Efficiency, Welfare and Ownership of Private Information*

Qihong Liu† Konstantinos Serfes‡

March 5, 2009

Abstract

Unrestricted flows of information usually improve efficiency. The recent growth of the Internet as a medium of communication and commerce, combined with the development of sophisticated software tools have paved the road for 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. We investigate whether consumers can claim some of the value of their own private information, while at the same time efficient flows of information are guaranteed. We address this question in a principal-agent adverse selection model. Prior to the contracting stage, the agent (consumer) chooses how much (precision) of his private information to sell to the principal. This gives rise to a signaling game that precedes the adverse selection stage. We show that there exists a pooling *efficient* equilibrium, where both agent types sell all their information to the principal.

Keywords: Information sharing; Asymmetric information; Signaling game.

*We would like to thank Luís Cabral, Minghua Chen, Yutian Chen, Kevin Grier, Mike Hoy, Vibhas Madan, Theofanis Tsoulouhas, Yianis Varoufakis, Nikos Vettas, Randal Watson and seminar participants at the 2008 International Industrial Organization Conference, the Spring 2006 Midwest Theory meeting, the 2006 Far Eastern Meeting of the Econometric Society (FEMES), the University of Oklahoma, the University of North Dakota, the University of Athens and the Athens University of Economics and Business for helpful comments, discussions and suggestions. Qihong Liu acknowledges a faculty fellowship from the University of Oklahoma. 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

Unrestricted flows of information usually improve efficiency. The recent growth of the Internet as a medium of communication and commerce, combined with the development of sophisticated software tools have paved the road for 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, micromarketing and one-to-one marketing [e.g., Shaffer and Zhang (2002)]. Moreover, consumers can play an integral role in the generation of these data through their actions in the marketplace. However, consumers may have little incentives to facilitate the collection and exchange of information, unless they receive some compensation in return.[1]

Supermarkets offer shoppers discounts conditional on joining frequent shopper programs, e.g., Deighton (2002). Essentially, shoppers sell private information about their preferences to retailers. Moreover, shoppers have a choice about how much (and what kind of) information to give up through, for example, the selected use of their preferred cards. Retailers and/or manufacturers can use this information to tailor their offerings to individual consumers. Will strategic consumers find it in their best interest to use their preferred cards?

Automobile insurance companies have started offering innovative contracts which monitor the behavior of the drivers, Filipova and Welzel (2005) and Filipova (2007).[2] A monitoring device (black box) which is installed into the driver's car records mileage, speed, acceleration, braking and distance to other traffic, among other things. The driver chooses the *precision* of the monitoring and at the end of the billing period, he decides whether to upload the data to the insurance company. In other words, the driver makes a decision about how much of his private information (if any) to give up. The insurance company promises discounts in the next period. On top of the moral hazard problem, there is also an adverse selection problem. Even after each driver has exerted maximum effort, there is still inherent risk heterogeneity among drivers. Indeed, the "black box" contains

[1] 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)]. This literature argues that personal information is an asset and consumers should get compensated when companies use their personal information. According to the November 21, 2002 Business Week article "Wanna See My Personal Data? Pay Up," 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.

[2] See also, "For a price, would you let car insurer along for the ride?" *USA TODAY*, August 3, 2005.

information about the type (i.e., inherent risk probability) of the driver. Although such a practice is not widespread among drivers and insurance companies yet, the improvement of facilitating technologies and the desire of insurance companies to better assess risk and tailor their offers will make it more popular in the future. A question that arises is: Will a strategic driver have enough incentives to use the "black box" and if so how much of the information contained in the "black box" will he/she choose to upload?

This paper makes an attempt to formally address some of the issues that were raised in the above motivating examples. We ask the following question: Can consumers/agents claim some of the value of their own private information, while at the same time efficient flows of information are guaranteed? This is an important question in our information technology era. We are interested in two specific aspects of the problem: i) economic efficiency and ii) benefits to consumers/agents.

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. 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 (nonlinear pricing). The outcome is not efficient because the principal optimally distorts his offer to the "bad" type (low demand or inefficient) agent. This may mean, for example, that product quality is below its first best level.[3] At the same time the "good" type (high demand or efficient agent) receives a surplus.

We extend this standard adverse selection 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, past behavior, 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. 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.

The agent decides about how much of his data to put up for sale and at what price. Following the agent's action, the principal makes his offer to the agent. Because the bad type receives zero surplus (in the absence of any information exchanges) he will have an incentive to sell all his information at the highest possible price. The good type on the other hand will never want to sell all his information, even if he gets compensated (up to the maximum the principal is willing to pay). This creates an interesting tension. If signaling was not an issue, then the inefficient type

[3]For instance, the principal 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.

would sell all his PD and the efficient type would sell either none or only a fraction of his PD. But signaling emerges naturally in this context and cannot be ignored. In particular, the agent's action 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.[4] Thus, signaling in our context may be detrimental to the agent, regardless of his type. As a consequence, the agent may choose to hold on to his PD, although such an information exchange may have improved efficiency.[5] We offer a discussion about how our modeling assumptions fit with the motivating examples we presented earlier in section 4.

We show that a separating equilibrium does not exist. In a separating equilibrium, the agent's offer reveals his type and the principal has no incentive to purchase the agent's PD. The contract the principal designs leaves the agent with zero surplus. Hence, it is not difficult to see that in this case the efficient type always wants to mimic the inefficient type. We then search for a pooling equilibrium. We focus on the efficient outcome, where the agent, regardless of his type, sells all his PD. We show that the efficient pooling perfect Bayesian equilibrium exists (with off-the-equilibrium beliefs that satisfy the Banks and Sobel (1987) universal divinity refinement). The principal purchases the agent's PD, draws a perfect signal and offers to the agent the *efficient* (first-best) contract. The intuition about why neither agent type wishes to deviate is as follows. A deviation entails an offer that will generate a noisy signal. The efficient type is 'more likely' to deviate than the inefficient. Roughly speaking, this is because the efficient type's surplus from the contract, but *not* the inefficient's surplus, is strictly positive when the signal he sells is less than perfect. Hence, according to the Banks and Sobel universal divinity refinement, the principal should attach probability one that such a deviation has come from the efficient type. Given this belief, neither type wishes to deviate. We also demonstrate that under reasonable assumptions on the off-the-equilibrium beliefs the efficient equilibrium is unique.

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

[4] We assume that the firm (principal) knows how to target the consumer (agent) even if the firm does not purchase the consumer's private data. For example, the firm knows the consumer's e-mail or address of residence.

[5] In the context of "cars with black boxes" example we introduced earlier, each driver, after the black box has collected a certain amount of information, can decide whether to upload these data to the insurance company (the driver can also control the precision (quality) of the information that is collected). This action can signal the driver's type.

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.

Our framework can also fit other similar contexts. For example, consider a principal, such as the Department of Defense, who 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.

The rest of the paper is organized as follows. The next section contains a literature review. We lay out the model in section 3. Section 4, contains the analysis and our main result (Theorem 1). We conclude in section 5. The Appendix contains most of the proofs.

2 Literature review

Our paper is related to the following three distinct strands of literature: i) information sharing and ii) principal-agent models with endogenous information structures and iii) distribution of property rights and efficiency.

Information sharing: 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.[6] 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. Kim and Choi (2007) study oligopolistic firms' incentives to share information about whether a particular consumer views the products as complements or substitutes. This distinction provides a new and important way of examining the information sharing problem. Nevertheless, they maintain the assumption that consumers are passive. Instead, the present paper assumes that consumers are active and fully rational players in the information exchange process.

Bouckaert and Degryse (2006) also allow consumers to have a say in the information exchange process. More specifically, they consider three different privacy policies: i) "opt-out", ii) "opt-in"

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

and iii) "anonymity." Under the first regime a firm that wants to share information with third parties must give consumers the opportunity to deny them permission. Under the second regime firms must obtain a consumer's explicit consent before personal information is shared. These first two policies allocate the property rights to consumers and they differ, in terms of how much information is shared, only if consumers find it costly to exercise the option. Otherwise, the two regimes are equivalent. In the absence of any costs, if consumers want to protect their privacy they will exercise their option under the opt-out policy but will not do so under the opt-in policy. If consumers want to release their information they will do the opposite. The third regime allows no information sharing whatsoever.

Principal-agent with endogenous information structures: 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.

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 and 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.[7] 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.

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.

Distribution of property rights and efficiency: Coase theorem states that if transaction costs are absent (or low), then private bargaining will lead to an efficient outcome, Coase (1960). But what can we say about efficiency in markets with asymmetric information? Ishiguro (2003) finds conditions under which the Coase theorem holds in a bilateral trade model with asymmetric information. Our result contributes to this literature by demonstrating that agents will not have any incentives to withhold their private information from the principal. Hence, efficiency is achieved regardless of who has the property rights to the agents' private data.

[7] 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).

3 The description of the model

A principal wants to delegate to an agent the production of q units of a good.[8] 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 $V = S(q) - t_i$.

We assume that the agent has the property rights to its own private data (PD) and can sell any subset of the PD to the principal.[9] 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 $r < 1$.[10]

[8] 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.

[9] We assume that the information seller cannot manipulate the information contained in the PD, nor can he sell selective fractions of the data, e.g., "good behavior" vs. "bad behavior."

[10] The upper bound of r does not have to be 1. It can be less than 1, if, for example, even if the agent sells all his PD the signal will not be perfect. The analysis will not change if we relax this assumption.

We assume that the database with the agent's information has already been developed and the only issue is how and how much of this information will be transmitted to the principal. This rules out the possibility that the agent may try to influence the type and quality of the data through his actions in the market place before information exchanges take place.

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 when the signal is s_i, $i = \ell, h$.

The game we will study can be described as follows (see also figure 1 where we present the sequence of events).

- 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 offers to the principal his PD of quality r at a price $F \geq 0$. The agent 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.

We denote by $\psi : [0,1] \times \left[\frac{1}{2}, 1\right] \times \mathbb{R}_+ \to [0,1]$ the *interim* belief held by the principal in stage 2 that the agent is efficient (π in stage 2 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 $\mu : [0,1] \times \{s_h, s_\ell\} \to [0,1]$ held by the principal in stage 3 that the agent is efficient as follows,

$$\mu = \begin{cases} \sigma_\ell (r|\psi), & \text{if } s = s_\ell \\ 1 - \sigma_h (r|\psi), & \text{if } s = s_h \\ \psi, & \text{if } s = \varnothing. \end{cases} \tag{3}$$

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

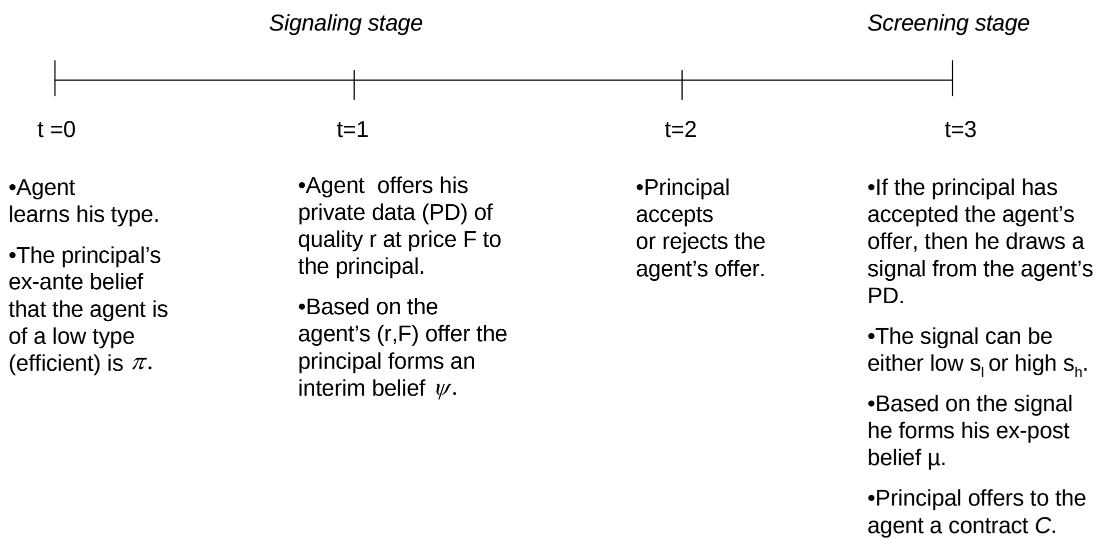

Figure 1: Sequence of events

2; otherwise (i.e., if the principal rejects the agent's offer, $s = \varnothing$), the ex-post belief coincides with the interim belief.

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)).[11]

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?

4 Analysis

We now provide a road map about how we solve the game (consult also figure 1). 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 4.2, we divide the (r, F) space into an acceptance and a rejection region. These regions are not fixed,

[11] 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-the-equilibrium paths. We restrict the off-equilibrium beliefs by insisting that they satisfy the Banks and Sobel (1987) universal divinity refinement.

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 4.1 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 impose on the off–the-equilibrium beliefs the universal divinity refinement. Given these beliefs, we show that there exists a unique pooling equilibrium where the agent succeeds in selling all his information to the principal.

Remark (On the principal's ability to make commitments). In the motivating examples we offered in the introduction, the principal (e.g., supermarket, insurance company) promises a reward when the agent sells his PD (e.g., uses his preferred shopping cards, uploads the black box). In our model there is a lack of commitment. The principal can refuse to buy the agent's PD if the agent's offer signals the type of the agent. We can reconcile the motivating examples with our modeling assumptions as follows. In reality, the principal interacts with the agent repeatedly and the principal needs many informative signals from the agent in order to learn the type of the agent with certainty. In each period, the principal promises the agent a reward in case the agent decides to sell his PD for that period. The principal, however, cannot commit to any future rewards. For example, a retailer promises certain shoppers a percentage discount when they use their preferred shopping cards over a given time period, but cannot make a credible promise over longer horizons. If the agent deviates from the candidate equilibrium offer, and this deviation reveals his true type, he still receives the promised reward for the current period but gives up any future rewards. So, signaling is costly because the principal cannot make long-run commitments. Our one period model captures this trade-off in a much simpler way.

4.1 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,
$$S'\left(q_h^{SB}\right) = \theta_h + \frac{\mu}{1-\mu}\Delta\theta. \tag{4}$$
- 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.[12] 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.

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

4.2 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 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 result is expected and its proof is straightforward and it is omitted. As r increases the principal gives up less information rent and therefore he is willing to pay a higher price for a more accurate signal. In figures 2 and 3 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.

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

The agent decides which offer (r, F) to make.[13] 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 $\bar{F}(r|\psi)$.[14] 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 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,

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

where $q_h^{SB}(\cdot)$ 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 $\Delta\theta q_h^{SB}(\sigma_\ell)$ and with probability $1-r$ he will draw a high signal in which case the efficient type's utility will be $\Delta\theta q_h^{SB}(1-\sigma_h)$.

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., $\frac{dF}{dr} = 0$.*

Proof. See Appendix. ■

Suppose, for example, that $S(q) = \sqrt{q}$. 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 2). On the other hand,

[13] 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.

[14] This assumption is not critical. We could have instead assumed that both players have bargaining power at this stage. An offer made by the agent would take the distribution of bargaining powers into account and would give some surplus to the principal. The results would not change qualitatively.

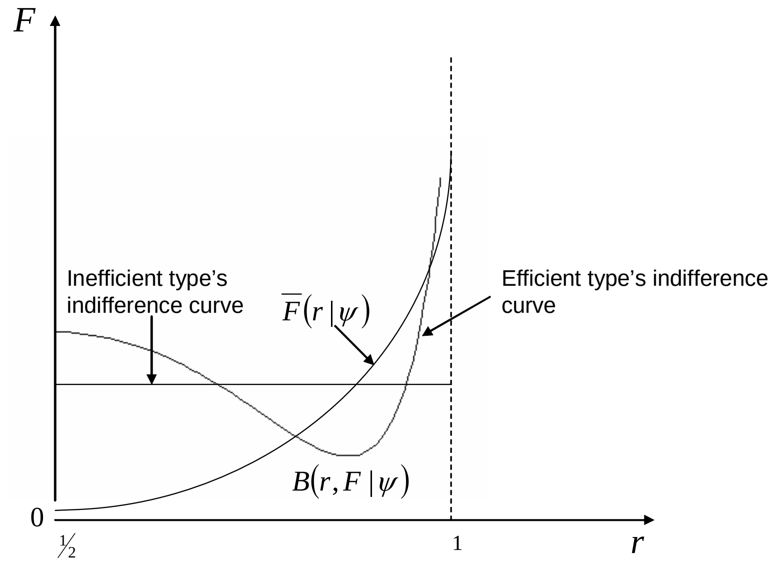

Figure 2: Indifference curves when ψ is high

if $\psi = .3$, then the indifference map is strictly increasing (see figure 3). (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,

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

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 $S\left(q\right) = \sqrt{q}$, for instance, the third derivative is indeed positive. Moreover, 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,[15]

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

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

[15] 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., $E\left[\Pr\left(\ell|r\right)\right] = \psi$.

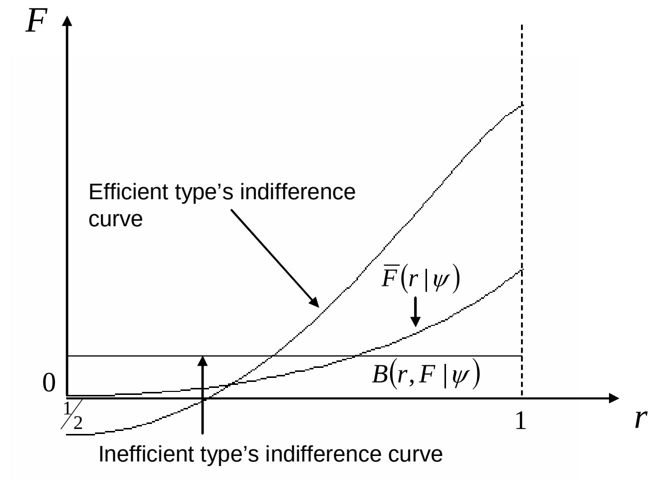

Figure 3: Indifference curves when ψ is low

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 $E[\Pr(\ell|r)]$. It can be easily shown that $E[\Pr(\ell|r)]$ is strictly increasing in r, with $E[\Pr(\ell|r)] = \psi$, at $r = \frac{1}{2}$ and $E[\Pr(\ell|r)] = 1$, 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 $\Pr(\ell|r)$ 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 $\Pr(\ell|r)$ 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 off–the-equilibrium beliefs.

4.4 Main results

First, we search 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 $\psi(r_\ell, F_\ell) = 1$. 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 $U_\ell = U_h = 0$.

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. ■

The main point is that the efficient type will always want to pool himself with the inefficient type. Next, we look for a pooling equilibrium, where both types offer (r', F'). The efficient outcome is when the agent sells all his PD, i.e., $r = 1$. Below, we state that this outcome is an equilibrium.

Theorem 1 (Efficient PBE). *The following is a pooling equilibrium with beliefs that satisfy the Banks and Sobel universal divinity refinement,*

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

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

- *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.[16]

[16] The above discussion, however, does not imply that the efficient type maximizes his expected utility even when

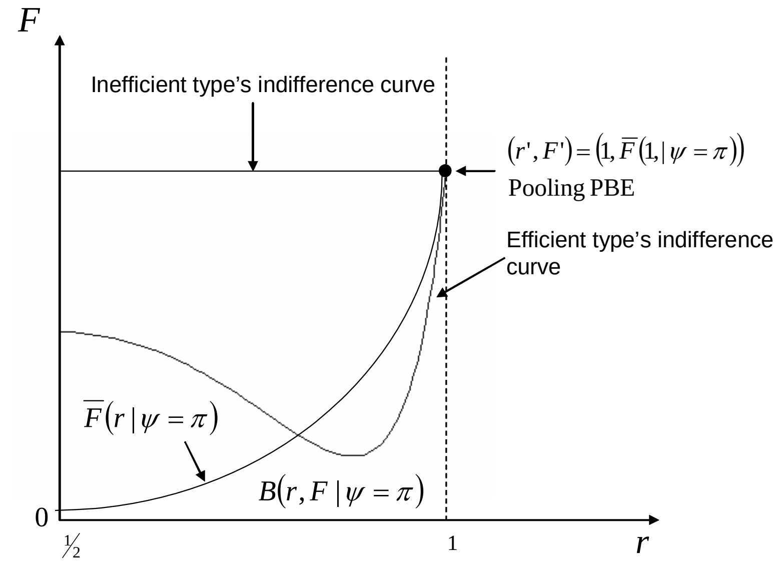

Figure 4: Pooling perfect Bayesian equilibrium

The beliefs that support the efficient pooling PBE satisfy the Banks and Sobel (1987) universal divinity refinement. The intuition is as follows. The efficient type is more likely to deviate from $(r', F') = \left(1, \bar{F}\left(1|\psi = \pi\right)\right)$ than the inefficient type. In other words, the set of beliefs under which the inefficient type has a profitable deviation is a subset of the set of beliefs under which the efficient type has a profitable deviation (see section A.3). Then, according to the universal divinity refinement, the principal must assign probability 1 to a deviation coming from the efficient type. (The Cho-Kreps (1987) intuitive criterion does not have a bite in our model because *both* types, under some beliefs, have a profitable deviation). Because $r' = 1$, both types receive zero surplus from the contract in stage 3, but only the efficient type, if he deviates to $r^{dev} < 1$, gets a surplus in stage 3. That is why, roughly speaking, the efficient type is more likely to deviate. But if the probability of an efficient type is 1, following a deviation, no type would want to deviate. The reason is that, given such an accurate belief, the principal does not need to purchase the information. So both types loose the price F and the efficient type receives no surplus from the contract in stage

$\pi \geq \tilde{\pi}$. Theorem 1 says that the efficient type is better off at the pooling PBE with information selling than without information selling. But it is possible that the efficient type would have enjoyed an even higher expected utility if he could sell 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 4 depicts the equilibrium when π is high (if π is low, then the efficient type's indifference map is increasing, but the equilibrium is not affected; see Lemma 2 and figure 3). From figure 4 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 (given that the principal's belief is 1), making both types worse off relative to the candidate equilibrium offer.

Uniqueness. Now consider an inefficient pooling offer, $r'' < 1$. At that offer, the principal's interim belief ψ (the probability of an efficient type) is equal to the ex-ante belief π, i.e., $\psi = \pi$. Following a deviation, the interim belief ψ does not have to be equal to π. We would like to show under what conditions, regarding the off-the-equilibrium belief ψ, such an offer is not an equilibrium. Then, we will end up with a *unique* PBE, the efficient one. (More details can be found in section A.3). The idea is as follows. We examine a specific deviation from $r'' < 1$ to $r' = 1$. We will demonstrate that the sets of beliefs of the principal that justify this deviation for each agent type are not nested. Therefore, according to the universal divinity refinement, neither type is 'more likely' to deviate and hence neither type can be eliminated. We then assume that the principal is agnostic and does not update his ex-ante belief, following a deviation from any $r'' < 1$ to $r' = 1$. Because the beliefs are not updated the inefficient type would deviate. This eliminates all the inefficient equilibria and we obtain a unique (efficient) PBE.

Let's see now why the two off-the-equilibrium belief sets cannot be nested. At $r'' < 1$ the efficient type receives a strictly positive surplus from the contract in stage 3, while at $r' = 1$ his surplus is zero. The inefficient type always receives zero surplus in stage 3. When ψ is very high, both types are going to be worse off because the price the principal is willing to pay, given his very accurate belief, is very low. But for a relatively high ψ (but not too high), the inefficient type is willing to deviate but not the efficient. The reason is that the efficient type wants a higher compensation than the inefficient type in order to offer a perfect signal and give up his strictly positive surplus. So, there always exists a region of F's at $r' = 1$, where an offer with $r' = 1$ and an F' in that region makes the inefficient type better off and the efficient worse off, relative to the initial inefficient pooling offer, $r'' < 1$. This logic does not apply when the candidate equilibrium is the efficient one, $r' = 1$, because the efficient type's surplus in stage 3 is already zero. On the other hand, for a very low ψ, it is the efficient type who wants to deviate but not the inefficient. Very low ψ means that the principal attaches a very high probability that the offer has come from the inefficient type. So, the contract in stage 3 involves very little distortion which benefits the efficient type, even if the offer is rejected by the principal (see section A.2). Because the offer is rejected the inefficient type, who receives zero surplus in stage 3 anyway, is worse off.

5 Conclusion

Consumers, through their actions in the market place, can generate databases with detailed information about their preferences and habits. This information is valuable to firms, but consumers, unless they get compensated, they may have little incentives to facilitate this process.

Our main purpose in this paper is to examine the incentives of consumers/agents to generate and sell information about their private type. 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 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 (but nonverifiable) 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 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.

When supermarkets offer shoppers discounts conditional on joining frequent shopper programs, they pay the shoppers for the asset, Deighton (2002). The more consumers use their preferred cards the more information about their preferences they give up, but at the same time the compensation they receive (in the form of discounts) rises. This positive relationship between the compensation consumers receive and how much of their private information they sell to retailers can be viewed as the $\bar{F}(r|\psi)$ curve in our model. Then, consumers choose how much information to sell to retailers through the selected use of their preferred cards. Our model predicts that strategic consumers will *not* find it in their best interest to withhold any information by refraining from using their shopping cards (or using them only a fraction of the time). Efficient flows of information will be observed in equilibrium.

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 topic. A number of assumptions made in this paper should be relaxed by future research. First, information gathering can be endogenized. In this paper, we have implicitly assumed that the signal cannot be manipulated by the agent (i.e., it is too costly for the agent to do so). If this assumption is relaxed, then agents will behave strategically knowing that their behavior will affect the type of data that will be generated. Second, one could introduce competition among multiple principals in a common agency framework. Finally, the number of agent types can be assumed to be more than two.

A Appendix: Proofs

A.1 Proof of Lemma 2

First, we characterize the map of indifference curves for the efficient type. By total differentiation of EU_ℓ, see (10), we obtain,

$$dEU_\ell = \left[\Delta\theta q_h^{SB}(\sigma_\ell) + r\Delta\theta \frac{dq_h^{SB}(\sigma_\ell)}{d\sigma_\ell}\frac{d\sigma_\ell}{dr} - \Delta\theta q_h^{SB}(1-\sigma_h) + (1-r)\Delta\theta \frac{dq_h^{SB}(1-\sigma_h)}{d\sigma_h}\frac{d\sigma_h}{dr} \right] dr + dF.$$

Now we use (8), where μ is replaced by σ_ℓ or $1-\sigma_h$,

$$\Rightarrow dEU_\ell = \left[\begin{array}{c} \Delta\theta q_h^{SB}(\sigma_\ell) + \frac{r(\Delta\theta)^2}{S''\left(q_h^{SB}(\sigma_\ell)\right)(1-\sigma_\ell)^2}\frac{d\sigma_\ell}{dr} - \Delta\theta q_h^{SB}(1-\sigma_h) - \\ (1-r)\frac{(\Delta\theta)^2}{S''\left(q_h^{SB}(1-\sigma_h)\right)\sigma_h^2}\frac{d\sigma_h}{dr} \end{array} \right] dr + dF.$$

Next, we use $\frac{d\sigma_\ell}{dr} = \frac{\psi(1-\psi)}{(1-r-\psi+2\psi r)^2}$, $\frac{d\sigma_h}{dr} = \frac{\psi(1-\psi)}{(r+\psi-2\psi r)^2}$, and we set $dEU_\ell = 0$ to obtain,

$$\Rightarrow \frac{dF}{dr} = -\left[\begin{array}{c} \underbrace{\Delta\theta\left[q_h^{SB}(\sigma_\ell) - q_h^{SB}(1-\sigma_h)\right]}_{(-)} + \\ \frac{\psi(\Delta\theta)^2}{(1-\psi)}\left(\underbrace{\overbrace{\frac{r}{S''\left(q_h^{SB}(\sigma_\ell)\right)(1-r)^2}}^{(-)} - \overbrace{\frac{1-r}{S''\left(q_h^{SB}(1-\sigma_h)\right)r^2}}^{(+)}}_{(-,+)} \right) \end{array} \right]. \quad (11)$$

The first term on the r.h.s. of (11) is negative because $\sigma_\ell \geq 1-\sigma_h$ which further implies $q_h^{SB}(1-\sigma_h) \geq q_h^{SB}(\sigma_\ell)$. The second term could be either positive or negative. When $r = \frac{1}{2}$, then $\sigma_\ell = 1-\sigma_h$, which implies that $S''\left(q_h^{SB}(\sigma_\ell)\right) = S''\left(q_h^{SB}(1-\sigma_h)\right)$ and the second term on the r.h.s. of (11) is zero. Hence, $\frac{dF}{dr} = 0$, since the first term on the r.h.s. of (11) is also zero at $r = \frac{1}{2}$. When $r = 1$, the second term on the r.h.s. of (11) is strictly negative (since $\frac{1-r}{S''\left(q_h^{SB}(1-\sigma_h)\right)r^2} = 0$ at $r = 1$, $q_h^{SB}(1) = 0$ and $q_h^{SB}(0) = q_h^* > 0$), which implies (by continuity) that $\frac{dF}{dr} > 0$ for r's in a neighborhood of 1. For every ψ, we can find a $\hat{r}(\psi)$ such that $\frac{dF}{dr} > 0$, $\forall r > \hat{r}(\psi)$. Then, $\bar{r} = \arg\max_\psi \hat{r}(\psi)$.

When $\psi = 0$, then $q_h^{SB}(\sigma_\ell) = q_h^{SB}(1-\sigma_h) = q_h^* > 0$ and therefore the second term on the r.h.s. of (11) is zero for all $r \in \left[\frac{1}{2}, 1\right]$. For each $r \in \left[\frac{1}{2}, 1\right]$ find the lowest ψ denoted by $\psi(r)$ such that $\frac{dF}{dr} > 0$. Such a ψ exists since the terms $S''\left(q_h^{SB}(\sigma_\ell)\right)$ and $S''\left(q_h^{SB}(1-\sigma_h)\right)$ are finite for low ψ's (because $q_h^{SB} > 0$ and increasing as ψ decreases) and consequently the term $\left(\frac{r}{S''\left(q_h^{SB}(\sigma_\ell)\right)(1-r)^2} - \frac{1-r}{S''\left(q_h^{SB}(1-\sigma_h)\right)r^2}\right)$ is also finite for low ψ's. Then, $\underline{\psi} = \arg\min_r \psi(r)$. Thus $\frac{dF}{dr} \geq 0$ for all $r \in \left[\frac{1}{2}, 1\right]$ provided that $\psi \leq \underline{\psi}$.

Finally, it is obvious that $\frac{dF}{dr} = 0$ for the inefficient type, since $EU_h = F$. ■

A.2 Lemma 3

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

Lemma 3. *Consider a pooling offer* (r, F) *on* $\bar{F}(r|\psi = \pi)$*, where* $r < 1$*. There exists a set of beliefs* $A = \left[0, \bar{\psi}\right]$*, such that if* $\psi \in A$*, then the efficient type has a profitable deviation, but not the inefficient type. The principal will reject the 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 $\psi\left(r', F'\right)$ 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.

The inefficient type will find such a deviation unprofitable, since, given that the offer will be rejected, $F = 0$.

Proof of Lemma 3. Suppose that $\psi\left(r', F'\right) = 0$. 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., $\bar{F}\left(r|\psi = \pi\right)$. At this price the principal is indifferent between accepting and rejecting the agent's offer and the price can be expressed as follows,

$$\begin{aligned}
\bar{F}(r|\psi=\pi) \;&\equiv\; \Pr(s_\ell|r)\,EV(\sigma_\ell)+[1-\Pr(s_\ell|r)]\,EV(1-\sigma_h)- \\
&\quad \left[\pi\left[S(q_\ell^*)-\theta_\ell q_\ell^*\right]+(1-\pi)\left[S\left(q_h^{SB}\right)-\theta_h q_h^{SB}\right]-\pi\Delta\theta q_h^{SB}(\pi)\right] \\
&= [\pi r+(1-\pi)(1-r)]\begin{bmatrix}\sigma_\ell\left[S(q_\ell^*)-\theta_\ell q_\ell^*\right]+\\ (1-\sigma_\ell)\left[S\left(q_h^{SB}(\sigma_\ell)\right)-\theta_h q_h^{SB}(\sigma_\ell)\right]-\\ \sigma_\ell\Delta\theta q_h^{SB}(\sigma_\ell)\end{bmatrix}+ \\
&\quad [1-\pi r-(1-\pi)(1-r)]\begin{bmatrix}(1-\sigma_h)\left[S(q_\ell^*)-\theta_\ell q_\ell^*\right]+\\ \sigma_h[S\left(q_h^{SB}(1-\sigma_h)\right)-\theta_h q_h^{SB}(1-\sigma_h)]-\\ (1-\sigma_h)\Delta\theta q_h^{SB}(1-\sigma_h)\end{bmatrix}- \\
&\quad \begin{bmatrix}\pi\left[S(q_\ell^*)-\theta_\ell q_\ell^*\right]+\\ (1-\pi)\left[S\left(q_h^{SB}\right)-\theta_h q_h^{SB}\right]-\pi\Delta\theta q_h^{SB}(\pi)\end{bmatrix} \\
&= (1-r-\pi+r\pi)\left[S\left(q_h^{SB}(\sigma_\ell)\right)-\theta_h q_h^{SB}(\sigma_\ell)\right]-r\pi\Delta\theta q_h^{SB}(\sigma_\ell)+ \\
&\quad r(1-\pi)\left[S\left(q_h^{SB}(1-\sigma_h)\right)-\theta_h q_h^{SB}(1-\sigma_h)\right]-\pi(1-r)\Delta\theta q_h^{SB}(1-\sigma_h)- \\
&\quad (1-\pi)\left[S\left(q_h^{SB}(\pi)\right)-\theta_h q_h^{SB}(\pi)\right]+\pi\Delta\theta q_h^{SB}(\pi).
\end{aligned}$$

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

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

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

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

The term in the first curly brackets is negative. Let's see why. First, we assume that the benefit function $S(\cdot)$ is sufficiently concave in q, so that the function $S\left(q_h^{SB}(\cdot)\right)-\theta_h q_h^{SB}(\cdot)$ is concave in the probability that the agent is of a high type (inefficient).[17] Second, the probability that the

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

agent is of a high 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.3 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 off-the-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 5. For instance, when we consider a deviation $\left(r^{dev},F^{dev}\right)\in\mathcal{D}_h(r',F')$, we allow for $\psi\neq\pi$.

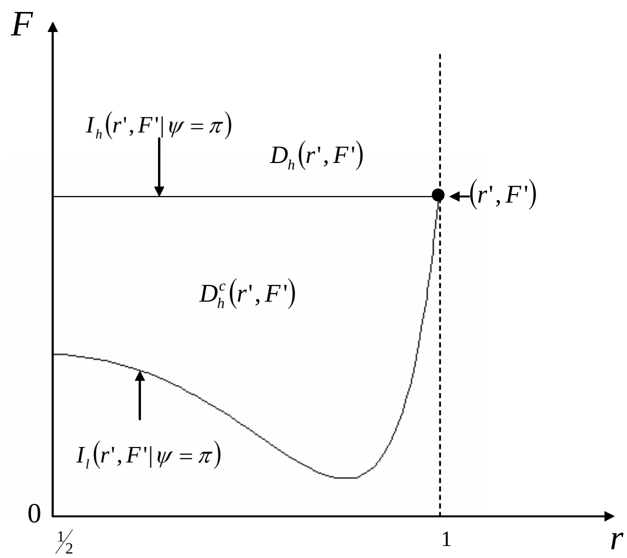

Figure 5: Various sets in the (r, F) region starting from (r', F')

To increase the predictive power of signaling games, restrictions must be made on the off-the-equilibrium beliefs. To this end, define the following belief function (see also figure 5),

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

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 how we obtain uniqueness.

Existence. First, we show that given (12), 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_ℓ^* and $t_\ell^* = \theta_\ell q_\ell^*$. 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 (12), 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. (The Cho-Kreps intuitive criterion does not apply because both types have a profitable deviation, under certain beliefs).

Uniqueness. We show that under some reasonable assumptions $(r', F') = \left(1, \bar{F}(1|\psi = \pi)\right)$ is the unique pooling universally divine PBE. By way of contradiction, consider a pooling offer (r'', F'') on $\bar{F}(r|\psi = \pi)$, where $r'' < 1$. Consider a deviation to $r' = 1$, see figure 6. We will show that neither type is 'more likely' to deviate to $r' = 1$, in the sense that the sets of the off-the-equilibrium beliefs of the principal that can justify such a deviation on part of each type of agent cannot be nested. More specifically, we will find some off-the-equilibrium beliefs where the efficient type has a profitable deviation but not the inefficient and some other off-the-equilibrium beliefs where the opposite is true. Then, the universal divinity refinement cannot eliminate either type. Since neither type is more likely to deviate, we will assume that the principal is agnostic and does not change his ex-ante belief, i.e., $\psi(r' = 1, F') = \pi$, following a deviation to $r' = 1$. Given this, the inefficient type has a profitable deviation from $r'' < 1$ to $r' = 1$.

Consider a candidate pooling equilibrium offer (r'', F'') and the three indifference curves that pass through it, assuming that the off-the-equilibrium belief ψ is equal to the ex-ante belief π: i) the efficient type's indifference curve $I_\ell(r'', F''|\psi = \pi)$, ii) the inefficient type's indifference curve

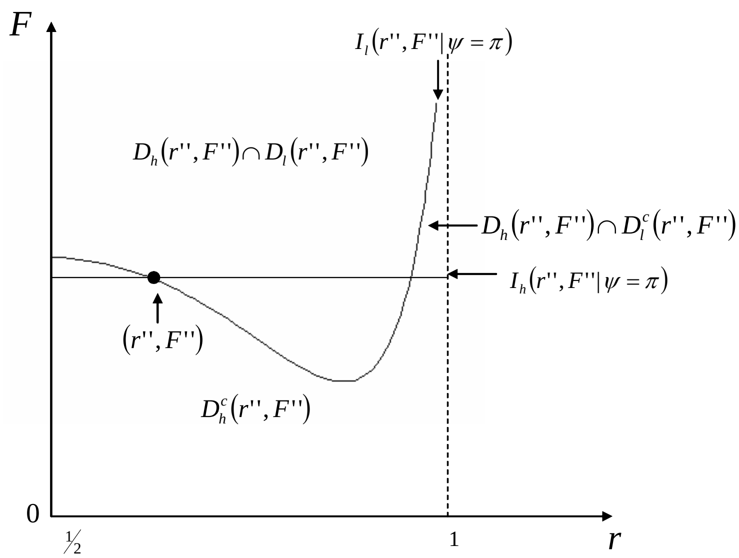

Figure 6: Various sets in the (r, F) region starting from (r'', F'')

$I_h(r'', F''|\psi = \pi)$ and iii) the principal's indifference curve $\bar{F}(r|\psi = \pi)$. Figure 6 depicts the first two indifference curves. Now we examine a deviation to $r' = 1$ on the principal's indifference curve, allowing for the off-the-equilibrium belief ψ to be different from π. The principal's indifference curve is increasing (see Lemma 1) in r, holding ψ fixed, and therefore the inefficient type is always better off at $r' = 1$, relative to $r'' < 1$, if $\psi = \pi$. For the efficient type this is not necessarily true, as at $r' = 1$ his surplus from the contract in stage 3 is zero. There are two cases: either the efficient type is worse off or he is better off at $r' = 1$, relative to $r'' < 1$, when $\psi = \pi$. The first case is when the F on the indifference curve of the efficient type at $r' = 1$ is above $\bar{F}(r' = 1|\psi = \pi)$, the maximum the principal is willing to pay for information of perfect quality when $\psi = \pi$. By continuity this is also true in the neighborhood of π, i.e., when $\psi \in B = [\underline{\pi}, \bar{\pi}]$). The second case is when the indifference curve of the efficient type at $r' = 1$ is below $\bar{F}(r' = 1|\psi = \pi)$.

First, assume that the efficient type is worse off at $r' = 1$. Hence, if the principal's off-the-equilibrium belief ψ is in a neighborhood of π, i.e., $\psi \in B = [\underline{\pi}, \bar{\pi}]$, then the inefficient type prefers to deviate, from $r'' < 1$ to $r' = 1$, but not the efficient. Moreover, $\bar{\psi} < \bar{\pi}$, because by construction of the set A, $\bar{\psi} < \pi$.[18] 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, neither type is more likely to deviate, i.e., $A \nsubseteq B$ and $B \nsubseteq A$. As we mentioned above, we assume that $\psi = \pi$.

[18]The set A was derived in Lemma 3, see section A.2. The efficient type, but not the inefficient, would want to deviate if the principal's interim belief is $\psi \in A = [0, \bar{\psi}]$.

Given this, the inefficient type will deviate and $r'' < 1$ is not an equilibrium.

Second, assume that the efficient type is better off at $r' = 1$, when $\psi = \pi$. Now, we increase ψ above π, so that we make the efficient type worse off, while the inefficient remains better off. An increase of ψ affects the principal's willingness to pay for a perfect signal. In particular, the $\bar{F}(r' = 1|\psi > \pi)$ falls below $\bar{F}(r' = 1|\psi = \pi)$, after ψ exceeds a threshold $\tilde{\psi}$. This is because the principal's willingness to pay is low when his belief is very accurate. In the extreme, the principal's willingness to pay is zero when $\psi = 1$. We also want ψ to be lower than a threshold $\hat{\psi}$, so that an offer with $r' = 1$ and F on the principal's indifference curve makes the inefficient type better off. Hence, for $\psi \in B = \left[\tilde{\psi}, \hat{\psi}\right]$, the efficient type is worse off and the inefficient type better off following a deviation to $r' = 1$ on $\bar{F}(r' = 1|\psi)$. This case is equivalent to the case above with the difference that the set B is now $B = \left[\tilde{\psi}, \hat{\psi}\right]$. Since $\tilde{\psi} > \pi$, the sets A and B are not nested and therefore the inefficient type would want to deviate. Given this, $r'' < 1$ is not an equilibrium.

Finally, we show that the efficient type may be worse off when the option of information selling becomes available. 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] Bouckaert, J. and H. Degryse (2006) "Opt in versus opt out: A free-entry analysis of privacy policies," TILEC Discussion paper, DP 2006-024.

[4] 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.

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

[6] Coase, R.H. (1960) "The problem of social cost," *The Journal of Law and Economics* 3, 1-44.

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

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

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

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

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

[12] Filipova, L. and P. Welzel. (2005) "Reducing Asymmetric Information in Insurance Markets: Cars with Black Boxes," working paper, Institute für Volkswirtschaftslehre.

[13] Filipova, L. (2007) "Monitoring and Privacy in Automobile Insurance Markets with Moral Hazard," working paper, Institute für Volkswirtschaftslehre.

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

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

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

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

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

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

[20] Kim, B-C. and J-P Choi (2007) "Customer information sharing: Strategic incentives and new implications," working paper, Michigan State University.

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

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

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

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

[25] 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.

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

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

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

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

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

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