no s can be a part of a profitable deviation from the equilibrium outcome after any finite sequence of pruning. Two cases are to be considered: deviations in $[0, s^*)$ and deviations in (s^*, s_u). Having shown that (I, s^*, s_u) can never be erased, any deviation including an $s' > s_u$ is dominated by a strategy where s' is replaced by s_u.

Case 1: $s \in (s^*, s_u)$. If for some $s \in (s^*, s_u)$ the strategy (U, s) is pruned away in some stage

$l < k$, then it must be the case that in Γ_l the receiver does not assign positive probability to θ_H after observing s, i.e., the strategies (I, s', s) for every s', except for possibly $s' = s$, have been pruned at some earlier stage $l' < l$. To see this, suppose by a way of contradiction that there exists some $s' \neq s$ such that (I, s', s) is present in Γ^l. Then, the receiver can attach a positive probability to s being played by the high type and the set of all equilibria that leads to outcome o contains an equilibrium with $r(s) = \tilde{r}$, where $(s, \tilde{r})$ is on IC_U. But this contradicts that (U, s) is never a best reply in any equilibrium that leads to outcome o in Γ^l.

The above implies that in Γ^k, either the strategies (U, s) and (I, s', s) exist for some $s' \neq s$, or all such strategies are pruned away prior to stage k. In either case, some receiver strategy in which $r(s) \leq \tilde{r}$ is consistent with the receiver's sequential rationality constraint, and hence the sender does not have a profitable deviation to (U, s). He also does not have a profitable deviation to acquiring information and one of the two types playing s, since on the interval under study IC_U is strictly below the low type's indifference curve through $(s^*, E[\theta])$ and the high type's through (s_u, θ_H).

Case 2: $s < s^*$. On this interval IC_U is below $E[\theta]$. Therefore, for a profitable deviation to occur one would need to prune all the strategies in which the low type plays s, with the possible exception of (I, s, s), while leaving at least one strategy where some other type plays it; that would force $r(s) \geq E[\theta]$ and in particular above IC_U. However, as we show in the following paragraph, if all the strategies (I, s, s'), $s' \neq s$, have been pruned at or before game Γ_l, $l < k$, then the strategy (U, s) and all strategies (I, s'', s) must have been pruned by some earlier stage $l' < l$. Thus, if all the strategies in which θ_L plays s have been pruned, no strategies in which any type plays s are left, and s cannot represent a profitable deviation.

We now prove the assertion that if all the strategies (I, s, s'), $s' \neq s$, have been pruned at or before game Γ_l, $l < k$, then the strategy (U, s) and all strategies (I, s'', s) must have been pruned by some earlier stage $l' < l$. Suppose on the way to a contradiction that at stage l, a strategy (I, s, s'), for some $s' \neq s$, is available along with either (U, s) or (I, s'', s) for some s''. Then, in stage l, the set of all equilibria that leads to outcome o is nonempty, and includes a receiver strategy in which $(s, r(s))$ is on IC_L, $r(s) \geq \theta_L$, and $r(s_u) = \theta_H$. The last property follows from our initial observation that (I, s^*, s_u) cannot be pruned, and the penultimate property from $s^* \leq \bar{s}^*$. Because the sender is indifferent between her payoff in o, and (I, s^*, s_u) when $r(s^*) = E(\theta)$ and $r(s_u) = \theta_H$, in stage l, the strategy (I, s, s_u) is not pruned, leading to a contradiction. □

Proof of Lemma 5: s_u is defined by the indifference condition:

$$E[\theta] - s^{*2}\left(\frac{\lambda}{\theta_H} + \frac{1-\lambda}{\theta_L}\right) = \theta_H - s_u^2\left(\frac{\lambda}{\theta_H} + \frac{1-\lambda}{\theta_L}\right).$$

Which can be rewritten as

$$s_u^2 = s^{*2} + \frac{\theta_H - E[\theta]}{\frac{\lambda}{\theta_H} + \frac{1-\lambda}{\theta_L}}.$$

Plugging the above expression into the formula for $b(s^*)$ yields

$$b(s^*) = \lambda(\theta_H - E[\theta])\left(1 - \frac{1}{\theta_H\left(\frac{\lambda}{\theta_H} + \frac{1-\lambda}{\theta_L}\right)}\right).$$

Defining c_N to be equal to the right hand-side of the above expression establishes the first part of our result.

We showed that equilibria with no information acquisition and strictly costly signaling exist only when the cost of information is precisely c_N. We are only left to consider the equilibrium outcome where the sender does not acquire information and chooses $s^* = 0$. As above, one can define s_u and argue that actions above s_u are NWBR for the low type or the uninformed sender. The difference is that the uninformed sender cannot deviate to actions below s^*, thus we only need that $c \geq c_N$; if $c < c_N$ then the sender can profitably deviate towards acquiring information in the game obtained after pruning actions above s_u for the uninformed and the low type.

Finally, we verify that the equilibrium outcome without information acquisition and $s^* = 0$, henceforth outcome o, survives the refinement for $c \geq c_N$. Towards that goal, let IC_U be the uninformed type's indifference curve that passes through the point $(0, E(\theta))$. Any receiver's strategy with $r(0) = E(\theta)$ that never goes above IC_U supports the outcome o as an equilibrium. Indeed, let s_H be the action at which IC_U intersects the ray θ_H. As long as the receiver's response is not above IC_U the low type would choose $s^* = 0$ due to single-crossing, while the high type's best possible option would be to choose s_H and the receiver to reply with θ_H. The cost of information c_N is precisely the cost at which the sender is indifferent between his prescribed equilibrium strategy and the deviation $(I, 0, s_H)$. Since the deviation is an optimal deviation with information acquisition the sender does not have a profitable deviation towards acquiring information.

Suppose on the way to a contradiction that there was a sequence of pruning such that in Γ^k, the outcome o would fail to be an equilibrium outcome. For this to be the case, there would have to exist some $s < s_H$ such that $r(s)$ could not be sustained as an off equilibrium receiver action in Γ^k. The discussion in the previous paragraph implies that $r(s)$ would need to be above IC_U.The only way this could occur is if action s was available only to the high type and, therefore, $r(s) = \theta_H$. The strategies in the game Γ_k in which the high type plays s would also have had to be available in all the games Γ_l, $l < k$; that is, they could not have been pruned earlier.

For strategies where the sender remains uninformed and plays s to be pruned away, $r(s)$ would need to be strictly below IC_U, or differently, the receiver should not be able to assign the belief that he is facing θ_H. For that, all the strategies where θ_H plays s, with a possible exception of (I, s, s), should have been previously pruned away. In summary, to arrive at the game Γ_k where s is available only to θ_H, one would have had to prune away the strategies where θ_U plays s, which, in turn, would require that s was pruned for θ_H even earlier. This, of course, can not be. □

Proof of Lemma 6: The proof of the result is somewhat lengthy and, thus, broken down in steps: i) we argue that only separating equilibria after information acquisition can potentially survive the refinement at $c = 0$ (strictly speaking, we argue that all the other equilibria with information acquisition can be refined away); ii) only the equilibria with information acquisition in which the sender is indifferent between acquiring information (equilibrium play) and deviating to not acquiring information followed by pretending to be the high type can potentially survive the refinement; iii) we define a c_I and argue that no equilibrium with information acquisition survives the refinement for $c > c_I$; iv) we verify that the remaining equilibria—one for each $c \leq c_I$—survive the refinement.

Step 1: Let $c = 0$ and fix an equilibrium outcome with information acquisition in which the two types pool with positive probability. Denote such pooling action by $\tilde{s}$, and the receiver's response with $\tilde{r}$; notice that $\tilde{r} < \theta_H$. For $\tilde{s}$ to be a part of equilibrium, the receiver's response to $s' > \tilde{s}$ cannot be above the high type's indifference curve through $(\tilde{s}, \tilde{r})$. Due to single-crossing nature of the preferences both the low type's as well as uninformed sender's indifference curves through $(\tilde{s}, \tilde{r})$ are strictly above the high type's for $s' > \tilde{s}$. More precisely, for $s' > \tilde{s}$ every equilibrium with the fixed outcome is such that:

$$\begin{aligned} r(s') - \frac{s'^2}{\theta_L} &< r(s') - E\left[\frac{s'^2}{\theta}\right] \\ &< r(s') - \frac{s'^2}{\theta_H} \\ &\leq r(\tilde{s}) - \frac{\tilde{s}^2}{\theta_H}, \end{aligned}$$

where the last inequality is guaranteeing that the high type does not have a profitable deviation from his equilibrium action $\tilde{s}$. Thus, actions above $\tilde{s}$ are NWBR for the low type or the uninformed sender. After removing all the strategies where the uninformed sender or the low type play actions above $\tilde{s}$, one obtains a game in which in any equilibrium the receiver should respond to an action $s' > \tilde{s}$ with θ_H. But then the high type would have a profitable deviation just above $\tilde{s}$.

Step 2: Fix a $c \geq 0$. We claim that if an equilibrium with information acquisition is to survive the equilibrium refinement it has to be the case that the sender is indifferent between following the equilibrium strategy and deviating to not acquiring information followed by the high type's equilibrium action; if an equilibrium outcome at c exists to start with. Lemma 2 established that for $c > 0$ in any equilibrium with information acquisition the two types separate themselves and, moreover, the low type chooses $s_L = 0$ and strictly prefers his action to the high type's. In the first step we showed that an analogous statement must hold for $c = 0$ if the equilibrium is to not be refined away. Fix an equilibrium outcome with information acquisition and suppose, contrary to our claim, that the sender strictly prefers equilibrium play to not acquiring information followed by pretending to be the high type. Consider the set of all equilibria that lead to the outcome. Because the low type strictly prefers his equilibrium play, $s_L = 0$, to the high type's, s_H, and because the sender strictly prefers acquiring information, there exists an $\epsilon > 0$ such that $s \in (s_H - \epsilon, s_H)$ are NWBR for the low type nor for the uninformed sender; that is, any strategy with no information acquisition and $s \in (s_H - \epsilon, s_H)$ yields a strictly smaller payoff than the equilibrium strategy. Hence, these strategies can be pruned away. In the game obtained after the pruning, actions in $(s_H - \epsilon, s_H)$ can only be played by the high type, therefore in any equilibrium that leads to the outcome the receiver should respond to $s \in (s_H - \epsilon, s_H)$ with $r = \theta_H$. But then, the outcome is not an equilibrium of the reduced game, namely, type θ_H has a profitable deviation in $(s_H - \epsilon, s_H)$.

The above establishes that if an outcome with information acquisition is to survive an equilibrium refinement, then it has to be the case that the sender is indifferent between acquiring and not acquiring information. We also know that any such outcome is separating, therefore we can conclude that this is the most efficient separating outcome—any separating outcome where the high type burns less surplus, at the fixed c, would have sender deviating to not acquiring information. We are yet to establish the range of c for which such equilibria survive the refinement.

Step 3: The most efficient (separating) equilibria with information acquisition have the following form. At $c = 0$ the outcome where the agent is indifferent between acquiring and not acquiring information is the Riley outcome. As c increases so does s_H, as can be seen from the sender's indifference between acquiring and not acquiring information (and pretending to be the high type):

$$\theta_L - \frac{c}{1-\lambda} = \theta_H - \frac{s_H^2}{\theta_L}.$$

In particular, to each c corresponds an s_H. Next we establish that such equilibria with s_H above some threshold can be refined away, or equivalently, such equilibria can be refined away for c above some threshold.

Fix a $c > 0$ and the equilibrium outcome with information acquisition such that the sender is indifferent between acquiring information and not acquiring information followed by mimicking the high type. Denote the intersection of the low type's indifference curve through $(0, \theta_L)$ and the uninformed sender's indifference curve through (s_H, θ_H) by (s_{LU}, r_{LU}). We argue that an equilibrium can be refined away if $r_{LU} < E[\theta]$. This establishes an upper bound on s_H and thus on c.

If $r_{LU} < E[\theta]$, then there exist an $\epsilon > 0$, such that for actions in $(s_{LU}, s_{LU} + \epsilon)$ the receiver is responding with $r < E[\theta]$ in every equilibrium; in equilibrium his responses cannot be above the uninformed agent's indifference curve. However, the actions in question are strictly below the low type's indifference curve by the definition of (s_{LU}, r_{LU}) and the single-crossing property of our environment, and as such, NWBR for the low type. But then, in a game obtained after pruning the said NWBR strategies, in any equilibrium the receiver would have to reply to $(s_{LU}, s_{LU} + \epsilon)$ with at least $E[\theta]$, making a deviation to not acquiring information followed by one of those actions profitable. Therefore, if the equilibrium outcome is to survive the NWBR criterion, it must be the case that $r_{LU} \geq E[\theta]$. Since r_{LU} is decreasing in c, the highest value of information acquisition cost where equilibria with indifference and information acquisition could possibly survive NWBR, c_I, is such that $r_{LU} = E[\theta]$.

Step 4: Verification. Fix a $c \leq c_I$ and the most efficient (separating) equilibrium outcome with information acquisition at c. We argue that there exists no finite sequence of pruning of the original game with respect to the outcome that leads to a game in which the outcome is not an equilibrium outcome.

Suppose on the way to a contradiction that there is a sequence of pruning such that in the game Γ^k o fails to be an equilibrium outcome; but not in any Γ^j, for $j < k$. Then in all the games Γ^l, for $l < k$, in all equilibria that give rise to outcome o the receiver's response function is weakly below the indifference curve of θ_u passing through (s_H, θ_H) (which we call IC_U), and the indifference curve of the low type passing through the point $(0, \theta_L)$ (which we call IC_L). The first condition must hold due to the nature of equilibrium outcome being that the sender is indifferent between acquiring information and not acquiring information followed by the high type's action. For o not to be an equilibrium outcome in Γ^k, there should exist some $s \neq \{0, s_H\}$ such that $r(s)$ is strictly above the minimum of IC_U and IC_L in every candidate for an equilibrium. It should also be noted, that it cannot be the case that up to Γ_k all the strategies in which the sender plays s are pruned, otherwise s could not be a part of a profitable deviation.

Let the intersection of IC_U and IC_L be denoted (s_i, r_i). Because $c \leq c_I$, $r_i \geq E(\theta)$. Let also s_μ be such that IC_L crosses the ray $E(\theta)$ at s_μ. There are three cases to consider: $s \leq s_\mu$, $s \in (s_\mu, s_i]$, $s > s_i$.

Case 1: $s \leq s_\mu$. Since in Γ_k the receiver's response $r(s)$ is strictly above IC_L and the latter is in this case below $E[\theta]$, it has to be that the receiver cannot assign a belief to θ_L, or equivalently, all the strategies in which θ_L plays s have been pruned, except for possibly the strategy (I, s, s). For s to be NWBR for the low type in some earlier game, the receiver's beliefs should have been restricted to θ_L, i.e., all the strategies where the other two types play s should have been pruned away even earlier. But then after pruning away also the strategies where the low type plays s, s is not available for any type and, therefore, cannot represent a profitable deviation.

More precisely, since Γ^k is the first game in which $r(s)$ is above IC_L, it has to be the case that in Γ^{k-1} there exists some strategy in which θ_L chooses s. If θ_L is the only type in Γ^{k-1} that can choose s, $r(s)$ cannot be above θ_L in any subsequent games obtained after pruning. On the other hand, if θ_L is not the only type in Γ_{k-1} who can choose s, any $r(s) \in [\theta_L, E(\theta)]$ that is weakly below IC_L can be sustained as an off equilibrium receiver action, hence (I, s, s_H) is a best reply to *some* receiver strategy that induces the outcome

o, and is not pruned in Γ^k. Therefore, in Γ^k, there exists some strategy in which θ_L chooses s, so $r(s) = \theta_L$ is consistent with the requirement that the receiver best responds to some belief; contradicting the idea that $r(s)$ must be above the minimum of the two indifference curves.

Case 2: $s \in (s_\mu, s_i]$. First, given the definitions of s_μ and s_i, on the interval under the consideration IC_L is below IC_U, and moreover, IC_L is above $E[\theta]$. Given that $r(s)$ is above IC_L in Γ^k, it has to be the case that $r(s) = \theta_H$. Moreover in Γ^{k-1} there exists some strategy in which type θ_L or θ_U chooses s. Let Γ^l, $l < k$, be the last game in which the strategies in which the low type plays s are removed. In Γ^l the receiver can still assign positive probability to the low type, but if he can also assign a positive probability to θ_U or θ_H, then one can construct an equilibrium where s is a best response for a low type. Thus, if s is to be NWBR for θ_L, it must be the case that the receiver is in Γ^l assigning positive probability only to the low type, i.e., all the strategies where θ_U or θ_H play s have been previously pruned away. If that is the case, s cannot be a part of a profitable deviation in Γ_k.

Case 3: $s > s_i$. For $s > s_i$ the indifference curve IC_U is below IC_L. Moreover, for $s \geq s_H$, $r(s) \leq \theta_H$ implies that the receiver's strategy is always strictly below IC_U. Therefore the only potential deviation actions are in (s_i, s_H). The idea is: to arrive at a game where only the high type can play s one would need to prune away all the strategies in which θ_U plays s. But for s to be NWBR for θ_U one would need to prune away the strategies where the high type plays s beforehand. One can not have it both ways.

More formally, suppose there is some $s \in (s_i, s_H)$ such that $r(s)$ is above IC_U in Γ^k. For this to be the case, in Γ^k, only type θ_H should have a strategy in which he chooses s. Let Γ^l, $l < k$, be the game where the last strategy in which θ_U plays s is pruned, that is, such that in no equilibrium with the outcome o is s a best response for θ_U. Since in Γ^l the high type has at least one strategy in which he plays s, that would mean that no other type can play it; otherwise the requirement that the receiver has to best respond to some belief would not restrict him in $[E[\theta], \theta_H]$ and s would be a best response for θ_U. But if only θ_H can play s in Γ_l, it must be the case that $r(s) = \theta_H$, contradicting the supposition that Γ_k is the first game in which o is not an equilibrium outcome. □

Proof of Lemma 7: First we argue that any equilibrium with randomization over information acquisition decisions must be separating—each type (low, high, uninformed) undertakes a different amount of signaling—if it is to survive the refinement.

We start by showing that the high type and the low type cannot pool in an equilibrium with information acquisition. When $c > 0$, we showed this in Lemma 2. We will now show this when $c = 0$. Suppose on the way to a contradiction that the low and high types choose a signaling action s with positive probability. Then, the sender's equilibrium payoff is equal to the payoff she would get by not acquiring information and choosing s. Drawing the indifference curves of all the three types that pass through the action s and the equilibrium wage at s, we obtain that the high type's indifference curve is the flattest, hence it crosses the ray θ_H at some s_H that is further above from those at which the other two types' indifference curves intersect the ray θ_H. Hence, for some $\epsilon > 0$, applying the pruning procedure we obtain that we can erase all strategies in which the uninformed or the low type chooses an action above $s_H - \epsilon$. But then, in the new game obtained after the pruning, the wages for these actions have to be θ_H, which makes the initial outcome not a Nash equilibrium of the new game, a contradiction.

To show that the uninformed sender cannot pool with one of the two types, fix an equilibrium outcome with randomization over information acquisition in which the low type and the uninformed sender pool with positive probability on some action $\tilde{s}$, but not the high type. Then the receiver must respond with an $\tilde{r} < E[\theta]$. For $s' > \tilde{s}$ the uninformed sender's indifference curve through $(\tilde{s}, \tilde{r})$ is below the low type's indifference curve through the same point. For the prescribed outcome to be an equilibrium the receiver's response to actions above $\tilde{s}$ must, therefore, not be above the uninformed's indifference curve. Actions above $\tilde{s}$ are then NWBR for the low type. In the game obtained after pruning the strategies where the low type plays actions above $\tilde{s}$ the receiver should in every equilibrium respond to an $s' > \tilde{s}$ with an $r \geq E[\theta]$. But then the sender could profitably deviate to not acquiring information and choosing an s' just slightly above $\tilde{s}$. Equilibrium outcomes in which the uninformed sender and the high type pool are refined away similarly.

The above allows us to focus on equilibrium outcomes with randomization over information acquisition decisions and separation. Fix an outcome in which the uninformed agent strictly prefers his equilibrium action to the high type's. Due to single-crossing then so does the low type. The actions just below the high type's are then NWBR for the low type or the uninformed sender. After removing all the strategies where the uninformed sender or the low type play the mentioned actions one obtains a game where the actions could only be played by the high type, and therefore the receiver responds to them in any equilibrium with $r = \theta_H$. But then the sender has a profitable deviation. Likewise, if the low type were to strictly prefer his own action to the uninformed sender's the actions just below the uninformed's would be NWBR for the low type. In the game obtained after pruning the sender would have an incentive to deviate.

The only remaining equilibrium outcomes are the ones in which the low type chooses $s_L = 0$, the uninformed sender chooses an action that makes the low type indifferent and the high type an action that makes the uninformed sender indifferent. All the actions are, thus pinned down by the indifference curves. Since the sender must be indifferent between acquiring and not acquiring information, there is only one cost of information at which such an equilibrium outcome can exist. It should be noted that there is a continuum of equilibrium outcomes, as the probability with which the information is acquired is not pinned down.

Finally, we argue that the outcome with randomization over information acquisition, outcome o, survives the refinement. On the way to a contradiction, suppose there is a sequence of pruning with respect to the outcome o such that in the stages $i = 1, ..., k-1$, o is an equilibrium outcome in Γ^i, and in Γ^k, o is not an equilibrium outcome. Observe that, in the proof of Lemma 4, we only used the property that $c = b(s^*)$ to argue that the outcome survives the refinement. In proving the lemma, we used the claims that it cannot be that at stage k, $r(s)$ is above IC_U for $s > s^*$ in all equilibria, neither can it be the case that $r(s)$ is above IC_L in all equilibria. These claims continue to hold in the equilibrium outcome under consideration here. But if the claims are true, then in Γ^k, o is an equilibrium, which is a contradiction.

□

Proof of Lemma 8: Differentiating the value of information $b(s^*)$, given by (7), with respect to s^* gives

$$\frac{db}{ds^*} = -\lambda \left[g_s(s_H, \theta_H) \frac{ds_H}{ds^*} - g_s(s^*, \theta_H) \right].$$

On the other hand, differentiating the uninformed sender's indifference condition between $(s^*, E[\theta])$ and (s_H, θ_H) results in

$$\lambda \left[g_s(s_H, \theta_H) \frac{ds_H}{ds^*} - g_s(s^*, \theta_H) \right] + (1-\lambda) \left[g_s(s_H, \theta_L) \frac{ds_H}{ds^*} - g_s(s^*, \theta_L) \right] = 0,$$

or

$$\frac{ds_H}{ds^*} = \frac{\lambda g_s(s^*, \theta_H) + (1-\lambda) g_s(s^*, \theta_L)}{\lambda g_s(s_H, \theta_H) + (1-\lambda) g_s(s_H, \theta_L)}.$$

Combining the two equations yields

$$\begin{aligned}\frac{db}{ds^*} &= -\lambda \left[g_s(s_H, \theta_H) \frac{\lambda g_s(s^*, \theta_H) + (1-\lambda) g_s(s^*, \theta_L)}{\lambda g_s(s_H, \theta_H) + (1-\lambda) g_s(s_H, \theta_L)} - g_s(s^*, \theta_H) \right] \\ &= \lambda(1-\lambda) \frac{g_s(s_H, \theta_L) g_s(s^*, \theta_H) - g_s(s^*, \theta_L) g_s(s_H, \theta_H)}{\lambda g_s(s_H, \theta_H) + (1-\lambda) g_s(s_H, \theta_L)}.\end{aligned}$$

The sign of $\frac{db}{ds^*}$ is, therefore, determined by the sign of $g_s(s_H, \theta_L) g_s(s^*, \theta_H) - g_s(s^*, \theta_L) g_s(s_H, \theta_H)$, which, in turn, coincides with the sign of

$$-\frac{d^2 \log(g_s(s, \theta))}{ds d\theta}.$$

□

Proof of Lemma 9: The cross-partial derivative $\frac{d^2 \log(g_s(s,\theta))}{dsd\theta} = 0$ implies that the first derivative $\frac{d \log(g_s(s,\theta))}{d\theta}$ is constant in s. Therefore, there exist real functions $\tilde{f}$ and $\tilde{g}$ such that $log(g_s(s, \theta)) = \tilde{f}(s) + \tilde{g}(\theta)$, or $g_s(s, \theta) = e^{\tilde{f}(s)+\tilde{g}(\theta)}$. Define $f(s) = e^{\tilde{f}(s)}$ and $g(s) = e^{\tilde{f}(s)}$, which delivers the result.

The other direction is established via a simple computation. □

Proof of Lemma 10: Differentiating equation (11) with respect to s yields:

$$h'(s) \left(\lambda f(\theta_H) + (1-\lambda) f(\theta_L)\right) = h'(s_u) \left(\lambda f(\theta_H) + (1-\lambda) f(\theta_L)\right) \frac{ds_u}{ds},$$

which together with $\lambda f(\theta_H) + (1-\lambda) f(\theta_L) \neq 0$ delivers

$$h'(s) = h'(s_u) \frac{ds_u}{ds}. \tag{12}$$

On the other hand, differentiating $\tilde{b}(s)$ gives:

$$\tilde{b}'(s) = -p_h h'(s_u)\frac{ds_u}{ds}\left(\pi_h f(\theta_H) + (1-\pi_h)f(\theta_L)\right) - p_l h'(s)\left(\pi_l f(\theta_H) + (1-\pi_l)f(\theta_L)\right) \tag{13}$$

$$+ h'(s)\left(\lambda f(\theta_H) + (1-\lambda)f(\theta_L)\right). \tag{14}$$

Combining (13) and (12) culminates in

$$\tilde{b}'(s) = -p_h h'(s)\left(\pi_h f(\theta_H) + (1-\pi_h)f(\theta_L)\right) - p_l h'(s)\left(\pi_l f(\theta_H) + (1-\pi_l)f(\theta_L)\right)$$
$$+ h'(s)\left(\lambda f(\theta_H) + (1-\lambda)f(\theta_L)\right).$$

The inequality $h'(s) < 0$, and an implication of Bayes' rule, $p_h\pi_h + p_l\pi_l = \lambda$, yield $b'(s) = 0$. □

Proof of Lemma 12: Fix the outcome where the sender acquires information and each type chooses the action corresponding to the Riley outcome (the most efficient separating equilibrium). In particular, type θ_1 chooses $s = 0$ and each type θ_i is indifferent between his own action and the action taken by θ_{i+1}. By the single-crossing property, each type θ_j, then strictly prefers action s_j to any action s_k with $k < j$ and also, each type θ_j strictly prefers s_j to any s_k with $k > j+1$.

To support the Riley outcome, assume that the receiver replies to any action $s \in (s_i, s_{i+1})$ with $r = \theta_i$, with convention $s_{n+1} = \infty$. The sender would, therefore, after deviating to not acquiring information choose one of the actions s_i, denote it s_l. Since every type θ_i, $i \notin \{l, l+1\}$ strictly prefers action s_i to s_l, and there are at least three types, there is a type who prefers his own action to the one the uninformed sender would choose. Therefore, if c is small enough, the sender is better off acquiring information. □

Proof of Lemma 13: Fix the outcome with information acquisition and the Riley outcome. Each type is indifferent between his own action and the action of the upward-adjacent type. The indifference curves connecting all the type's option present an upper bound the receiver's response in any equilibrium with the outcome.

To argue that the outcome cannot be refined away we can argue that no s that is not played in equilibrium can be a part of a profitable deviation after a finite sequence of deletion of NWBR strategies. Actions above s_n can clearly never be a a part of profitable deviation. Let's focus on some $s \in (s_i, s_{i+1})$. The receiver responds to s_i with θ_i and s_{i+1} with θ_{i+1}. For action s to become a profitable deviation it would have to be the case that after pruning it can be played only by types θ_{i+1} and above or the uninformed sender if $E[\theta] > \theta_i$.

Suppose that s indeed represented a profitable deviation after some number of rounds of deletions. Then s must have been NWBR for type θ_i at some earlier stage. However, to have been a NWBR for θ_i the action should have been eliminated for all the types above θ_i at an even earlier stage. But then s cannot be a profitable deviation. $\square$