Should Consumers be Given Control of their Own Private Information?[*]

Qihong Liu[†] Konstantinos Serfes[‡]

February 12, 2007

****Preliminary Draft****

Abstract

Consumers, through their actions in the market place, have generated a vast amount of information about their preferences, habits and characteristics. Nevertheless, this information is exploited by other parties and consumers do not have much say in the process. In this paper, we are mainly interested in examining the consequences of a regulation that would transfer to consumers all the property rights to their own private information. Under such a regulation, the act of information selling induces a signaling game that may be detrimental to the consumer. This suggests that flows of information may be restrained which will impact efficiency (and perhaps consumer surplus) negatively. Nevertheless, we show that no matter who has the property rights to consumer private information, the outcome will be efficient. In other words, the Coase theorem holds in an asymmetric information environment. Moreover, consumers become better off under the regulation.

JEL Classification Codes: C72; D82.

Keywords: Information sharing; Asymmetric information; Coase theorem; Signaling game.

---

[*] A previous version of the paper was titled "Trading private information in a principal agent adverse selection model." We would like to thank Minghua Chen, Vibhas Madan, Theofanis Tsoulouhas and the seminar participants at the Spring 2006 Midwest Theory meeting at Michigan State University and the 2006 Far Eastern Meeting of the Econometric Society (FEMES) at Tsinghua University, China for helpful comments, discussions and suggestions. The usual caveat applies.

[†] Department of Economics, University of Oklahoma, 729 Elm Ave, Norman, OK 73019. Tel: (405) 325-5846. Fax: (405) 325-5842. E-mail: qliu@ou.edu.

[‡] Department of Economics and International Business, Bennett S. LeBow College of Business, Drexel University, Matheson Hall, 32nd and Market Streets, Philadelphia PA 19104. E-mail: ks346@drexel.edu. Phone: (215) 895-6816. Fax: (215) 895-6975.

“Wanna See My Personal Data? Pay Up.” *Business Week*, November 21, 2002

1 Introduction

Information about consumer preferences, habits and characteristics is very valuable to firms. The growth of the Internet as a medium of communication and commerce, combined with the development of sophisticated software tools have contributed to the collection and analysis of a vast amount of data about consumers. Firms who possess such information can target individual consumers (or certain groups of consumers) more effectively. This practice is facilitated by many information intensive marketing approaches such as database marketing, target marketing, micro-marketing and one-to-one marketing [e.g., Shaffer and Zhang (2002)]. Consumers, however, do not have the property rights to their own information. Rather, various agencies (e.g., credit bureaus) can sell these customer data to firms without consumers having much say in the whole process. There exists a growing marketing literature which argues that the solution to privacy invasion is to create institutions that allow consumers to build and claim the value of their marketplace identities [e.g. Deighton (2002)]. According to this literature, personal information is an asset and consumers should get compensated when companies use their personal information. Some small steps have been taken in this direction recently. Based on the Online Privacy Protection Act consumers should give their consent to firms before they share customer information with a third party. Nevertheless, firms make every effort to safeguard valuable consumer information and their option to sell it to third parties.[1]

This paper makes an attempt to address some of the important issues that were raised in the above paragraph. We ask the following question: what would happen if regulators gave consumers all the property rights to their own private information? We are interested in two specific aspects of the problem: i) economic efficiency and ii) benefits to consumers. As far as we know this is the first paper that studies this problem formally.[2]

We cast the problem in a principal agent model, where the principal can be thought of as being a firm and the agent as being a consumer (or a group of consumers) of an unknown type.[3] Without any additional information about the type of the agent, the principal will make a menu of offers available to the agent who in turn will self-select (e.g., second-degree price discrimination).

---

[1] See, for example, “A very public battle over privacy,” *Business Week*, May 23, 2002. According to the article “...most companies bury the opt-out notices within masses of legal jargon at the bottom of monthly mailings.”

[2] Our main objective is to investigate the pure economic consequences of consumer information sharing. We completely abstract from any privacy issues, which are nevertheless very important in practice.

[3] For instance, the firm can be a financial institution (bank) who has incomplete information about the creditworthiness of its clientele. The presence of asymmetric information leads to an inefficient outcome where the size of the loan is distorted. Better information about the type of the borrower will lead to a more efficient outcome.

The outcome is not efficient because the principal optimally distorts his offer to the bad type (low demand) agent. This may mean, for example, that product quality is below its first best level. At the same time the good type (high demand) receives a surplus. We extend this standard model by assuming that information about the type of the agent, in the form of a database, is available. This database contains information about the agent's characteristics (e.g., past purchasing history, credit card statements, etc.) and it can be used to provide an informative (but nonverifiable) signal about the agent's type. The quality (accuracy) of the signal is in direct relationship with how much of the agent's private data (PD) is sold. If the entire database is sold, then we assume that the signal that will be generated is perfect, otherwise it is noisy. We analyze two different cases depending on who has the property rights to the agent's private data (PD). In the first case a third player, whom we call an intermediary (e.g., credit bureau) has all the property rights and in the second case it is the agent himself who has all the property rights. We call the first case an *unregulated* regime and the second a *regulated* regime. Information is valuable to the principal because it improves his ability to make "better" offers to the agent and as a consequence more information leads to a more efficient outcome.

In the unregulated regime, the intermediary decides about how much information to sell to the principal and at what price. We assume that it is possible to sell any fraction of the PD. After the exchange of the PD the principal makes an offer to the agent (contract stage).

In the regulated regime, a regulator transfers all the property rights of the agent's PD from the third party to the agent. The intermediary now disappears.[4] The agent decides about how much of his data to put up for sale and at what price.[5] Following the agent's action, the principal makes his offer to the agent. The agent's action, however, about how much of his PD to offer for sale and at what price may signal the agent's true type. If the principal perceives the agent as being of a specific type with sufficiently high probability, then the agent's PD is of little value to the principal. The principal then may choose to rely exclusively on the information transmitted by the agent's offer and to decline the purchase of the agent's PD. The agent in that case has essentially divulged information about his type without getting compensated in return. Thus, signaling in our context may be detrimental to the agent, regardless of his type. (This possibility does not arise in the unregulated regime, due to the impersonal nature of the exchange). As a consequence, the

---

[4] The question that immediately arises is: who collects the information in the absence of the intermediary? Our implicit assumption is that the information has already been collected and the main question is how the transfer of property rights will affect efficiency and agent (consumer) welfare. Alternatively, we can imagine a technology (e.g., software) which collects information on behalf of each consumer. If consumers can benefit directly from their own information, then facilitating technologies will most likely be developed.

[5] According to the article "Wanna See My Personal Data? Pay Up," in Business Week, November 21 2002, a consumer auctioned-off on e-Bay 800 pages of his personal information and he walked off with $240. The designer's auction was the trial balloon for a new conceptual framework for privacy called Loome, which weaves together businesses' desire for customer data with an individual's wish to be compensated for sharing personal details of one's life.

agent may choose to hold on to his PD, although such an information exchange may have improved efficiency.

We are mainly interested in answering the following two questions:

1. What fraction of the PD will be traded in each one of the two regimes (regulated vs. unregulated)?
2. How will the surplus be shared among the involved parties in each one of the two regimes?

The first question is important from an efficiency point of view. For example, the signaling concern that arises in the regulated regime and was discussed above, may prevent the agent from selling part or all of his PD. This will lead to a less efficient outcome, where flows of information are restrained. Put it differently: does the Coase theorem hold? The answer is not obvious because some of the qualifications of the Coase theorem are not satisfied in our model. Specifically, there is no full information among all parties about the costs and benefits (due to the asymmetry of information).[6] The second question deals with agent (consumer) surplus which is one of the most important elements to regulatory authorities.

In the unregulated regime the intermediary has incentives to sell all the PD to the principal. The price will equal to the principal's maximum willingness to pay for the information (we assume that whoever makes an offer has all the bargaining power). The outcome will be efficient, since the principal upon acquiring the PD will learn the agent's type with certainty (i.e., signal is perfect). The intuition behind why all of the agent's PD will be sold is straightforward, because information is valuable to the principal (i.e., the more information the better) and the intermediary does not participate in the next stage of the game. Also, it is not difficult to see that the agent's surplus will be zero (not matter what his type is).[7]

Let's now turn our focus on the regulated regime analysis which is considerably more challenging than the analysis of the unregulated regime.[8] We formulate a three-stage stage. Nature chooses the agent's type (observable only by the agent) which can be either efficient (or high demand) "good" or inefficient (or low demand) "bad." In stage 1 (signaling stage) the agent makes an offer

---

[6] See, for example, Ishiguro (2003) where the focus is on the validity of Coase theorem under asymmetric information. Under some conditions, but not always, the Coase theorem holds.

[7] In reality there are some restrictions about how much (and what kind) of data firms can sell (or share). Our model can easily generate this outcome. If a regulator (in the unregulated regime) cares about efficiency and consumer surplus, then it can be shown that some restrictions on the flow of information will be imposed. This is because a regulator that cares about consumer surplus will not find optimal an outcome where consumers receive zero surplus. Some substitution between efficiency and agent surplus will be allowed. Alternatively, agents may receive positive surplus if there is competition among principals (common agency).

[8] To avoid possible confusion, we would like to clarify the following concepts. We use the word "signaling" to refer to the information that may be transmitted by the agent's action in stage 1 and the word "signal" to refer to the information that the principal receives in stage 3 if he purchases the agent's PD.

to the principal. The offer entails a fraction (or all) of the agent's PD and a price for his PD. The principal accepts or rejects the agent's offer in stage 2. Given the information that has been transmitted, the principal offers a menu of quantities and transfers to the agent in stage 3 (screening stage). If the agent signals his type through his choice of the quality of his PD and the price in stage 1, then the principal will not purchase the agent's PD. Rather, he will rely entirely on the information contained in the agent's action. It is then natural to search for a pooling equilibrium, where the agent does not reveal his type without receiving a monetary reward in return. We show that a unique pooling perfect Bayesian equilibrium exists where the agent sells *all* his PD to the principal.[9] The principal then draws a perfect signal and offers to the agent the efficient (first-best) contract. Because the agent gets compensated for selling his PD, he ends up with positive surplus regardless of his type. The "bad" type is always better off when he can sell his PD than when such an option is not available (and no information is sold whatsoever). The "good" type becomes better off (relative to the no information selling case) if and only if the ex-ante probability of the "good" type is high enough. So, in a market with very few efficient types, the efficient type is worse off when the option of selling his information becomes available relative to when such an option is not present. However, in equilibrium, they must pool themselves with the low types and offer their private data for sale.

The regulated equilibrium, as we discussed above, is characterized as follows: i) both agent types make the same offer (pooling equilibrium) and ii) the offer entails all of the agent's PD. As we have already alluded to a separating equilibrium would be detrimental to the agent and that is why both types, in equilibrium, make the same offers. The second result is less obvious. We show (in the appendix) that any pooling outcome where agents sell only a fraction of the PD is not an equilibrium. First, note that the inefficient type always receives zero surplus (excluding the price for his PD) from the contract that is offered in stage 3. This is a standard result in screening models. Hence, the inefficient type always has incentives to deviate by offering more of his PD for sale at a higher price. What may prevent him from doing so is the possibility that his action may change the principal's ex-ante beliefs drastically. Nevertheless, this is *not* the case. Let's briefly explain why. (A more detailed intuition is offered after theorem 1). Under some set of beliefs held by the principal, about the likelihood of an efficient type following a deviation, the efficient type has an incentive to offer more of his PD for sale but not the inefficient type, while under some other set of beliefs it is the inefficient type who wishes to sell more of his PD but not the efficient type. These two sets of beliefs cannot be nested and according to the Banks and Sobel universal divinity refinement both types are *likely* to deviate. This implies that the principal does not update his ex-ante beliefs following a deviation (in a specific region) and the inefficient type would always want to deviate from an offer that entails only a fraction of the PD by offering more of his PD for sale

---

[9] Moreover, the PBE has beliefs that satisfy the Banks and Sobel (1987) *universal divinity refinement*.

at a higher price.

If we compare the result from the regulated regime with the one from the unregulated regime, we will see that both regimes lead to an efficient outcome. This says that who holds the property rights does not matter for efficiency (Coase's theorem), even under asymmetric information. Moreover, the agent is better off in the regulatory regime relative to the unregulated regime, since under the regulation his surplus is strictly positive, regardless of his type, due to the payment he receives for his PD.

Our framework can also be applied in other contexts. For example, when a principal, such as the Department of Defense, deals with contractors (agents) to develop a new weapon, the principal does not know with certainty the agent's type, i.e., whether the agent is efficient or not. Nevertheless, it may very well be the case that the contractor has data (e.g., accounting or engineering data from past projects) that can generate an informative signal about his type. These data are the property of the agent and can be sold to the principal prior to the contracting stage.

1.1 Literature review

Our paper mainly contributes to the following two distinct strands of literature: i) the literature on information sharing and ii) the literature on principal agent models with endogenous information structures.

Early literature on information sharing among rival firms had mainly focused on two types of information exchanges: i) firms share - directly or indirectly - their private signals about demand conditions, or ii) firms exchange cost data.[10] Information sharing of consumer data has become an important phenomenon only recently due to the development of technologies that can track consumer behavior, and store and analyze vast amounts of information. Liu and Serfes (2006) study the issue of consumer information sharing among rival firms. The main assumption that is made in that paper is that consumers are passive and consequently they are not compensated directly when information about them is traded. Instead, the present paper assumes that it is possible for consumers to be key and active players in the information exchange process.

There exists a large body of principal-agent literature which investigates how the asymmetry of information can be mitigated. This may lead to information structures that are determined endogenously. In that literature, information about the private type is not readily available, but it can be gathered (at some cost) by either the principal or the agent. The main issue is how the contract should be structured to provide adequate incentives for information gathering. In contrast, our paper assumes that information about the type of the agent has already been collected and it belongs to the agent, and the main issue is how this information can be transmitted to the principal.

---

[10] See, for example, Gal-Or (1985), Shapiro (1986), Vives (1990), Villas-Boas (1994).

Riordan and Sappington (1988) consider a model where the agent and the court observe ex-post a verifiable signal that is correlated with the type of the agent. The first-best can be implemented. Agents receive zero rents. In Boyer and Laffont (2003) the principal receives an informative nonverifiable signal, which improves the contracting ability of the principal. Baron and Besanko (1984) and Khalil (1997) investigate the role of audit mechanisms with and without commitment. Myerson (1983) and Maskin and Tirole (1990 & 1992) formulate models where the principal learns his type before he makes an offer to the agent. With common values there is some allocative inefficiency, which disappears when we move to the private values paradigm.

In Laffont & Martimort (2002, pp.395-398) the agent chooses whether to learn his type after he has accepted the contract. In Crémer and Khalil (1992a) and Crémer, Khalil and Rochet (1998a) the agent may gather information about his type before the contract is offered. Information is gathered for rent-seeking purposes since the type will be learned costlessly ex-post.

Information gathering may also be done for productive purposes, if the type will not be learned costlessly after the contract is signed, e.g., Crémer and Khalil (1992b) and Crémer, Khalil and Rochet (1998b). Mezzetti and Tsoulouhas (2000) examine the agent's incentives to gather information after the contract is offered (but before it is signed), in a moral hazard model with a privately informed principal. The principal may also gather information after he signs the contract, e.g., Finkle (2005).

In the signaling game of Maskin and Tirole (1990) it is the principal who has private information and offers the contract. In our model, the signaling comes from the agent's side and the contract is offered by the principal. Our model contributes to the literature of endogenous information structures. When the principal offers the contract (in stage 3) the information he possesses has emerged endogenously based on the agent's decision in stage 1 about how much of his private information to sell. First, our game can be viewed as a one where the agent gathers information before the contract is offered, assuming that the agent already knows his type. In other words, the agent incurs a cost to organize his PD in a way that can be easily transferred to the principal. In the paper, we assume that this cost is zero, but our results should go through as long as the cost is sufficiently small.[11] Second, after information is gathered (organized), the agent sells his PD to the principal before the contract is offered, leading to a signaling game. Both of these features of our model are novel and as we argued above there exist real-world situations where this model can be insightful.

The agent's PD in our set-up generates a nonverifiable signal, as in Boyer and Laffont (2003). The difference is that in Boyer and Laffont the signal is exogenous, while in our model the quality of the signal the principal gets to observe in stage 3 is endogenously determined.

---

[11] Of course, the results will probably change if the agent is initially uninformed and the choice is whether to acquire information about his type or not, as in Crémer, Khalil and Rochet (1998a).

Our paper also touches on the literature of strategic information transmission in principal-agent models. For example, Kahn and Tsoulouhas (1999) study that issue in a repeated principal-agent relationship where the agent produces information (output) that is useful to the principal. They show that full disclosure of information, on part of the agent, occurs provided that the parties are patient enough.

The rest of the paper is organized as follows. We lay out the model in section 2. Section 3, contains the analysis of the regulated regime and our main result (theorem 1). We conclude in section 4. The appendix contains most of the proofs.

2 The description of the model

A principal wants to delegate to an agent the production of q units of a good.[12] In this section, we do not make any assumption about who has the property rights to the agent's private data, PD. It can be either the agent (regulated regime) or an intermediary (unregulated regime). The value to the principal of these q units is given by $S(q)$ where $S' > 0$, $S'' < 0$ and $S(0) = 0$. The production cost of the agent is unobservable to the principal, but it is common knowledge that the marginal cost belongs to the set $\Theta = \{\theta_\ell, \theta_h\}$, with $\theta_h \geq \theta_\ell$. The agent can be efficient (θ_ℓ) or inefficient (θ_h) with respective *ex-ante* probabilities π and $(1-\pi)$. Types are indexed by i, i.e., $i = h, \ell$. We denote by $\Delta\theta = \theta_h - \theta_\ell \geq 0$ the difference in the marginal costs. In other words, the agent's cost function is,

$$C(q, \theta_\ell) = \theta_\ell q, \text{ with ex-ante probability } \pi$$

or

$$C(q, \theta_h) = \theta_h q, \text{ with ex-ante probability } (1-\pi).$$

The economic variables of the problem are the quantity produced q and the transfer t received by the agent. Let $\mathcal{A}$ be the set of feasible allocations. Formally, we have,

$$\mathcal{A} = \{(q,t) : q \in \mathbb{R}_+, t \in \mathbb{R}\}.$$

These variables are observable by a third party (e.g. a court), and hence they can be viewed as a *contract*. The utility of type i agent is denoted by $U_i = t_i - \theta_i q_i$ and the utility of the principal by $V = S(q) - t_i$.

---

[12] We chose this specific set-up of the principal agent model because it is the most common in the literature [see Laffont and Martimort (2002)]. This model is useful in many settings such as, regulation, non-linear pricing by a monopolist, quality and price discrimination, financial and labor contracts. All the results would remain qualitatively the same if we formulated the problem in terms of a consumer of an unknown willingness to pay. In this case, $S(q;\theta)$ is type θ consumer's willingness to pay for quantity or quality q. The firm's cost function is cq and the surplus function is $S(q;\theta) - cq$, where c is a known marginal cost.

We assume that whoever is selling the PD (i.e., the intermediary of the agent) he can sell any subset of the PD to the principal.[13] The information that is contained in the agent's PD is in the form of an imperfect (nonverifiable) signal $s \in \{s_\ell, s_h\}$ correlated with the true type of the agent, that is,

$$\Pr(s_h|h) = \Pr(s_\ell|\ell) = r \geq \frac{1}{2}$$
$$\Pr(s_\ell|h) = \Pr(s_h|\ell) = 1 - r.$$

As r increases, the informativeness of the signal increases, in which case we say that the agent's PD is of a higher quality. So, if all of the agent's PD is sold to the principal, then we assume that $r = 1$ (perfectly informative signal). But also a fraction of the PD can be sold, in which case $r < 1$.

The probability that the agent is efficient if the signal is s_ℓ can be expressed via Bayes' rule as follows,

$$\sigma_\ell(r|\pi) = \Pr(\ell|s_\ell) = \frac{r\pi}{r\pi + (1-r)(1-\pi)}. \tag{1}$$

The probability that the agent is inefficient if the signal is s_h can be expressed via Bayes' rule as follows,

$$\sigma_h(r|\pi) = \Pr(h|s_h) = \frac{r(1-\pi)}{r(1-\pi) + (1-r)\pi}. \tag{2}$$

Hence, $\sigma_i(r|\pi)$ is the probability that the agent is of type i given that the ex-ante probability is π, $i = \ell, h$.

The game we will study can be described as follows (see also figure 1 where the extensive form game tree is depicted and figure 2 where we present the sequence of events, assuming that we are in the regulated regime).

- Stage 0. Nature chooses the type of the agent. With *ex-ante* probability π the agent is efficient (low type, ℓ) and with probability $1 - \pi$ is inefficient (high type, h). The type of the agent is private information.
- Stage 1. The agent (or the intermediary in the unregulated regime) offers to the principal his PD of quality r at a price $F \geq 0$. The agent (or the intermediary) has all the bargaining power at this stage.
- Stage 2. The principal accepts (A) or rejects (R) the agent's (r, F) offer.
- Stage 3. The principal offers a contract $\mathcal{C} = \{(t_h, q_h), (t_\ell, q_\ell)\}$ to the agent. The principal has all the bargaining power at this stage.

---

[13] We assume that the information seller cannot manipulate the information contained in the PD.

We denote by $\psi : [0,1] \times \left[\frac{1}{2}, 1\right] \times \mathbb{R}_+ \rightarrow [0,1]$ the *interim* belief held by the principal in stage 2 that the agent is efficient (π now is replaced by ψ, although it is possible that $\pi = \psi$). This belief depends on the ex-ante probability π and on the information transmitted by the agent's offer (r, F). The dependence of the interim belief on the agent's offer will become more clear later. Further, we define the ex-post belief $\mu : \left[\frac{1}{2}, 1\right] \times \mathbb{R}_+ \times \{s_h, s_\ell\} \rightarrow [0,1]$ held by the principal in stage 3 that the agent is efficient as follows,

$$\mu = \begin{cases} \sigma_\ell \left(r|\psi\right), & \text{if } s = s_\ell \\ \sigma_h \left(r|\psi\right), & \text{if } s = s_h \\ \psi, & \text{if } s = \varnothing. \end{cases} \tag{3}$$

These probabilities are obtained from (1) and (2) after π is replaced by ψ. The ex-post belief μ depends on the interim belief ψ and on the realized signal s, if the principal has accepted in stage 2; otherwise (i.e., if the principal rejects the agent's offer, $s = \varnothing$), the ex-post belief coincides with the interim.

We search for a perfect Bayesian equilibrium (PBE), in pure strategies. A PBE is a strategy profile and a system of beliefs such that the strategies are sequentially rational given the beliefs and the beliefs are updated via Bayes' rule, whenever possible (e.g., Fudenberg and Tirole (1991, pp.325-326)).[14]

The main questions we would like to address are: i) will players succeed in trading private information? and ii) if the answer to the previous question is affirmative, how much information will be traded?[15]

3 Analysis of the regulated regime

We now provide a road map about how we solve the game (consult also figures 1 and 2), assuming that the agent has all the property rights to his PD.[16] We begin with stage 3. Given the ex-post belief μ of the principal, this is a standard principal-agent adverse selection game. The solution gives us the equilibrium (terminal) payoffs as a function of the ex-post belief. Given these payoffs, stages 1 and 2 constitute a standard signaling game. The agent makes an offer, (r, F), in stage 1 and the principal accepts or rejects in stage 2. In section 3.2, we divide the (r, F) space into an acceptance and a rejection region. These regions are not fixed, but they depend on the interim belief ψ held by the principal when he decides whether to accept or to reject the offer. In section

---

[14] When possible means at all information sets reached with positive probability according to the given joint strategy. Beliefs are unrestricted by the concept of PBE on off-equilibrium paths. We restrict the off-equilibrium beliefs by insisting that they satisfy the Banks and Sobel (1987) universal divinity refinement.

[15] Glycopantis, Muir and Yannelis (2003) address similar questions (i.e., whether Pareto trades will take place among players) in a differential information economy. In particular, they are interested in how contracts are implemented as perfect Bayesian equilibria in extensive form games.

[16] The analysis of the unregulated regime is straightforward and it is not pursued further. As we argued in the introduction a third party will have incentives to sell all the PD. The outcome will be efficient.

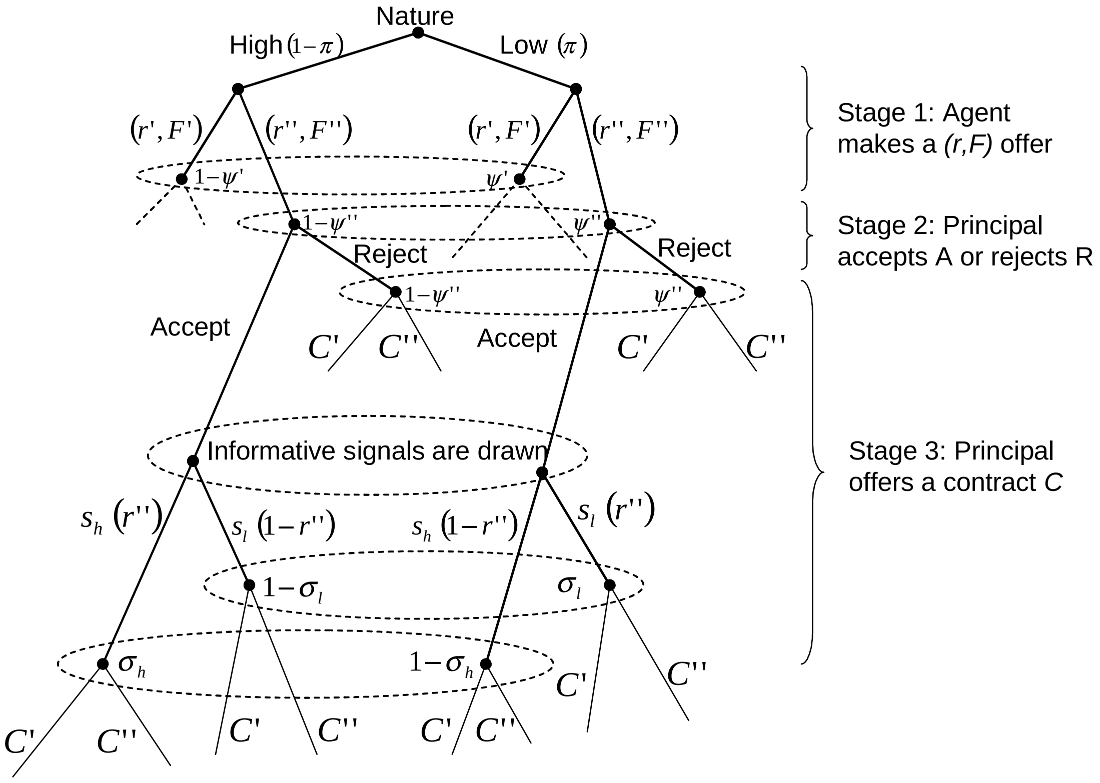

Figure 1: The extensive form game for the regulated regime. The figure is complete except that it shows only two offers (r', F') and (r'', F'') and two contracts C' and C'' when there are infinitely many choices available.

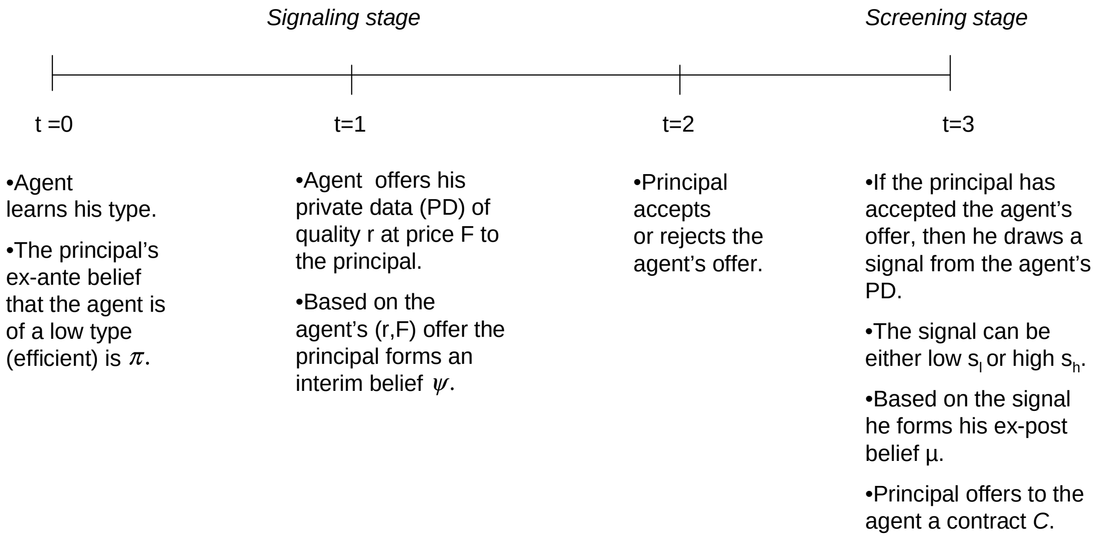

- Agent learns his type.
- The principal's ex-ante belief that the agent is of a low type (efficient) is π.

- Agent offers his private data (PD) of quality r at price F to the principal.
- Based on the agent's (r,F) offer the principal forms an interim belief ψ.

- Principal accepts or rejects the agent's offer.

- If the principal has accepted the agent's offer, then he draws a signal from the agent's PD.
- The signal can be either low s_l or high s_h.
- Based on the signal he forms his ex-post belief µ.
- Principal offers to the agent a contract *C*.

Figure 2: Sequence of events for the regulated regime

3.1 we determine the equilibrium of this game. First, we show that a separating equilibrium (where each type makes a different offer) does not exist. This is due to the fact that in a separating equilibrium the offers made by the agent reveal the agent's type. In this case the principal learns the agent's type with probability one and as a result the informative signal that is contained in the agent's private data does not provide any additional information. The principal's best response is to reject the agent's offer and to design the first-best contract which extracts the agent's surplus, regardless of his type. Therefore, the agent ends up with a zero payoff and consequently does not want to signal his type. Hence, we search for a pooling equilibrium. Since both types make the same offer (in a pooling equilibrium), the ex-ante belief π is not updated in equilibrium (so, along the equilibrium path, the ex-ante belief coincides with the interim belief, $\pi = \psi$). The ex-post belief μ does not need to coincide with the interim belief ψ, because the information the agent sells, if it gets accepted, generates an informative signal which updates the interim belief (see (3)). In equilibrium, the agent's offer does get accepted. Moreover, off-the-equilibrium path beliefs are important since they may affect the equilibrium and also help us to reduce the number of equilibria. We allow for off-equilibrium beliefs that satisfy the universal divinity refinement.[17] Given these beliefs we show that there exists a unique pooling equilibrium where the agent succeeds in selling all his information to the principal.

3.1 Stage 3: Principal offers a contract $\mathcal{C}$ (screening stage)

The ex-post belief held by the principal that the agent is efficient is μ, see (3). Under complete information the contract offered by the principal entails the efficient (first-best) quantities, q_h^* and q_ℓ^*, that satisfy $S'(q_h^*) = \theta_h$ and $S'(q_\ell^*) = \theta_\ell$. Under incomplete information the principal maximizes his expected profits subject to the usual incentive compatibility (IC) and individual rationality (IR) constraints, that is,

$$\max_{(q,t)\in\mathcal{A}} \mu\left[S(q_\ell) - \theta_\ell q_\ell\right] + (1-\mu)\left[S(q_h) - \theta_h q_h\right] - \left[\mu U_\ell + (1-\mu) U_h\right]$$

subject to: (i) $U_\ell \geq U_h + \Delta\theta q_h$, (ii) $U_h \geq U_\ell - \Delta\theta q_\ell$, (iii) $U_h \geq 0$ and (iv) $U_\ell \geq 0$,

where (i) and (ii) are the (IC) constraints and (iii) and (iv) are the (IR) constraints. This is a standard principal-agent adverse selection problem [e.g. Laffont and Martimort (2002), chapter 2]. It is well-known that the optimal contract is characterized as follows:

- No output distortion of the efficient type with respect to the first best, $q_\ell^{SB} = q_\ell^*$. A downward output distortion (second-best, SB) for the inefficient type, $q_h^{SB} < q_h^*$ with,

$$S'\left(q_h^{SB}\right) = \theta_h + \frac{\mu}{1-\mu}\Delta\theta. \tag{4}$$

---

[17] A more extensive discussion on the off-the-equilibrium beliefs is offered at the end of section 3.

- Only the efficient type gets a positive information rent given by,

$$U_\ell = \Delta\theta q_h^{SB}(\mu). \tag{5}$$

- The second best transfers are respectively given by,

$$t_\ell^{SB} = \theta_\ell q_\ell^* + \Delta\theta q_h^{SB}(\mu) \text{ and } t_h^{SB} = \theta_h q_h^{SB}(\mu). \tag{6}$$

So, the contract is $\mathcal{C} = \left\{\left(q_h^{SB}(\mu), \theta_h q_h^{SB}(\mu)\right), \left(q_\ell^*, \theta_\ell q_\ell^* + \Delta\theta q_h^{SB}(\mu)\right)\right\}$. The principal's expected utility is given by,

$$EV(\mu, \theta_h, \theta_\ell) = \mu\left[S(q_\ell^*) - \theta_\ell q_\ell^*\right] + (1-\mu)\left[S\left(q_h^{SB}(\mu)\right) - \theta_h q_h^{SB}(\mu)\right] - \mu\Delta\theta q_h^{SB}(\mu). \tag{7}$$

Further, we assume that $S'(0) = \infty$ and $\lim_{q\to 0} S'(q)\, q = 0$, which imply that the principal never finds it optimal to shut-down the inefficient agent.[18] To ensure that production is always finite we assume that $\lim_{q\to+\infty} S'(q) = 0$. By differentiating (4) with respect to q and μ we can derive the effect of an increase in the probability of the efficient type on the output of the inefficient type,

$$S'' dq_h - \frac{\Delta\theta d\mu}{(1-\mu)^2} = 0 \Rightarrow \frac{dq_h^{SB}}{d\mu} = \frac{\Delta\theta}{S''\left(q_h^{SB}(\mu)\right)(1-\mu)^2} < 0. \tag{8}$$

The downward output distortion of the inefficient type becomes more pronounced when the probability of an efficient type increases. This also lowers the efficient type's rent.

3.2 Stage 2: The principal accepts or rejects the agent's (r, F) offer

We wish to characterize the set of offers made by the agent in stage 1 that are acceptable to the principal, conditional on his beliefs. Suppose that the agent offers (r, F). The principal's interim belief that the agent is efficient is given by ψ, which may or may not be equal to the ex-ante probability π (so, in general, in stage 2 π is replaced by ψ). The interim belief may not coincide with the ex-ante belief, because the agent's (r, F) offer in stage 1 may transmit additional information about the type of the agent.

If the principal accepts the agent's offer, then his ex-post beliefs in stage 3 are given by (3), where s is either s_ℓ or s_h. If the signal is low, the principal's expected utility is given by (7) with the difference that the ex-post probability of an efficient type μ is replaced by $\sigma_\ell(r|\psi)$, i.e., $EV(\sigma_\ell)$, where its dependence on θ's, r and ψ is suppressed. If the signal is high, the principal's expected utility is given by (7) with the difference that the ex-post probability of an efficient type μ is replaced by $1 - \sigma_h(r|\psi)$, i.e., $EV(1 - \sigma_h)$. The probability that the principal will receive a

---

[18] For example, the function $S(q) = \sqrt{q}$ satisfies these conditions.

low signal conditional on the quality of the PD being r is given by $\Pr(s_\ell|r) = \psi r + (1-\psi)(1-r)$. Therefore, the principal's expected utility in stage 2, if he accepts and conditional on r, can be expressed as follows,

$$EV(r, F|\theta_h, \theta_\ell, \psi) = \Pr(s_\ell|r)\, EV(\sigma_\ell) + [1 - \Pr(s_\ell|r)]\, EV(1-\sigma_h) - F. \tag{9}$$

If the principal rejects the agent's offer, then $s = \varnothing$ and his ex-post belief coincides with the interim belief, i.e., $\mu = \psi$. His expected utility is given by (7) where μ is replaced by ψ, i.e., $EV(\psi, \theta_h, \theta_\ell)$. In this case the principal relies exclusively on his interim belief.

The set of acceptable offers is

$$\mathcal{B}(r, F|\psi) = \left\{ (r, F) \in \left[\frac{1}{2}, 1\right] \times \mathbb{R}_+ : EV(r, F|\theta_h, \theta_\ell, \psi) \geq EV(\psi, \theta_h, \theta_\ell) \right\}.$$

The next lemma characterizes the set $\mathcal{B}(r, F|\psi)$.

Lemma 1. *The set of acceptable offers for the principal is*

$$\mathcal{B}(r, F|\psi) = \left\{ (r, F) \in \left[\frac{1}{2}, 1\right] \times \left[0, \bar{F}(r|\psi)\right] \right\},$$

where $\bar{F}(r|\psi)$ *is an increasing function with* $\bar{F}\left(\frac{1}{2}|\psi\right) = 0$, for all ψ.

The proof is straightforward but lengthy and it is omitted.[19]

The fact that $\bar{F}(r|\psi)$ is an increasing function suggests that the principal is willing to pay a higher price for the agent's PD, only if the information is of a higher quality. In figures 3 and 4 we graph $\bar{F}(r|\psi)$, assuming that $S(q) = \sqrt{q}$. The principal will reject any offer $(r, F) \notin \mathcal{B}(r, F|\psi)$. The set of acceptable offers $\mathcal{B}(r, F|\psi)$ depends on the principal's interim belief ψ. If, for example, ψ is very high or very low, then the agent's PD is of little value and the set $\mathcal{B}(r, F|\psi)$ shrinks. Most offers in that case would be rejected.

3.3 Stage 1: The agent sells his private data (PD) to the principal (signaling stage)

The agent decides which offer (r, F) to make.[20] The agent has all the bargaining power at this stage and therefore we assume that any offer is on the boundary of the set $\mathcal{B}(r, F|\psi)$, that is, for any r the agent asks for the maximum possible price $\bar{F}(r|\psi)$. The principal then is indifferent between accepting and rejecting the agent's offer and we assume he accepts. The inefficient type's expected surplus (excluding the price for his PD) is always zero. This is because the contract $\mathcal{C}$ that

---

[19] It can be found at http://faculty.lebow.drexel.edu/SerfesK/Working_papers_and_work_in_progress.htm

[20] To avoid confusion, we would like to draw the reader's attention on the following matter. The signaling that may take place in stage 1 is different from the signal s that the principal draws from the agent's PD if he accepts the agent's offer in stage 2. For lack of better terminology, we use similar terms to refer to these two distinct cases.

is offered in stage 3 always leaves the inefficient type with zero surplus. Therefore, the expected utility of the inefficient type, if the principal accepts the agent's offer, is $EU_h = F$. The efficient type's expected utility, if the principal accepts the agent's offer, is given by,

$$EU_\ell\left(r, F|\theta_h, \theta_\ell, \psi\right) = r\Delta\theta q_h^{SB}\left(\sigma_\ell\right) + \left(1-r\right)\Delta\theta q_h^{SB}\left(1-\sigma_h\right) + F. \tag{10}$$

where $q_h^{SB}\left(\cdot\right)$ satisfies (4) and its argument is the probability of the efficient type given the signal that the principal has drawn. With probability r the principal will draw a low signal in which case the efficient type's utility will be $\Delta\theta q_h^{SB}\left(\sigma_\ell\right)$ and with probability $1-r$ he will draw a high signal in which case the efficient type's utility will be $\Delta\theta q_h^{SB}\left(1-\sigma_h\right)$.

The next lemma characterizes the indifference map of the two types,

Lemma 2.

- *The indifference map of the efficient type is increasing for high r's, i.e., $\frac{dF}{dr} > 0$, for all $r \geq \bar{r}$.*
- *If $\psi \leq \underline{\psi}$, then the indifference map of the efficient type is increasing, i.e., $\frac{dF}{dr} \geq 0$ for all $r \in \left[\frac{1}{2}, 1\right]$.*
- *The indifference map of the efficient type may be non-monotonic for high ψ's, i.e., $\frac{dF}{dr} \lesseqgtr 0$, as the example below illustrates.*
- *The indifference map of the inefficient type is horizontal, i.e., $\frac{dF}{dr} = 0$.*

The proof is straightforward and it is omitted.[21]

Suppose, for example, that $S\left(q\right) = \sqrt{q}$. Then it follows (details are omitted) that the efficient type's indifference map exhibits a U-shape if $\psi = .9$ and $\Delta\theta = 2$ (see figure 3). On the other hand, if $\psi = .3$, then the indifference map is strictly increasing (see figure 4). (The agents become better off when they move in the northwest direction while the principal becomes better off when he moves in the southeast direction.) Therefore, the indifference curves may not satisfy the *Spence-Mirrlees single-crossing property*.

The intuition behind the non-monotonicity of the efficient type's indifference map is as follows. The efficient type's utility (before any information is traded) is given by (5), where now $\mu = \psi$. It can be readily verified that,

$$\frac{d^2 U_\ell}{d\psi^2} = \frac{2\left(\Delta\theta\right)^2}{S'''\left(q^{SB}\left(\psi\right)\right)\left(1-\psi\right)^3}.$$

This suggests that the efficient agent's utility function is convex in the probability of the efficient type ψ if and only if $S''' > 0$. When $S\left(q\right) = \sqrt{q}$, the third derivative is indeed positive. Moreover,

---

[21] It can be found at http://faculty.lebow.drexel.edu/SerfesK/Working_papers_and_work_in_progress.htm

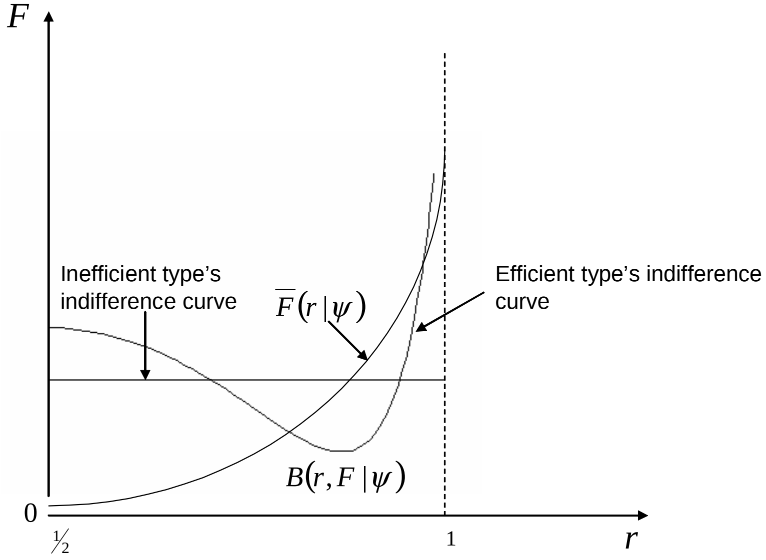

Figure 3: Indifference curves when ψ is high

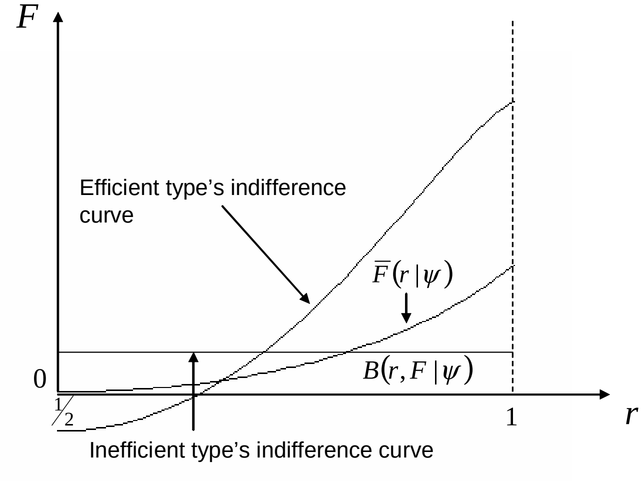

Figure 4: Indifference curves when ψ is low

the utility is decreasing in ψ. Now let's look at the effect of information of quality r on the efficient type's expected utility. The expected probability that the agent is efficient is given by,[22]

$$E\left[\Pr\left(\ell|r\right)\right]=\Pr\left(\ell|s_{\ell}\right)\Pr\left(s_{\ell}|r\right)+\Pr\left(\ell|s_{h}\right)\Pr\left(s_{h}|r\right)=\frac{\psi(4r^{2}\psi-3r^{2}+3r-4r\psi-1+\psi)}{(2r\psi+1-\psi-r)(2r\psi-\psi-r)}.$$

The probability the principal will attach on the agent being efficient [denoted by $\Pr\left(\ell|r\right)$], given that the agent is indeed efficient, is a random variable (from the perspective of the efficient type agent), which can take either a high value (if the signal is s_{ℓ}) or a low value (if the signal is s_{h}) with mean $E\left[\Pr\left(\ell|r\right)\right]$. It can be easily shown that $E\left[\Pr\left(\ell|r\right)\right]$ is strictly increasing in r, with $E\left[\Pr\left(\ell|r\right)\right]=\psi$, at $r=\frac{1}{2}$ and $E\left[\Pr\left(\ell|r\right)\right]=1$, at $r=1$. Two things happen as the informativeness r of the signal increases: i) the spread between the probability of an efficient type when the signal is low and that when the signal is high increases; a direct consequence of the fact that the signal has become more accurate and ii) the mean probability increases as well, because the signal is more accurate and the agent is efficient. This gives rise to two opposing forces that affect the efficient agent's expected utility. If we hold the mean of $\Pr\left(\ell|r\right)$ fixed at (say) ψ, an increase in r increases the agent's expected utility (*positive effect*). This follows from *Jensen's inequality* and the fact that the agent's utility is convex in the probability of the efficient type. Now allow the mean probability to vary. As r increases and the mean of $\Pr\left(\ell|r\right)$ increases as well the expected utility decreases (*negative effect*). The efficient agent looses when the principal identifies him more often on average. When the positive effect dominates the negative the efficient agent's expected utility is increasing in r (in which case the indifference map is decreasing) and vice versa. This happens when r is low and ψ is high.

The inefficient type is always better off with a higher F regardless of r and vice verse. This is because his surplus in stage 3 is always zero and therefore his expected utility in stage 1 equals the price for his PD. We will exploit this property later in the paper when we construct the out-of-equilibrium beliefs.

We continue by searching for a separating equilibrium.

Proposition 1. *A separating equilibrium does not exist.*

Proof. Consider a candidate equilibrium where the efficient type offers (r_{ℓ},F_{ℓ}) and the inefficient type offers (r_{h},F_{h}) such that $(r_{\ell},F_{\ell})\neq(r_{h},F_{h})$. Since the outcome is separating, we must have $\psi\left(r_{\ell},F_{\ell}\right)=1$. Because the principal knows the type of the agent who makes the offer, he does not have to purchase the agent's PD. There is no asymmetric information anymore and the principal can offer the efficient quantities q_{h}^{*} and q_{ℓ}^{*}. The agent's utilities are $U_{\ell}=U_{h}=0$.

---

[22] The probabilities below are the probabilities of an efficient type, conditional on a signal of quality r, that the principal will use when he designs a contract, *given* that the agent is indeed efficient. The principal, of course, does not know that at this stage, but the agent, who is the one who chooses the signal quality, does. In contrast, from the principal's perspective the signal is unbiased, i.e., $E\left[\Pr\left(\ell|r\right)\right]=\psi$.

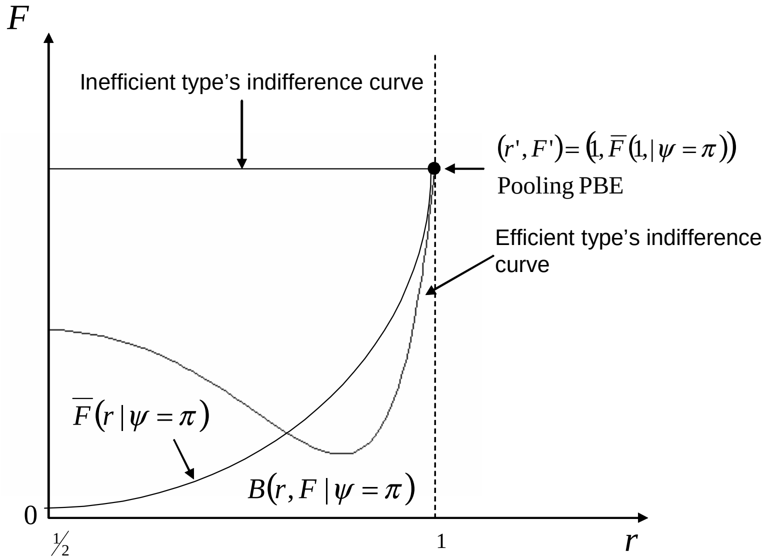

Figure 5: Pooling perfect Bayesian equilibrium

Now we check whether the efficient type has a profitable deviation. In particular, we examine a deviation to (r_h, F_h). Then, the efficient type is perceived as an inefficient type with probability 1, which gives him surplus of $\Delta\theta q_h^* > 0$. Thus, such a deviation is always profitable. ■

Next, we look for a pooling equilibrium, where both types offer (r', F').

Theorem 1 (Main result). *The following is the unique pooling universally divine PBE of the three-stage game:*

- *Stage 1: The agent offers* $(r', F') = \left(1, \bar{F}(1|\psi = \pi)\right)$ *regardless of his type, where*

$$\bar{F}(1|\psi = \pi) = (1 - \pi)\left([S(q_h^*) - \theta_h q_h^*] - \left[S\left(q_h^{SB}(\pi)\right) - \theta_h q_h^{SB}(\pi)\right]\right) + \pi\Delta\theta q_h^{SB}(\pi).$$

- *Stage 2: The principal is indifferent between accepting and rejecting* (r', F') *and he accepts.*
- *Stage 3: The principal offers the efficient (first-best) quantities:* q_ℓ^* *and* $t_\ell^* = \theta_\ell q_\ell^*$ *if the signal is low* $(s = s_\ell)$*, or* q_h^* *and* $t_h^* = \theta_h q_h^*$ *if the signal is high* $(s = s_h)$*.*

The inefficient type always becomes better-off when he can sell his PD. The efficient type becomes better-off when he can sell his PD if and only if the ex-ante probability of an efficient type is high, i.e., $\pi \geq \tilde{\pi}$*.*

Proof. See appendix. ■

Both agent types make the same offer in equilibrium. The offer entails *all* of the agent's private data (PD) and a positive price. Moreover, the offer itself does not reveal any information about the agent's type. The principal upon purchasing the agent's PD learns the agent's type with probability 1 (since the agent sold all his PD) and offers an efficient (first-best) contract. Because the agent gets compensated for selling information about himself, he ends up with positive surplus. Moreover, when the ex-ante probability of the efficient type is high, the low cost (efficient) agent is better off when he is able to sell his PD relative to when he cannot; otherwise he is worse off. The inefficient type is always better off.

The above discussion, however, does not imply that the efficient type maximizes his expected utility even when $\pi \geq \tilde{\pi}$. Theorem 1 says that the efficient type is better off at the pooling PBE with information selling than without information selling. But he would have enjoyed an even higher expected utility if he could have sold only a fraction of his PD, i.e., $r \in \left(\frac{1}{2}, 1\right)$. In the discussion after lemma 2, we offer an example where the efficient type's indifference map is U-shaped, for high π's (recall that in the pooling equilibrium $\pi = \psi$). From lemma 1, we know that the principal's indifference map is increasing. It follows, that the r that maximizes the efficient type's expected utility, subject to the constraint that the principal is indifferent between accepting and rejecting the agent's offer, must be strictly less than 1. But, of course, as theorem 1 has demonstrated, such an offer is not an equilibrium. Figure 5 depicts the equilibrium when π is high (if π is low, then the efficient type's indifference map is increasing; see lemma 2 and figure 4). From figure 5 it is clear that both types become better off when they can sell their PD relative to when they cannot, in which case the indifference curves begin at $(r, F) = \left(\frac{1}{2}, 0\right)$. Furthermore, it is apparent, from the same graph, that the r that maximizes the efficient type's expected utility is not 1, but rather it is in $\left(\frac{1}{2}, 1\right)$.

3.3.1 Intuition for the main result

As the theorem above states, the regulated equilibrium is characterized as follows: i) both agent types make the same offer (pooling equilibrium) and ii) the offer entails all of the agent's PD. Why the two different agent types pool their offers in equilibrium has been explained before. The second result is less obvious. Any pooling outcome where agents sell only a fraction of the PD is not an equilibrium. First, it is a standard result in screening models that the inefficient type always receives zero surplus. Hence, in our context, the inefficient type always has incentives to deviate by offering more of his PD for sale at a higher price, since that price is equal to his surplus. What may prevent him from doing so is the possibility that his action may change the principal's ex-ante beliefs drastically. Nevertheless, as we explain below, this is *not* the case.

Under some set of beliefs held by the principal (call this set A), about the likelihood of an efficient type following a deviation, the efficient type has an incentive to offer more of his PD for sale but not the inefficient type. To see this, suppose that the principal believes, after a deviation, that the probability of an efficient type is arbitrarily close to zero (and hence the probability of an inefficient type is arbitrarily close to one). The principal's best response is not to accept the offer because his beliefs are very accurate, and hence the information is useless. This will make the inefficient type worse off because his surplus from the contract in stage 3 is always zero and on top of that he also foregoes the revenue from the information selling. The efficient type, however, will become better off because he receives more surplus from the contract in stage 3 (since the principal believes that his type is extremely unlikely), which, as we show, outweighs the foregone revenue from the information selling. Under some other set of beliefs (call this set B) it is the inefficient type who wishes to sell more of his PD but not the efficient type. To see this, suppose that beliefs are not updated much after a deviation. Clearly, the inefficient type would want to deviate and offer more of his PD for sale at a higher price. But more accurate information, as a result of more PD being sold, will lead to lower surplus for the efficient type, if the deviation entails an offer with a very high fraction of the PD (and hence a very accurate signal). Moreover, there does not exist a price for the PD that would compensate the efficient type and at the same time keep the principal indifferent between accepting and rejecting the offer. This is because the principal's maximum willingness to pay for the PD is a weighted average (across the two agent types), whereas the efficient agent knows his type for sure. These two sets of beliefs (A and B) cannot be nested and according to the Banks and Sobel universal divinity refinement both types are *likely* to deviate. This implies that the principal does not update his ex-ante beliefs following a deviation (in a specific region) and the inefficient type would always want to deviate from an offer that entails only a fraction of the PD by offering more of his PD for sale at a higher price.

Now consider the pooling equilibrium offer (where all the information is sold). Any deviation from $(r', F') = \left(1, \bar{F}\left(1|\psi = \pi\right)\right)$ is viewed from the principal as if it has come from the efficient type with probability 1. Whenever the inefficient type wants to deviate so does the efficient type, but not necessarily the other way around (the reason is offered in the proof of theorem 1). This implies that the efficient type is "more likely" to deviate. So, according to the universal divinity refinement any deviation should be perceived as if it has come from the efficient type with probability 1. This belief also satisfies the Cho and Kreps (1987) *intuitive criterion*. However, if we insist on beliefs that only satisfy the intuitive criterion, then we may end up with a continuum of equilibria. It is not difficult to see why. Consider any interior point $(r'', F'') \neq \left(1, \bar{F}\left(1|\psi = \pi\right)\right)$ on $\bar{F}\left(r, |\psi = \pi\right)$. If any deviation is viewed as if it has come from the efficient type, then no type would find a deviation profitable; hence we obtain a continuum of PBE. These beliefs do satisfy the intuitive criterion, because the efficient type actually has a profitable deviation (under some beliefs). However, they are *not* universally divine because, as we argued in the above paragraph, there exists a region (which does not exist when $(r', F') = \left(1, \bar{F}\left(1|\psi = \pi\right)\right)$) where *both* types are *likely* to deviate.[23]

---

[23] This region is denoted by $\mathcal{D}_h\left(r'', F''\right) \cap \mathcal{D}_\ell^c\left(r'', F''\right)$ in figure 7 in the Appendix.

By insisting that the beliefs must pass the universal divinity refinement, then we obtain a unique pooling PBE.

4 Conclusion

Our main purpose in this paper is to examine the efficiency and re-distributive consequences of a regulation that would give consumers control of their own private information. We cast this problem in a principal agent adverse selection framework. We can think of the principal as a firm who will make product offers to the agent (consumer, or group of consumers). The agent's private information is assumed to be contained in a database and the agent, if he has control of his information, can choose how much private data (PD) to offer for sale and the corresponding price for his PD. The agent's PD generates an informative signal that improves the principal's contracting abilities. The quality (accuracy) of the signal is in direct relationship with how much PD the agent chooses to sell. If he sells all his PD, then the signal is perfectly informative. The main concern is that if agents have the property rights to their own information, flows of information may be restrained, with adverse effects on efficiency. The agent runs the risk of revealing, through his choice of how much of his private information to sell, information about his type without getting compensated for it. If, for example, the principal perceives the agent as being efficient with sufficiently high probability, then the agent's PD is of little value to the principal. Rather, the principal will rely *exclusively* on his belief when he designs the contract. In this case, the agent may choose not to sell his private information. We show, however, that this will not be the case. A unique pooling universally divine PBE exists where the agent succeeds in selling *all* his PD. The principal then draws a perfect signal and offers the first-best contract. Efficiency is also achieved if, instead, a third party has control of the agent's information. Our results, therefore, indicate that the Coase theorem holds in an asymmetric information environment. Finally, the agent is better off if he has control of his own information.

Further work needs to be done before we can draw a more clear picture about the impact of consumer private information on market efficiency. We hope that this paper will spawn more research on this important issue. One interesting extension is to allow for competing principals (common agency).

A Appendix: Proofs

A.1 Lemma 3

Lemma 3 below presents a key result which is used repeatedly in the proof of theorem 1.

Lemma 3. *Consider a pooling offer* (r, F) *on* $\bar{F}(r|\psi = \pi)$*. Then, the efficient type always finds it profitable to deviate to* (r', F')*, if he is viewed as an inefficient type with a very high probability (i.e.,* $\psi(r', F') \in [0, \bar{\psi}]$*), in which case the principal will reject the efficient type's offer and he will rely exclusively on his interim belief* ψ*.*

The above lemma says the following. Suppose that the efficient type can make an offer (r, F), without signaling his type, to the principal which is accepted. The principal's ex-post beliefs are given by μ, i.e., they depend on which signal the principal will draw in stage 3. The efficient type is better off by deviating to another offer (*any offer*) if such a deviation is perceived as if it has come from the inefficient type with a sufficiently high probability. In this case the principal does not accept the efficient agent's offer (because information has little value since ψ is very close to zero) and he offers a contract which entails very little distortion to the output designed for the inefficient type, i.e., $\left(q_h^{SB}, t_h^{SB}\right)$ is very close to (q_h^*, t_h^*), because $\psi(r', F')$ is very close to zero. This clearly increases the efficient type's utility relative to his pre-deviation utility from the contract $\mathcal{C}$ (excluding the price F), but in the process he also foregoes the price F that he would have received had he not deviated. Lemma 3 demonstrates that the price F is lower than the benefit the efficient type derives when he deviates.

Proof of lemma 3. Suppose that $\psi(r', F') = 0$. The efficient type's utility before deviation is given by (10). His utility after deviation is given by,

$$EU_\ell^{dev} = \Delta\theta q_h^*.$$

The principal does not distort the inefficient type's quantity, since he believes that the probability that the agent is efficient is zero. Assume that the price F is the maximum possible price that the principal will accept for any r, i.e., $\bar{F}(r|\psi = \pi)$. At this price the principal is indifferent between accepting and rejecting the agent's offer and the price can be expressed as follows,

$$
\begin{aligned}
\bar{F}\left(r|\psi=\pi\right) &\equiv \Pr\left(s_{\ell}|r\right)EV\left(\sigma_{\ell}\right)+\left[1-\Pr\left(s_{\ell}|r\right)\right]EV\left(1-\sigma_{h}\right)- \\
&\left[\pi\left[S\left(q_{\ell}^{*}\right)-\theta_{\ell}q_{\ell}^{*}\right]+\left(1-\pi\right)\left[S\left(q_{h}^{SB}\right)-\theta_{h}q_{h}^{SB}\right]-\pi\Delta\theta q_{h}^{SB}\left(\pi\right)\right] \\
&= \left[\pi r+\left(1-\pi\right)\left(1-r\right)\right]\left[\begin{array}{c}\sigma_{\ell}\left[S\left(q_{\ell}^{*}\right)-\theta_{\ell}q_{\ell}^{*}\right]+ \\ \left(1-\sigma_{\ell}\right)\left[S\left(q_{h}^{SB}\left(\sigma_{\ell}\right)\right)-\theta_{h}q_{h}^{SB}\left(\sigma_{\ell}\right)\right]- \\ \sigma_{\ell}\Delta\theta q_{h}^{SB}\left(\sigma_{\ell}\right)\end{array}\right]+ \\
&\left[1-\pi r-\left(1-\pi\right)\left(1-r\right)\right]\left[\begin{array}{c}\left(1-\sigma_{h}\right)\left[S\left(q_{\ell}^{*}\right)-\theta_{\ell}q_{\ell}^{*}\right]+ \\ \sigma_{h}\left[S\left(q_{h}^{SB}\left(1-\sigma_{h}\right)\right)-\theta_{h}q_{h}^{SB}\left(1-\sigma_{h}\right)\right]- \\ \left(1-\sigma_{h}\right)\Delta\theta q_{h}^{SB}\left(1-\sigma_{h}\right)\end{array}\right]- \\
&\left[\begin{array}{c}\pi\left[S\left(q_{\ell}^{*}\right)-\theta_{\ell}q_{\ell}^{*}\right]+ \\ \left(1-\pi\right)\left[S\left(q_{h}^{SB}\right)-\theta_{h}q_{h}^{SB}\right]-\pi\Delta\theta q_{h}^{SB}\left(\pi\right)\end{array}\right] \\
&= \left(1-r-\pi+r\pi\right)\left[S\left(q_{h}^{SB}\left(\sigma_{\ell}\right)\right)-\theta_{h}q_{h}^{SB}\left(\sigma_{\ell}\right)\right]-r\pi\Delta\theta q_{h}^{SB}\left(\sigma_{\ell}\right)+ \\
&r\left(1-\pi\right)\left[S\left(q_{h}^{SB}\left(1-\sigma_{h}\right)\right)-\theta_{h}q_{h}^{SB}\left(1-\sigma_{h}\right)\right]-\pi\left(1-r\right)\Delta\theta q_{h}^{SB}\left(1-\sigma_{h}\right)- \\
&\left(1-\pi\right)\left[S\left(q_{h}^{SB}\left(\pi\right)\right)-\theta_{h}q_{h}^{SB}\left(\pi\right)\right]+\pi\Delta\theta q_{h}^{SB}\left(\pi\right).
\end{aligned}
$$

Substitute $\bar{F}\left(r|\psi=\pi\right)$ into EU_{ℓ} as it is given by (10),

$$
\begin{aligned}
EU_{\ell}\left(r,F|\theta_{h},\theta_{\ell},\pi\right) &= \left(1-r\right)\left(1-\pi\right)\left[S\left(q_{h}^{SB}\left(\sigma_{\ell}\right)\right)-\theta_{h}q_{h}^{SB}\left(\sigma_{\ell}\right)\right]+ \\
&r\left(1-\pi\right)\left[S\left(q_{h}^{SB}\left(1-\sigma_{h}\right)\right)-\theta_{h}q_{h}^{SB}\left(1-\sigma_{h}\right)\right]- \\
&\left(1-\pi\right)\left[S\left(q_{h}^{SB}\right)-\theta_{h}q_{h}^{SB}\left(\pi\right)\right]+\left(1-r\right)\left(1-\pi\right)\Delta\theta q_{h}^{SB}\left(1-\sigma_{h}\right)+ \\
&r\left(1-\pi\right)\Delta\theta q_{h}^{SB}\left(\sigma_{\ell}\right)+\pi\Delta\theta q_{h}^{SB}\left(\pi\right).
\end{aligned}
$$

The change in expected utility, between before and after deviation, is given by,

$$
\begin{aligned}
\Delta EU_{\ell} &= EU_{\ell}\left(r,F|\theta_{h},\theta_{\ell},\pi\right)-EU_{\ell}^{dev} \\
&= \{\left(1-r\right)\left(1-\pi\right)\left[S\left(q_{h}^{SB}\left(\sigma_{\ell}\right)\right)-\theta_{h}q_{h}^{SB}\left(\sigma_{\ell}\right)\right]+ \\
&r\left(1-\pi\right)\left[S\left(q_{h}^{SB}\left(1-\sigma_{h}\right)\right)-\theta_{h}q_{h}^{SB}\left(1-\sigma_{h}\right)\right]- \\
&\left(1-\pi\right)\left[S\left(q_{h}^{SB}\left(\pi\right)\right)-\theta_{h}q_{h}^{SB}\left(\pi\right)\right]\}+ \\
&\left\{\left(1-r\right)\left(1-\pi\right)\Delta\theta q_{h}^{SB}\left(1-\sigma_{h}\right)+r\left(1-\pi\right)\Delta\theta q_{h}^{SB}\left(\sigma_{\ell}\right)+\pi\Delta\theta q_{h}^{SB}\left(\pi\right)-\Delta\theta q_{h}^{*}\right\}.
\end{aligned}
$$

The term in the first curly brackets is negative. First, we assume that the benefit function $S\left(\cdot\right)$ is sufficiently concave in q, so that the function $S\left(q_{h}^{SB}\left(\cdot\right)\right)-\theta_{h}q_{h}^{SB}\left(\cdot\right)$ is concave in the probability that the agent is of a high type (inefficient).[24] Second, the probability that the agent is of a high

---

[24] Note that $q_{h}^{SB}\left(\cdot\right)$ is not necessarily concave in the probability that the agent is of a high type, unless $S''' < 0$. For example, when $S\left(q\right)=\sqrt{q}$, the third derivative is positive, i.e., $S''' > 0$. Nevertheless, it turns out that $S\left(q_{h}^{SB}\left(\cdot\right)\right)-\theta_{h}q_{h}^{SB}\left(\cdot\right)$ is concave in the probability of the inefficient type.

type, given information of quality r, is a random variable with mean $1-\pi$, i.e.,

$$\begin{aligned}
\Pr(h|r) &= \Pr(h|s_\ell)\Pr(s_\ell|r) + \Pr(h|s_h)\Pr(s_h|r) \\
&= \left[\frac{(1-r)(1-\pi)}{r\pi + (1-r)(1-\pi)}\right][\pi r + (1-\pi)(1-r)] + \\
&\quad \left[\frac{r(1-\pi)}{r(1-\pi) + (1-r)\pi}\right][1 - \pi r - (1-\pi)(1-r)] \\
&= (1-r)(1-\pi) + r(1-\pi) = 1-\pi.
\end{aligned}$$

Then, from *Jensen's inequality*,

$$\begin{aligned}
E\left[S\left(q_h^{SB}(\sigma) - \theta_h q_h^{SB}(\sigma)\right)\right] &= (1-r)\left[S\left(q_h^{SB}(\sigma_\ell)\right) - \theta_h q_h^{SB}(\sigma_\ell)\right] + \\
&\quad r\left[S\left(q_h^{SB}(1-\sigma_h)\right) - \theta_h q_h^{SB}(1-\sigma_h)\right] \\
&< \left[S\left(q_h^{SB}(\pi)\right) - \theta_h q_h^{SB}(\pi)\right],
\end{aligned}$$

which implies, if we multiply both sides of the above inequality by $(1-\pi)$, that the term in the first curly brackets is negative.

The term in the second curly brackets in ΔEU_ℓ is also negative. This follows easily from the facts that: i) $q_h^* > q_h^{SB}(1-\sigma_h) > q_h^{SB}(\pi) > q_h^{SB}(\sigma_\ell)$ and ii) $(1-r)(1-\pi) + r(1-\pi) + \pi = 1$. Therefore, $\Delta EU_\ell < 0$ which implies that a deviation from any (r, F) on part of the efficient type as long as he is perceived as an efficient type with probability zero is always profitable. From continuity, the result should hold in a neighborhood of zero, say $\left[0, \bar{\psi}\right]$. ■

A.2 Proof of theorem 1

Fix a pooling outcome $(r', F') = \left(1, \bar{F}(1|\psi = \pi)\right)$; in other words, the offer is on the boundary of the set $\mathcal{B}(r, F|\psi = \pi)$. Denote by $I_h(r', F'|\psi = \pi)$ and $I_\ell(r', F'|\psi = \pi)$ the indifference curves (ICs) of the inefficient and efficient type respectively that pass through (r', F'). Next let,

$$\mathcal{D}_h\left(r', F'\right) = \left\{(r, F) : \text{that belong to IC's higher than } I_h\left(r', F'|\psi = \pi\right)\right\}$$

$$\mathcal{D}_\ell\left(r', F'\right) = \left\{(r, F) : \text{that belong to IC's higher than } I_\ell\left(r', F'|\psi = \pi\right)\right\}.$$

The complements of the above sets are given by $\mathcal{D}_h^c(r', F')$ and $\mathcal{D}_\ell^c(r', F')$. The indifference map depends on the principal's interim belief ψ. If $\psi \neq \pi$, which may happen when we consider out-of-equilibrium beliefs, then we obtain a different indifference map than the one when $\psi = \pi$. Nevertheless, the sets defined above assume that $\psi = \pi$ and this is done simply for the purpose of dividing the (r, F) space into different regions, starting from $(r', F') = \left(1, \bar{F}(1|\psi = \pi)\right)$, see figure 6. For instance, when we consider a deviation $\left(r^{dev}, F^{dev}\right) \in \mathcal{D}_h(r', F')$, we allow for $\psi \neq \pi$.

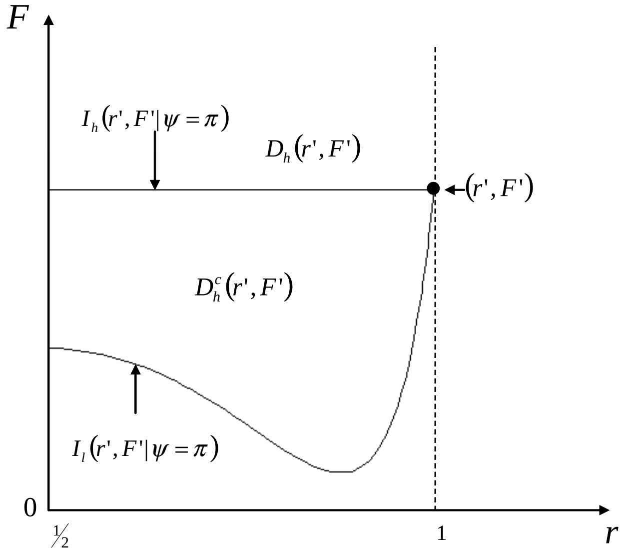

Figure 6: Various sets in the (r, F) region starting from (r', F')

To increase the predictive power of signaling games, restrictions must be made on the out-of-equilibrium beliefs. To this end, define the following belief function (see also figure 6),

$$\psi(r, F) = \begin{cases} \pi, & \text{if } (r, F) = (r', F') \\ 1, & \text{if } (r, F) \in \mathcal{D}_h(r', F') \\ 1, & \text{if } (r, F) \in \mathcal{D}_h^c(r', F') \,. \end{cases} \tag{11}$$

Because both types make the same offer, we assume that no additional information is revealed when the offer is (r', F'), i.e., the principal's interim belief ψ coincides with the ex-ante probability π. Thus, the above belief function is anchored around the pooling outcome $(r', F') = \left(1, \bar{F}(1|\psi = \pi)\right)$. According to this belief function, any deviation from $(r', F') = \left(1, \bar{F}(1|\psi = \pi)\right)$ is perceived as if it has come from the efficient type with probability 1. First, we show that the pooling outcome together with these beliefs constitute a universally divine pooling PBE. Second we show that this is the unique universally divine pooling PBE.

Existence. First, we show that given (11), the strategies followed by the principal and the agent are optimal. Since $\psi = \pi$ and $(r', F') \in \mathcal{B}(r, F|\psi = \pi)$ the principal's best response is to accept. Next, we show that it is each type's best response to offer (r', F'). Before deviation, each type enjoys strictly positive expected surplus, i.e., $EU_h = EU_\ell = \bar{F}(1|\psi = \pi) > 0$. The signal is perfect and each agent's surplus is only the price of the PD, as in stage 3 the principal will offer the efficient (first-best) quantities. The agent would not find any deviation profitable regardless of his type. Following any deviation, the agent is perceived as an efficient type with certainty. The principal in stage 2 rejects the agent's offer and in stage 3 he offers q_h^* and $t_h^* = \theta_h q_h^*$. If it is the efficient type who deviated his surplus is zero. If it is the inefficient type who deviated, he will not accept the contract in stage 3 and he also receives zero surplus. It follows that neither type would want to deviate.

Finally, the beliefs are derived using Bayes' rule from the equilibrium strategies.

We now demonstrate that the beliefs, given by (11), pass the Banks and Sobel (1987) *universal divinity refinement*. First, a deviation to $\mathcal{D}_h^c(r', F')$, must have come from the efficient type. The inefficient type obtains lower utility if he deviates that way, no matter how the principal responds. Second, consider a deviation to $\mathcal{D}_h(r', F')$. We will show that the efficient type is "more likely" to deviate in that region. Put simply, when the inefficient type wants to deviate so does the efficient type, but not always the other way around. Then, according to the universal divinity refinement, the principal must attach probability 1 that such a deviation has come from the efficient type. To begin with, the principal can either accept or reject the deviator's offer. If the offer gets accepted both types become better off. Why this is so is clear for the inefficient type since he receives a higher price for his PD and the surplus he receives from the contract $\mathcal{C}$ is always zero. Now let's turn our attention to the efficient type. The efficient type at $r = 1$ receives zero surplus from the contract $\mathcal{C}$ in stage 3. By deviating in $\mathcal{D}_h(r', F')$, he becomes strictly better off because the price for his PD increases and $r \leq 1$, which implies that he receives some surplus from the contract $\mathcal{C}$ (recall that the inefficient type's indifference map is horizontal). Next, suppose that the principal rejects the agent's offer. Clearly, the inefficient type always becomes worse off, since he foregoes the price for his PD. However, according to lemma 3, the efficient type may become better off if he is perceived as an inefficient type with a sufficiently high probability.

Uniqueness. We show that $(r', F') = \left(1, \bar{F}(1|\psi = \pi)\right)$ is the unique pooling universally divine PBE. By way of contradiction, consider an offer (r'', F'') on $\bar{F}(r|\psi = \pi)$ that is different from $\left(1, \bar{F}(1|\psi = \pi)\right)$. Define the following belief function (see also figure 7),

$$\psi(r, F) = \begin{cases} \pi, & \text{if } (r, F) = (r'', F'') \\ \pi, & \text{if } (r, F) \in \mathcal{D}_h(r'', F'') \cap \mathcal{D}_\ell^c(r'', F'') \\ [\pi, 1], & \text{if } (r, F) \in \mathcal{D}_h(r'', F'') \cap \mathcal{D}_\ell(r'', F'') \\ 1, & \text{if } (r, F) \in \mathcal{D}_h^c(r'', F''). \end{cases} \tag{12}$$

The difference between (11) and (12) is that (11) deals with deviations from $(r', F') = \left(1, \bar{F}(1|\psi = \pi)\right)$ whereas (12) deals with deviations from $(r'', F'') \neq (r', F')$. For example, the set $\mathcal{D}_h(r', F') \cap \mathcal{D}_\ell^c(r', F') = \varnothing$ (see figure 5), while, as we show below, $\mathcal{D}_h(r'', F'') \cap \mathcal{D}_\ell^c(r'', F'') \neq \varnothing$.

First note that (r'', F'') will necessarily entail $F'' < \bar{F}(1|\psi = \pi)$ and $r'' < 1$ (recall, from lemma 1, that $\bar{F}(r|\psi = \pi)$ is strictly increasing). The inefficient type can deviate to $\left(r^{dev}, F^{dev}\right) \in \mathcal{D}_h(r'', F'') \cap \mathcal{D}_\ell^c(r'', F'')$ and become strictly better off, since from (12) such a deviation does not change the priors and the inefficient type moves to a higher indifference curve. Moreover, $\mathcal{D}_h(r'', F'') \cap \mathcal{D}_\ell^c(r'', F'') \neq \varnothing$. At $r = 1$ the price $F(1|\psi = \pi)$ that the efficient type would like to receive in order to remain indifferent between $(1, F(1|\psi = \pi))$ and (r'', F'') must be higher than F''.

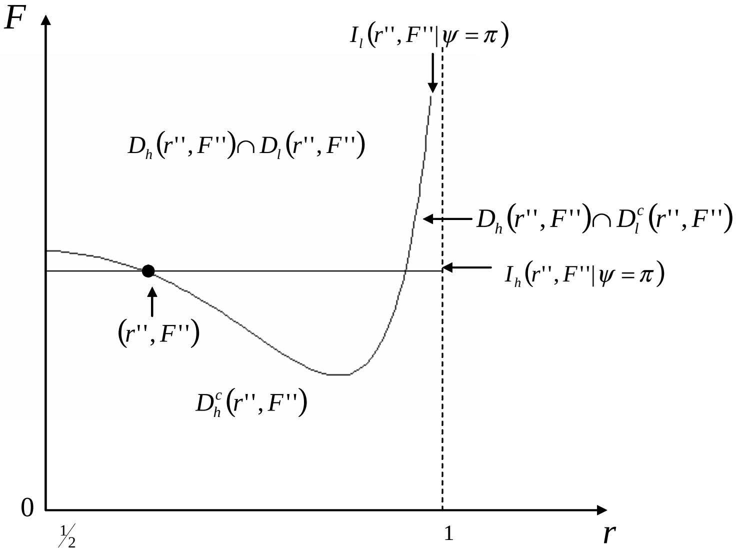

Figure 7: Various sets in the (r, F) region starting from (r'', F'')

This is true because at $r'' < 1$ the efficient type's surplus from the contract $\mathcal{C}$ is higher than that at $r = 1$, where his surplus is zero. Moreover, the inefficient type's indifference curve is horizontal. These imply that $\mathcal{D}_h(r'', F'') \cap \mathcal{D}_\ell^c(r'', F'') \neq \varnothing$ for r's that are close to 1, as it is also clear from figure 7.

We now show that (12) passes the Banks and Sobel (1987) universal divinity refinement. Consider a specific deviation from the candidate pooling equilibrium.

First, a deviation to $\mathcal{D}_h^c(r'', F'')$, must have come from the efficient type. The inefficient type obtains lower utility if he deviates that way, no matter how the principal responds.

Second, we examine a deviation in $\mathcal{D}_h(r'', F'') \cap \mathcal{D}_\ell(r'', F'')$. We will show that the inefficient type is never more likely to deviate in that set. This would suggest that the belief should attach at least probability π to such a deviation having come from the efficient type. The principal can either accept or reject the agent's offer. Suppose the principal rejects the offer. The inefficient type becomes worse off no matter what, while the efficient type (from lemma 3) may become better off. In particular, the efficient type becomes better off when the principal attaches a very high probability that the deviator is the inefficient type. Such a response on the part of the principal is indeed optimal if he believes that the deviator is inefficient with a very high probability. Additional information, in this case, is not valued much and the best response is not to acquire it. This suggests that the inefficient type cannot be more likely to deviate to $\mathcal{D}_h(r'', F'') \cap \mathcal{D}_\ell(r'', F'')$ than the efficient type. Hence, the beliefs must be in the range $[\pi, 1]$. We do not attempt to sharpen the beliefs more, since the non-existence of a PBE does not depend on the values the beliefs take in the $\mathcal{D}_h(r'', F'') \cap \mathcal{D}_\ell(r'', F'')$ region.

Third, we examine a deviation in $\mathcal{D}_h(r'', F'') \cap \mathcal{D}_\ell^c(r'', F'')$. According to the universal divinity refinement, if neither type is more likely to deviate, then the interim belief ψ should coincide with the ex-ante belief π. Again, from lemma 3, the efficient type would want to deviate (but not the inefficient type) if the principal's interim belief is $\psi \in A = \left[0, \bar{\psi}\right]$. If the principal's interim belief is in a neighborhood of π, i.e., $\psi \in B = [\underline{\pi}, \bar{\pi}]$, then the inefficient type prefers to deviate but not the efficient. This happens because the efficient type, by construction of the $\mathcal{D}_h(r, F)$ and $\mathcal{D}_\ell(r, F)$ sets, will move to a lower indifference curve if he deviates to $(r, F) \in \mathcal{D}_h(r'', F'') \cap \mathcal{D}_\ell^c(r'', F'')$ and $\psi = \pi$. By continuity, the efficient type would become worse off also in a neighborhood of π. Moreover, $\bar{\psi} < \bar{\pi}$, because by construction of the set A, $\bar{\psi} < \pi$. To see this suppose, by way of contradiction, that $\bar{\psi} \geq \pi$. Take a $\psi = \pi \in A \cap B$. Since $\psi \in A$, this implies that the efficient type finds such a deviation profitable even if his offer does not get accepted. On the other hand, since $\psi \in B$, this implies that the efficient type will become worse off if he deviates and his offer gets accepted, contradiction. Therefore, $A \nsubseteq B$ and $B \nsubseteq A$ and by the universal divinity refinement the interim belief must coincide with the prior.

Next, we derive the equilibrium price $\bar{F}(1|\psi = \pi)$ of the agent's PD. When $r = 1$, $q_h^{SB}(\sigma_\ell) = 0$, $q_h^{SB}(1 - \sigma_h) = q_h^*$ and the price of the PD can be expressed as follows,

$$\bar{F}(1|\psi = \pi) = (1 - \pi)\left[S(q_h^*) - \theta_h q_h^*\right] - (1 - \pi)\left[S\left(q_h^{SB}(\pi)\right) - \theta_h q_h^{SB}(\pi)\right] + \pi \Delta\theta q_h^{SB}(\pi).$$

The efficient type becomes better off when PD can be sold relative to when such an option is not available if and only if,

$$\bar{F}(1|\psi = \pi) \geq \Delta\theta q_h^{SB}(\pi) \Leftrightarrow (1 - \pi)\left[\underbrace{\left[S(q_h^*) - \theta_h q_h^*\right] - \left[S\left(q_h^{SB}(\pi)\right) - \theta_h q_h^{SB}(\pi)\right]}_{A} - \underbrace{\Delta\theta q_h^{SB}(\pi)}_{B}\right] \geq 0.$$

The term A is zero when $\pi = 0$ (because $q_h^{SB}(0) = q_h^*$) and is strictly increasing in π (because the surplus from the inefficient type $S\left(q_h^{SB}(\pi)\right) - \theta_h q_h^{SB}(\pi)$ decreases as π increases). The term B is strictly positive when $\pi = 0$, it becomes zero when $\pi = 1$ and it is strictly decreasing in π. Thus, there exists a $\tilde{\pi}$ such that $\forall \pi \leq \tilde{\pi}$, $\bar{F}(1) \leq \Delta\theta q_h^{SB}(\pi)$ and $\forall \pi \geq \tilde{\pi}$, $\bar{F}(1) \geq \Delta\theta q_h^{SB}(\pi)$. This indicates that the efficient type is better off when he can sell his PD if and only if the ex-ante probability of the low type (efficient) is high enough. ■

References

[1] Banks, J.S. and J. Sobel (1987) "Equilibrium selection in signaling games," *Econometrica* 55, 647-661.

[2] Baron, D. and D. Besanko (1984) "Regulation, asymmetric information and auditing," *RAND Journal of Economics* 15, 447-470.

[3] Boyer, M. and J.J. Laffont (2003) "Competition and the reform of incentive schemes in the regulated sector," *Journal of Public Economics* 87, 2369-2396.

[4] Cho, I-K. and D. Kreps (1987) "Signaling games and stable equilibria," *Quarterly Journal of Economics* 102, 179-222.

[5] Crémer, J. and F. Khalil (1992a) "Gathering information before signing a contract," *American Economic Review* 82, 566-578.

[6] Crémer, J. and F. Khalil (1992b) "Gathering information before the contract is offered," *European Economic Review* 38, 675-682.

[7] Crémer, J., F. Khalil and J.C. Rochet (1998a) "Strategic information gathering before a contract is offered," *Journal of Economic Theory* 81, 163-200.

[8] Crémer, J., F. Khalil and J.C. Rochet (1998b) "Contracts and productive information gathering," *Games and Economic Behavior* 25, 174-193.

[9] Deighton, J. (2002) "Market solutions to privacy problems?," working paper Harvard Business School.

[10] Finkle, A. (2005) "Relying on information acquired by a principal," *International Journal of Industrial Organization* 23, 263-278.

[11] Fudenberg, D. and J. Tirole (1991) *Game Theory*, The MIT Press.

[12] Gal-Or, E. (1985) "Information sharing in oligopoly," *Econometrica* 53, 329-344.

[13] Glycopantis, D., A. Muir and N.C. Yannelis (2003) "On extensive form implementation of contracts in differential information economies," *Economic Theory* 21, 495-521.

[14] Ishiguro, S. (2003) "Comparing allocations under asymmetric information: Coase theorem revisited," *Economics Letters* 80, 67-71.

[15] Kahn, C.M. and T. Tsoulouhas (1999) "Strategic transmission of information and short-term commitment," *Economic Theory* 14, 131-153.

[16] Khalil, F. (1997) “Auditing without commitment,” *RAND Journal of Economics* 28, 629-640.

[17] Laffont, J.J. and D. Martimort (2002) *The theory of incentives: The principal-agent model*, Princeton.

[18] Liu, Q. and K. Serfes (2006) “Customer information sharing among rival firms,” *European Economic Review* 50, 1571-1600.

[19] Maskin, E. and J. Tirole (1990) “The principal-agent relationship with an informed principal, I: Private values,” *Econometrica* 58, 379-410.

[20] Maskin, E. and J. Tirole (1992) “The principal-agent relationship with an informed principal, II: Common values,” *Econometrica* 60, 1-42.

[21] Mezzetti, C. and T. Tsoulouhas (2000) “Gathering information before signing a contract with a privately informed principal,” *International Journal of Industrial Organization* 18, 667-689.

[22] Myerson, R. (1983) “Mechanism design by an informed principal,” *Econometrica* 51, 1767-1798.

[23] Riordan, M. and D. Sappington (1988) “Optimal contracts with public ex-post information,” *Journal of Economic Theory* 45, 189-199.

[24] Shaffer, G. and Z.J. Zhang (2002) “Competitive one-to-one promotions,” *Management Science* 48, 1143-1160.

[25] Shapiro, C. (1986) “Exchange of cost information in oligopoly,” *The Review of Economic Studies* 53, 433-446.

[26] Villas-Boas, J.M. (1994) “Sleeping with the enemy: Should competitors share the same advertising agency?,” *Marketing Science* 13, 190-202.

[27] Vives, X. (1990) “Trade association disclosure rules, incentives to share information, and welfare,” *RAND Journal of Economics* 21, 409-430.