Collusion and Data Privacy

Bob Mungamuru
Stanford University
bobji@stanford.edu

Michael Padilla
Stanford University
mtp@stanford.edu

Hector Garcia-Molina
Stanford University
hector@cs.stanford.edu

April 19, 2007

Abstract

In certain situations where a firm has an interest in privacy, external parties can be in a position to collude with each other in a manner that violates the firm's privacy. These external parties often have a monetary incentive to collude against the firm. The threat of collusion is discussed within the context of a simple economic model. Although a simple reward scheme can be used to deter collusion, a deterrent based on a long-term business relationship is more effective. An optimal collusion-resistant mechanism is designed with which the firm can filter out parties with low discount factors.

1 Introduction

The possibility of collusion between external parties can sometimes pose a significant threat to our privacy interests. Consider the following two motivating examples.

Example 1 – Finance. Most large financial institutions have both an investment banking arm (A) and an equity trading arm (B). For various regulatory reasons generally aimed at preventing conflicts of interest and insider trading (e.g., the US SEC [13]), at most institutions a "Chinese Wall" exists between their banking and trading practices. That is, information sharing between A and B is prohibited. As such, both the bank and the regulators (F) want to prevent such collusion. On the other hand, collusion can be very profitable for both traders and bankers, who are typically compensated based on commissions and performance bonuses. How, then, can the bank and the regulatory bodies set up incentives so that collusion does not occur?

Example 2 – Supply Chain. In typical supply chains, a single supplier (A) will have multiple customers (say, B and F) who compete in the same industry (e.g., Tata supplying both Ford and GM). Orders often have large lead times, so customers B and F submit orders to A well in advance

of the required delivery date, according to their respective demand forecasts [14]. However, since B and F compete in the same industry, knowledge of F's demand forecasts would be quite valuable to B. Supplier A can therefore collude with B against F, by selling F's demand forecast data to B. Is it possible for F to deter such a collusion between A and B?

These two motivating examples have certain features in common. Firstly, we have a firm (F, in the examples) with an interest in keeping certain information private. Secondly, there are two external parties (A and B, in the examples) who have both the means and monetary incentive to collude with each other. Finally, it is in the firm's interest to prevent this collusion from occurring. In this paper, we will construct an economic model of this general scenario wherein collusion between external parties poses a threat to a firm's privacy interests. We suggest simple incentive-based solutions whereby the firm can actually deter such collusion.

We will analyze the following simplified model, which is intended to be an abstraction of the many real-world scenarios in which collusion may occur: Consider a *firm* that has some sensitive data whose storage has been outsourced to two *service providers* A and B. Each of A and B hosts only a subset of the firm's data. Neither of the subsets hosted by A or B is sensitive on its own. However, if A and B were to combine these subsets, it would constitute a violation of the firm's privacy, causing considerable harm to the firm. We choose to analyze this firm-provider model (fully described in Section 2) because it neatly captures the various incentives and collusive behaviours that can arise in real-world scenarios.

1.1 More Examples

Aside from the examples discussed above, there are many other scenarios captured by the firm-provider abstraction that arise both in research contexts and in practice. For example:

- Web-based email providers (e.g., Google) and online shopping sites (e.g., Amazon) would be able to generate highly targeted advertisements by jointly mining a user's email and shopping habits. However, such a collusion between the email provider and the shopping site would be an violation of the user's privacy, which a privacy-conscious user would wish to prevent.

- Privacy policy languages are an active area of research (e.g., [5]). Policy languages define rules and logic for which data can be shared between parties in a sensitive workflow (e.g., patient records in a hospital). However, the standard underlying assumption that these parties will not communicate (and collude) outside the context of this workflow is unrealistic.

- In early 2005, Apple filed a lawsuit against ThinkSecret.com (a website devoted to rumors about

Apple products), alleging that the site leaked proprietary trade secrets illegally obtained through an unnamed Apple employee [4]. Apple's suit was unsuccessful.

- In 2006, a Coca-Cola employee was turned in to authorities by Pepsi, when the employee attempted to sell Coca-Cola's trade secrets to Pepsi for a sum of $1.5 million [7].

While these and other scenarios can be modeled using the firm-provider abstraction, the relative values of the model parameters in each context are quite different. As a result, the behaviour we observe in each scenario can vary greatly.

1.2 Contributions

To summarize, we make the following contributions:

- An economic model is constructed, based on a firm and two service providers, which captures the threat of collusion as it relates to privacy (Sections 2 and 3).
- It is shown that collusion can be deterred in this model, in theory, by setting up a simple reward scheme for "honest" providers (Section 4).
- The same effect can be achieved with a much smaller reward, by instead using the provider's future revenue stream as a deterrent. Moreover, it is sufficient to ensure just one of the providers is "honest". However, this scheme relies on knowledge of the provider's private information (Section 5).
- An optimal, incentive-compatible, collusion-resistant mechanism is derived, where knowledge of provider's type is no longer required (Section 6).

2 Model

As described in the previous Section, our model is based on the canonical example of a *firm* F that over a series of periods wishes, for each single period, to outsource its data storage needs to two *service providers*, A and B[1]. We denote by I the original data, which is split into two shares I_A and I_B, hosted respectively by A and B. Each of the two shares has no value by itself.

[1]The firm-and-two-providers abstraction has garnered some recent research interest from computer scientists interested in providing "databases as a service" (e.g., [2, 11]). Simply encrypting the firm's data and hosting it at a single provider is not feasible, since query performance would suffer too greatly.

We assume that A and B are aware of each others' existence in this business relationship, and know an inverse function f that may be used to reconstruct I from both I_A and I_B i.e., $I = f(I_A, I_B)$. This is in contrast to the typical assumption in security research that the providers are unaware of each others' presence, obviating any concern for possible interactions between them. In exchange for providing their services, A and B receive from the firm payments of y_A and y_B per period, respectively. In providing the data hosting service to the firm, the two providers incur marginal costs of c_A and c_B. Finally, it is assumed that the firm and providers are all risk-neutral and have zero reservation utility[2].

The basic setup is illustrated in Figure 1. Observe that while, in general, the information I will be different for each time period t, this variability will not matter as we are assuming that the other economic parameters are the same for each period, and it is they that will determine the providers' behaviour. In addition, we make the assumption that the data set I, different for each time period, only has value to the agents and potential third parties during that specific period i.e., after one period the data becomes "stale" and is deemed worthless.

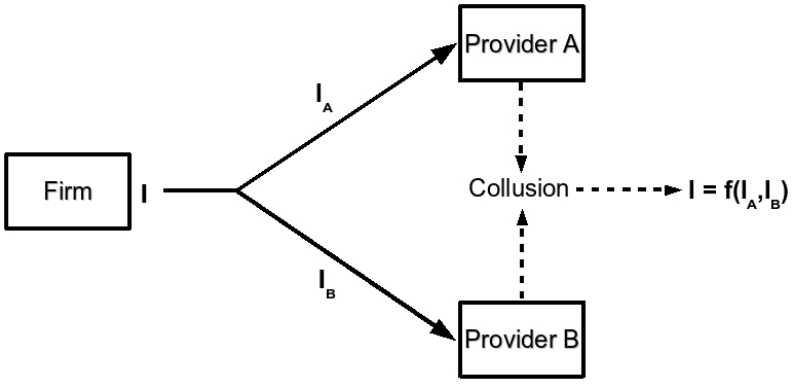

Figure 1: The threat of collusion in data outsourcing.

We assume that the firm's sole objective is to maintain the privacy of its proprietary data I, and not have providers A and B attempt to reconstruct the data (later, in Section 6, we ask whether it is worth the firm's effort to deter collusion). This objective is fundamentally due to a plethora of reasons such as the data being confidential (patients' medical records, financial account information, etc.), representing proprietary product specifications and/or business plans, etc.

We assume that the data is such that, were the providers A and B to collude against the firm and reconstruct the data I, then they would each be able to receive payoffs of v_A and v_B, respectively. For

[2]In this context, reservation utility refers to the minimum utility that would be necessary for an agent to engage in the proposed business transaction.

example, the data may be in some form directly valuable to the providers themselves. Alternatively, we might imagine that the data could be provided to a third party, in return for payments v_A and v_B, perhaps as equal shares of a payment v ($v_A = v_B = \frac{v}{2}$). The firm's goal, therefore, is to interact with the two providers according to a contract (mechanism) that will work to deter potential collusion between them.

There is an interesting difference between the two examples presented in Section 1. In the finance example, the firm F (the bank) has a business relationship with both providers A and B (investment banking and trading divisions). On the other hand, in the supply chain case, the firm F (the customer) only has business dealings with A (the supplier), but not B (the other customer). That is, $y_B = 0$ in the supply chain example. In Section 5, we will demonstrate that it is sufficient to deter just one of the providers from wishing to collude, meaning that collusion can be deterred even in the supply chain example.

2.1 Discount Factor and Reputation

The goal of each provider i is to maximize the sum of all future discounted payoffs, i.e., to maximize

$$\text{NPV}_i = \sum_{t=0}^{\infty} \delta_i^t \pi_{t,i},$$

where NPV stands for "net present value", $\pi_{t,i}$ is the payoff to provider i in period t, and δ_i is the discount factor. $\pi_{t,i}$ will be explained in more detail in the Sections to follow, but is a result of the providers' decisions to collude or not collude and depend on v_i, y_i, δ_i, etc. In many cases, δ_i is thought of as the time value of money if invested in a risk-free account. If the risk-free interest rate is r per period, then in this case one might set $\delta_i = \frac{1}{1+r}$. However, in a more general context, δ_i can in addition represent agent-specific time preferences and probabilities of continuation. In general, δ_i will be specific to each provider, and will influence the relative value it places on payoffs sooner as opposed to those later and hence will, in general, affect which strategies are optimal. In our model, we have a game of incomplete information in which each agent enters the game only knowing their own δ (i.e., their type). It is also assumed (to be discussed in more detail later) that agents' δ's are drawn from a distribution $p(\delta)$, which is common knowledge. As our analysis develops, it will be seen that the providers' different values for δ play a pivotal role in the development of our mechanism.

When considering whether to collude or not, providers need also consider the impact that collusion could have on their business relationships with other firms, if discovered. Such a discovery and subsequent disclosure would suggest to other firms in the industry that the colluding partner is untrustworthy. While a provider may gain if they decided to collude in the context of one business

relationship, the fear of having this discovered and their overall reputation soiled will in many cases act to discourage collusion. We will model this "fear" by assuming that each agent in the game attaches to their reputation a value $z(s_i, \delta_i)$, where s_i denotes the "size" of the provider, perhaps measured by its market capitalization if it is publicly traded, or its yearly revenue[3]. If a company has a large s_i, then it can be thought that it is fairly well-established and hence has worked hard to develop a reputation that represents a lot of its value and hence $\frac{\partial z}{\partial s_i} > 0$. Similarly, if a firm has a large δ_i, then it weighs future payoffs highly and thus places a large value on its reputation, as this is a large component in its future business opportunities. Hence, it similarly follows that $\frac{\partial z}{\partial \delta_i} > 0$. To simplify the notation, we will supress the dependence on s_i and δ_i, defining $z_i \equiv z(s_i, \delta_i)$. In the course of the agents' interactions, to be described in more detail to follow, if a given agent is found to have colluded (which will occur with probability p_1, the probability of detection to be discussed more later) then they will be considered to have lost this reputation value z.

3 Base Case

We begin by considering the base case of a game played over a single period. This simplified setting will introduce the flow of the game in our model and illustrate some of the basic ideas that will be further developed in the dynamic case to follow. It will be easy to see, in the base case, that some providers will indeed have an incentive to collude against the firm.

The order of play in the base case is as follows:

1. The firm pays y_A and y_B to A and B and hands I_A to A and I_B to B.
2. The providers host the firm's data[4], incurring marginal costs c_A and c_B.
3. Each provider decides, separately, whether or not to collude with the other provider (we call this the *collusion game*). The outcome of this game is observable by both providers, but not by the firm.
4. Collusion is said to occur only if both providers decide to collude, i.e. cooperate by sharing their data to profit from it.

If collusion does not occur, the game ends at this point. However, if collusion does occur, provider i receives v_i ($i \in \{A, B\}$), and the game continues:

[3]It is assumed that the function $z(\cdot, \cdot)$ is the same across providers in an industry.

[4]Is is assumed throughout that when providers are indifferent to providing their service and not providing their service, that they will opt to do so.

		Provider B	
		D	C
Provider A	D	$y_A - c_A, y_B - c_B$	$y_A - c_A, y_B - c_B$
	C	$y_A - c_A, y_B - c_B$	$y_A - c_A + \delta_A(v_A - p_1 p_2 K - p_1 z_A),\ y_B - c_B + \ldots$

Table 1: Payoff matrix for base case.

5. With probability p_1, the firm detects the providers' collusion[5].

6. If the collusion is detected, the firm is able (and willing) to prove it in court with probability p_2.

7. If collusion is proven in court, the firm extracts a penalty of K from each provider.

8. If collusion is detected, each provider's reputation also suffers; the value lost is z_i, $i \in \{A, B\}$.

The resulting normal form game is given in Table 1, which shows the expected payoffs to the two providers. Here C stands for "collude" and D stands for "don't collude". Unless (C, C) is played, the payoffs are $(y_A - c_A, y_B - c_B)$ since no collusion occurs and no further revenues are generated beyond the firm's payment to the providers. However, if (C, C) is played, collusion occurs and the payoff to provider i is $y_i - c_i + \delta_i(v_i - p_1 p_2 K - p_1 z_i)$. Hence, if $v_i > p_1(p_2 K + z_i)\ \forall i$, C weakly dominates D for both providers and (C, C) is a dominant strategy equilibrium – that is, collusion will always occur. Observe that in the base case, the firm will offer to pay each firm $y_i = c_i$, to cover their marginal cost.

It is important to note that, in many realistic scenarios, K will be of the same order as v, and the probabilities of both detection and successful litigation are very small, $p_1, p_2 << 1$. This latter point is based on the fact that these probabilities are based on an overall population of such scenarios, and that while in some cases colluders may be careless and easily detected and prosecuted, in most cases the complexity of corporate IT infrastructures and the sophistication of the providers make these events unlikely. Thus, the payoff to provider i is approximately $y_i - c_i + \delta_i(v_i - p_1 z_i)$ and it is seen that the value of z_i, the value provider i places on their reputation, will be critical to determining the behavior of the providers. Specifically, note that $y_i - c_i$ will be received regardless of what strategy is decided and hence the term $\delta_i(v_i - p_1 z_i)$ will dictate which strategy will be played. If $z_i >> v_i$ such that $p_1 z_i > v_i$, then the possible profit from collusion is not worth the possible loss of reputation. One such example is Coke/Pepsi case mentioned in Section 1, where Pepsi informed both the authorities and Coca-Cola about having been approached with valuable inside information by a Coca-Cola employee. In contrast, if $p_1 z_i < v_i$, then reputation alone is not enough to deter collusion, and the provider is willing to risk damage to their industry reputation.

[5]We assume no false positives.

		Provider B	
		D	C
Provider A	D	$y_A - c_A, y_B - c_B$	$y_A - c_A + \delta_A R, y_B - c_B - \delta_B z_B$
	C	$y_A - c_A - \delta_A z_A, y_B - c_B + \delta_B R$	$y_A - c_A + \delta_A(v_A - p_1 p_2 K - p_1 z_A),\ y_B - c_B + \ldots$

Table 2: Payoff matrix with a reward for snitching.

In light of the above discussion, and the possibility that collusion is a providers' optimal action, it is natural to ask: What incentives can the firm provide to deter collusion? We begin to address this issue in the next Section.

4 A Simple Reward Scheme

One approach the firm can take to deter collusion in the base case is to introduce a reward $R > 0$ to providers who, when presented with an offer to collude by the other firm, opt to remain "loyal" and report the attempted collusion to the firm. That is, if for example A wants to collude but B does not (i.e. (C, D) is played in the collusion game), the firm may reward B with R for snitching on A, and vice versa if (D, C) is played. In order for this to be feasible, we assume that both providers are able to credibly report the outcome of the collusion game to the firm; a firm that does not attempt collusion while the other provider does is able to provide appropriate convincing proof of this fact.

Employing this reward scheme modifies the payoff matrix to that of Table 2. We note that because $R > 0$, when given the opportunity to snitch, firms will always elect to do so. Now, if $R > v_B - p_1 p_2 K - p_1 z_B$, then B will prefer to play D whenever A plays C and similarly for A if $R > v_A - p_1 p_2 K - p_1 z_A$. In this case, (D, D) replaces (C, C) as the unique NE of the game and collusion will not occur. Therefore, we see that with a sufficiently large reward for loyalty, collusion can be avoided for all δ_i. This works even in the case where $y_B = 0$ – we can still reward B for snitching on A.

5 Long Term Contracts

In Section 3, an incentive problem arises because the expected penalty as a result of collusion, $p_1 p_2 K + p_1 z_i$, is simply not substantial enough to deter collusion between the providers. As an extreme example, suppose that $p_1 p_2 K + p_1 z_i \approx 0$ i.e., winning a battle in court is unlikely and the expected reputation value loss alone is not enough to deter collusion. The solution proposed in Section 4 was to offer a reward R to provider A for refusing B's offer to collude and snitching on B instead (vice versa

for B snitching on A). To deter collusion, therefore, the firm would have to offer A a reward $R > v_A - p_1 p_2 K - p_1 z_A \approx v_A$, which is roughly the value A would gain from collusion in the first place!

Thus, in many instances, using rewards alone to deter collusion is impractical. In cases where the reward is being offered to a party where no other business relationship exists (e.g., as in the supply chain example, with $y_B = 0$), it is still more impractical. Hence, we ask the question: Can we deter collusion using a smaller reward? The answer is "Yes", if we use long-term contracts.

Intuitively, rather than simply rewarding B for snitching on A, it would be more effective to punish A for attempting to collude with B as well (observe that, so far, we have leveraged neither y_A nor δ_A towards deterring collusion). The difficulty, of course, is how to actually "punish" A – it is difficult to sue a provider in court when collusion did not actually occur. However, the firm does have the recourse of depriving A of future revenues.

Suppose that the firm, instead, interacts with the providers repeatedly across multiple periods. For example, in the supply chain example, the customer might commit to ordering from the same supplier once per month, rather than using a different supplier each month. The sequence of events is as follows (the differences from the simple reward scheme are marked in **boldface**):

1. The firm pays y_A to A and y_B to B and hands I_1 to A and I_2 to B
2. The providers host the firm's data, incurring marginal costs c_A and c_B
3. At the end of the period, the providers play the collusion game – the outcome of this game is observable by the other provider, but not by the firm:
   (a) If both A and B wish to collude, then collusion occurs and provider i receives v_i
   (b) If A wishes to collude but B does not, then collusion doesn't occur and B decides whether to snitch on A.
   (c) If B wishes to collude but A does not, then collusion doesn't occur and A decides whether to snitch on B.
   (d) If neither A nor B wish to collude, then collusion doesn't occur **and play proceeds to the next period, with a new set of data**.
4. If A snitches on B, the firm pays A a reward R **and no longer deals with B in future periods (the firm finds a replacement for B)**; B suffers a loss in reputation valued at z_B; vice versa if B snitches on A.
5. If collusion occurs, the firm detects it with probability p_1.

6. If collusion is detected, the firm is able (and willing) to prove it in court with probability p_2, **and no longer deals with A or B in future periods (i.e., the firm replaces both A and B)**; providers A and B suffer reputation losses z_A and z_B, respectively.

7. If collusion is proven in court, the firm can extract a penalty of K from each provider.

8. **Play continues to the next period, with a new set of data.**

The resulting dynamic game[6] is represented in Figure 2. Without loss of generality, we have depicted A as making the first move in Figure 2. C stands for "collude", D stands for "don't collude", S stands for "snitch" and NS stands for "not snitch". The asterisk at each terminal indicates that the game restarts at the end of each period (possibly with new players A and B who may have been replaced because they were caught in the act of collusion). The payoffs are represented by the symbols $\pi_i^{j,k}$, which is the expected payoff to provider i when the outcome of the collusion game is j and the snitching decision is k (e.g., $\pi_B^{DC,S}$ is the payoff to provider B when A plays "don't collude", B plays "collude" and then A decides to snitch on B).

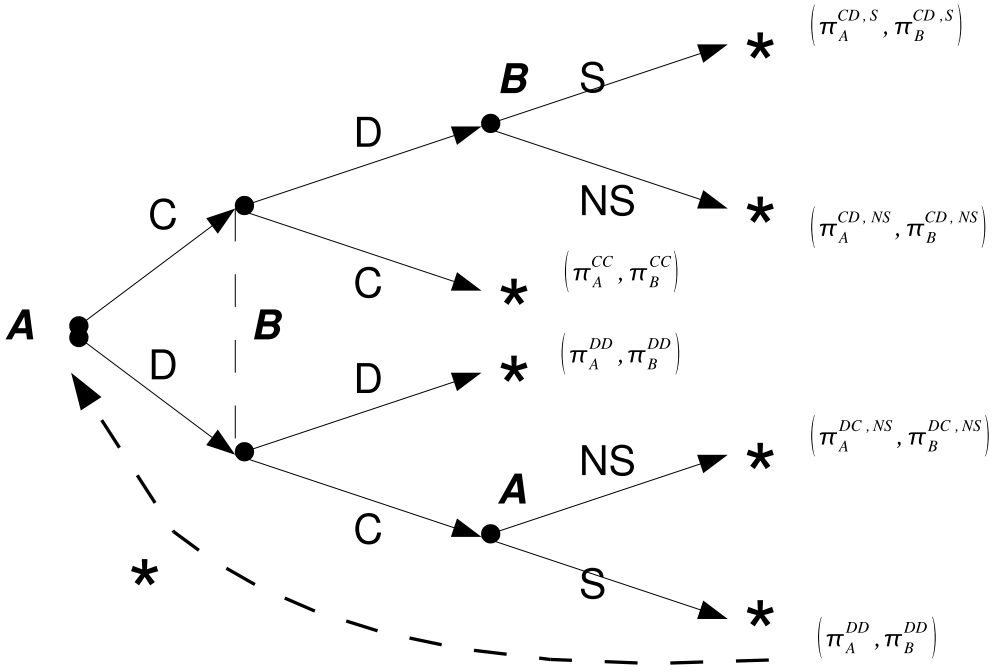

Figure 2: Game tree for long-term contract.

[6]We might have tried to model this repeated interaction using the theory of infinitely repeated games. However, the notion that a "dishonest" provider (i.e., one who attempts to collude) can be replaced by a new player is better captured in a (infinite) dynamic game framework.

From the sequence of events for the dynamic game, we have the following payoffs[7]:

$$\begin{aligned}
\pi_A^{CD,S} &= y_A - c_A - \delta_A z_A \\
\pi_B^{CD,S} &= y_B - c_B + \delta_B (R + G_B) \\
\pi_A^{CD,NS} &= y_A - c_A + \delta_A G_A \\
\pi_B^{CD,NS} &= y_B - c_B + \delta_B G_B \\
\pi_A^{CC} &= y_A - c_A + \delta_A (v_A - p_1 p_2 K - p_1 z_A + (1 - p_1) G_A) \\
\pi_B^{CC} &= y_B - c_B + \delta_B (v_B - p_1 p_2 K - p_1 z_B + (1 - p_1) G_B) \\
\pi_A^{DD} &= y_A - c_A + \delta_A G_A \\
\pi_B^{DD} &= y_B - c_B + \delta_B G_B \\
\pi_A^{DC,NS} &= y_A - c_A + \delta_A G_A \\
\pi_B^{DC,NS} &= y_B - c_B + \delta_B G_B \\
\pi_A^{DC,S} &= y_A - c_A + \delta_A (R + G_A) \\
\pi_B^{DC,S} &= y_B - c_B - \delta_B z_B
\end{aligned}$$

Here, G_A and G_B are the values to A and B, respectively, of continuing the game onto the next period. For example, observe that A does not receive the payoff G_A if B snitches on A (i.e., if C, D, and S are played). The following proposition suggests that, in this infinite dynamic game setting, collusion can be deterred using a much smaller reward.

Proposition 1. *Suppose* $R > 0$ *and either* $y_A > c_A + (1 - \delta_A)\left[\frac{v_A - R}{p_1} - p_2 K - z_A\right]$ *or* $y_B > c_B + (1 - \delta_B)\left[\frac{v_B - R}{p_1} - p_2 K - z_B\right]$. *Then, in the unique sub-game perfect equilibrium of the dynamic game,* (D, D) *will be played every period and collusion will never occur.*

See Appendix A for a proof for Proposition 1. The implication of Proposition 1 is that, if the firm enters into a long-term contract with the providers, as long as either y_A or y_B is sufficiently large, it only needs to offer a reward $R > 0$ to ensure that collusion never occurs[8]. The role of the reward here is mainly to ensure that in the event (D, C) or (C, D) is played, snitching will always occur. The prospect of receiving $y_i - c_i$ every period into the future is the primary factor that deters provider i from wanting to collude against the firm.

[7]Note that we might have considered a different representation of the dynamic game tree where the providers' decisions to collude are not made simultaneously. Instead, say, A would know B's decision before having to decide whether to collude. The conclusion that only one of the providers needs to be kept honest will also apply to this case.

[8]Here we assume that there is no "loyalty" between the providers. Technically, we also need R not too big to ensure that providers do not play (C, D) or (D, C), and then share the reward (this isn't much of a restriction)

Observe, from Proposition 1, that the firm must now pay an amount $y_A > c_A$ per period (i.e., a premium above and beyond the reservation utility) to ensure that collusion does not occur. This is in contrast to the base case and the simple reward scheme, where covering the providers marginal costs was sufficient i.e., $y_i = c_i$ (although collusion is not deterred in the former).

Another important conclusion from Proposition 1 is that it is sufficient to deter either one of A or B from colluding (i.e., by paying them a sufficient high y per period), but deterring both is unnecessary. That is, the firm can deter collusion by paying $y_A > c_A + (1 - \delta_A)\left[\frac{v_A - R}{p_1} - p_2 K - z_A\right]$ to provider A, and paying just the marginal cost $y_B = c_B$ to provider B. Intuitively, deterring just one provider is sufficient since the "honest" provider will always refuse to collude and snitch on the "dishonest" provider. Knowing this, the "dishonest" provider will never even offer to collude. The result that deterring just one provider is sufficient is very convenient. It allows us to deter collusion even in situations where the firm does not have a business relationship with one of the providers i.e., $y_B = 0$. The supply chain scenario is an example.

For the remainder of this Section and the next, we assume without loss of generality that the firm's goal is to deter provider A from wanting to collude with B. Provider B is always paid $y_B = c_B$. For a fixed y_A, define $\underline{\delta}(y_A)$ to be the value of δ such that $y_A = c_A + (1 - \delta)\left[\frac{v_A - R}{p_1} - p_2 K - z_A\right]$. Then, for a fixed y_A, we know that if provider A's type is $\delta_A > \underline{\delta}(y_A)$, then A will not be interested in colluding with B. The only remaining question, is how to ensure that $\delta_A > \underline{\delta}(y_A)$ i.e., how can a firm ensure that they are dealing with a provider whose discount factor is sufficiently high?

6 A Screening Mechanism

In Section 5, we concluded that the firm can deter collusion by entering into a long-term contract with provider A, and paying a sufficiently high amount y_A per period. However, we have not discussed how the firm enters into this contract in the first place. A plausible policy for the firm might be to first fix its budget y_A, and then find a sufficient "reputable" provider whose discount factor is $\delta > \underline{\delta}(y_A)$[9]. The problem, of course, is how to actually find out what provider A's discount factor is.

Thus far, we have assumed that δ_A is known. In reality, however, δ_A is private information known only to A. In fact, if A's discount factor was less than $\underline{\delta}(y_A)$, he would have incentive to lie about it, in order to get hired by the firm. The fact that δ_A is private information leads to some basic questions:

1. How can the firm get the provider to report δ truthfully?

[9]There are probably good reasons, other than collusion deterrence, for a firm to deal with providers who have a high δ. See the earlier discussion on δ and continuation probability

2. What does it mean for a provider to "report" his δ?

3. Is there an "optimal" policy for the firm?

4. Is it even worth it for the firm to try to deter collusion?

We can address the first two questions by designing an *incentive-compatible* contract. Recall from earlier that it is sufficient to deter collusion for any one of the providers we wish to deal with (say A). Suppose the firm is evaluating propsective provider A, and wishes to enter into a contract only if $\delta_A > \underline{\delta}(y_A)$. Incentive compatibility means that the firm offers a contract to A, which A would accept only if $\delta_A > \underline{\delta}(y_A)$. If, $\delta_A < \underline{\delta}(y_A)$, then it would be in A's interest to decline the offer. Thus, good providers are "self-selecting".

Consider what we call an *entry fee contract*, which is one of the simplest forms of incentive-compatible contracts. An entry fee contract involves an up-front lump sum payment e from the provider to the firm at the time the contract is entered, followed by payments of y_A per month by the firm to the provider A, for services rendered until the contract is terminated (e.g., because the provider goes out of business, or collusion is detected). It is easy to see that if the firm sets $e = \frac{y_A - c_A}{1-\delta'}$ for some δ' of its choosing, then only those providers with $\delta_A > \delta'$ will accept the contract[10].

Therefore, returning to the questions we posed earlier, a firm can use an entry fee contract to incentivize providers to "report" their discount factor truthfully. By "report", we mean that the provider signals whether $\delta_A > \delta'$ via his acceptance or rejection of the firm's contract offer. It remains, then, to determine what δ' the firm should choose in the entry fee contract.

The third and fourth questions posed above can be answered by designing a contract that is optimal for the firm. An optimal incentive-compatible contract for the firm will be one that maximizes its expected net present value (where the expectation is taken over A's type, which is unknown to the firm). However, we must first quantify how much the firm values its privacy[11]. Let V be the monetary loss to the firm if collusion between the providers is not prevented, and let δ_0 be the discount factor applied by the firm to future cashflows.

Proposition 2. *Fix y_A. Suppose $V \geq \frac{c_A}{1-\delta_0}$. Then, there exists a $\delta^*(y_A) > \underline{\delta}(y_A)$ such that an entry fee contract with $e = \frac{y_A - c_A}{1-\delta^*(y_A)}$ and a per-period payment of y_A maximizes the firm's net present value.*

[10] Assume $\delta_A < \delta'$ and that the provider's reservation utility is zero. Then, NPV $= e + \frac{y_A - c_A}{1-\delta_A} < e + \frac{y_A - c_A}{1-\delta'} = 0$, and so the provider will decline the contract.

[11] Firms such as Intel and IBM seem to value their privacy greatly (product secrecy, in this case), so they take measures such as imposing press embargoes on product releases. Apple's lawsuit against ThinkSecret seems to indicate the same. As suggested by studies such as [1], however, some users on the web do not value privacy as highly.

$\delta^*(y_A)$ is increasing in y. Only providers of type $\delta_A > \delta^*(y_A)$ accept the contract. If $V < \frac{c_A}{1-\delta_0}$, then the firm does not try to deter collusion.

A proof of Proposition 2 is provided in Appendix A. In the proof of Proposition 2, we assume the existence of a publicly-known probability density $p(\delta)$ over the provider's discount factor.

Proposition 2 yields several interesting insights on the optimal incentive-compatible contract. Firstly, since $\delta^*(y_A)$ is increasing in y_A, the higher the firm's budget, the more discriminating it is in selecting a provider i.e., "you get what you pay for". Secondly, collusion is deterred, since $\delta^*(y_A) > \underline{\delta}(y_A)$. Thirdly, for a given budget y_A, types $\delta \in (\underline{\delta}(y_A), \delta^*(y_A))$ decline the contract offered, even though they would not have colluded against the firm. This is an inefficiency the firm must accept in exchange for imposing incentive-compatibility.

The utility of a type $\delta^*(y_A)$ provider who accepts the contract is exactly zero. Since providers of all types are offered the same contract parameters, providers with $\delta > \delta^*(y_A)$ will have a strictly positive utility. The positive utility earned by these types is the *informational rent* extracted due to the fact that δ is private information. The optimal "cutoff" discount factor $\delta^*(y_A)$ expresses a balance between having the provider A accept the contract terms, and not having high-δ providers extract too much rent. Finally, the need for $V \geq \frac{c_A}{1-\delta_0}$ is easy to interpret – any provider that the firm deals with will require a payment of at least c_A per period, to recover marginal costs. If $V < \frac{c_A}{1-\delta_0}$, even this minimal investment is not worthwhile for the firm.

The "entry fee" in an entry fee contract could be any number of things. For example, it may be an up-front "deposit" that the provider pays the firm, which would be lost in the event of collusion being detected. It might also be a "due diligence" fee charged to the provider, in order for the firm to verify that the provider is trustworthy.

7 Discussion

7.1 Future Work

In this Section we review certain aspects of our model. Relaxations of some of our modeling assumptions may be useful directions for future work.

Providers in our current model are assumed to have no loyalty towards each other. We can extend our model to include the possibility that, say, provider B might retaliate against A for snitching or not-colluding. Alternatively, A may not snitch on B in exchange for B returning the favour in some other context. Strategic behavior across contracts with different firms may significantly alter a provider's actions. Similarly, groups of firms might band together to act in concert against a misbehaving

provider. Techniques from cooperative game theory may be more appropriate for modeling such strategic behavior between groups of firms or providers.

Another useful extension would be to extend the mechanism for negotiating a contract to multiple firms. Equivalently, we can include a non-zero reservation utility for the provider, to represent other opportunities. On the other hand, *firms* may collude with each other, against the providers, to obtain favorable contract terms. Analyses from the industrial organization and multi-agent mechanism design literature may be relevant here, specifically related to agents bidding for a contract.

We may wish to model the possibility that snitching is not necessarily credible. It may be interesting to capture the possibility of provider B falsely accusing provider A in spite of (D, D) being played, just to claim a reward – other mechanisms may need to be introduced to verify credibility.

Using a richer model of industry reputation (i.e., z) is likely to produce interesting results. For example, how do real firms quantify their reputation? What are the other primary factors that determine the value of reputation? We have made the simplifying assumption that z is lost entirely upon snitching. However, it is more realistic to model the erosion of reputation, and the actions a firm might take to restore it.

Although V, v_A and v_B are assumed to be known in our model, they may be better modeled as random variables. This distinction was not important in our model thus far, due to our assumption of risk neutrality. With risk averse agents, however, the actual distributions over V, v_A and v_B will affect each agents utility, and therefore must be taken into account. Moreover, we have modeled the repeated interaction between firm and provider as an infinite game. However, in some cases, a finitely repeated game may be more realistic.

7.2 Related Work

To the authors' knowledge, the general problem of collusion between agents leading to unwanted sharing of information has not been studied before. The threat of collusion modeled in this paper is an instance of security as an externality, as discussed in [3]. Collusion resistant cryptographic protocols for data communications do exist. A good example is [6]. However, the literature on cryptography and protocols typically addresses the secure transmission of data, whereas this paper is concerned with what happens after the data has been decrypted.

The topic of collusion has been well-studied in the economics literature, most commonly in the industrial organization literature and in auction theory. In IO, the central problem is organizations colluding to fix prices in a competitive market. A good survey of the topic can be found in [8]. In auction theory, the problem is somewhat the opposite – it is the buyers that pose the threat of

collusion. An approach to collusion-resistant auction design is [10].

An example of an approach to collusion deterrence similar in spirit to this paper is [12]. Here, the goal is to prevent collusion between an employee and a supervisor sent to monitor his progress – collusion, in this context, is the falsification of reports by the supervisor in exchange for a bribe from the employee. Although the problem setup is different, the solution technique is a prisoner's dilemma-style setup, similar to the snitching reward we propose here.

8 Conclusion

In this paper, we have studied collusion between external parties as a threat to a firm's privacy interests. While it is often in a party's interest to collude against the firm, we have shown that that the firm can actually set up incentives of its own to deter collusion from taking place.

Collusion is a problem in several other contexts beyond the examples we have presented in this paper. While our model captures a number of these contexts, extensions to our model may allow us to faithfully analyse a wider range of scenarios, such as when groups of providers or firms interact strategically amongst themselves.

References

[1] A. Acquisti and J. Grossklags. Losses, gains, and hyperbolic discounting: An experimental approach to information security attitudes and behaviors. In *2nd Workshop on the Economics of Information Security*, College Park, MD, USA, May 2003.

[2] G. Aggarwal, M. Bawa, P. Ganesan, H. Garcia-Molina, K. Kenthapadi, R. Motwani, U. Srivastava, D. Thomas, and Y. Xu. Two can keep a secret: A distributed architecture for secure database services. In *Proceedings of 2nd Biennial Conference on Innovative Data Systems Research*, Asilomar, USA, January 2005.

[3] R. Anderson and T. Moore. The economics of information security. *Science*, 314:610–613, Oct 2006.

[4] Apple slams rumor site with lawsuit, 2005. http://www.eweek.com/article2/0,1895,1748605,00.asp.

[5] A. Barth, A. Datta, J. C. Mitchell, and H. Nissenbaum. Privacy and contextual integrity: Framework and applications. In *Proceedings of the 27th IEEE Symposium on Security and Privacy*, Oakland, USA, May 2006.

[6] D. Boneh and B. Waters. A collusion resistant broadcast, trace and revoke system. In *Proceedings of the 13th ACM Conference on Computer and Communications Security*, Alexandria, VA, USA, Nov 2006.

[7] Coke requests court secrecy, 2005. http://www.cnn.com/2006/LAW/07/06/coke.secrets/index.html.

[8] R. A. . S. H. (ed.). *Handbook of Game Theory with Economic Applications*, volume 3. Elsevier, 1 edition, 2002.

[9] D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, Cambridge, MA, USA, 1991.

[10] A. Goldberg and J. Hartline. Collusion-resistant mechanisms for single-parameter agents. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms*, Vancouver, BC, Canada, Jan 2005.

[11] H. Hacigumus, B. Iyer, C. Li, and S. Mehrotra. Executing sql over encrypted data in the database-service-provider model. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*, Madison, USA, June 2002.

[12] F. Kofman and J. Lawarree. A prisoner's dilemma model of collusion deterrence. *Journal of Public Economics*, 59(1):117–136, 1996.

[13] The laws that govern the securities industry, 2006. http://www.sec.gov/about/laws.shtml.

[14] H. Lee and S. Whang. Information sharing in a supply chain. *International Journal of Technology Management*, 20(3):373–387, 2000.

A Proofs

A.1 Proof of Proposition 1

Since $R > 0$, $\pi_A^{DC,S} - \pi_A^{DC,NS} = \delta_A R > 0$ and $\pi_B^{CD,S} - \pi_B^{CD,NS} = \delta_B R > 0$. Therefore, given the opportunity, both A and B will always choose S over NS. Moreover, $\pi_B^{DD} - \pi_B^{DC,S} = \delta_B(G_B + z_B) > 0$, so B would prefer to play D if A plays D. Therefore, if $\pi_B^{CD,S} > \pi_B^{CC}$, then D would be preferable to C at all nodes in B's information set. That is, for B, playing D would be a dominant strategy over playing C, and A would respond by playing D as well. Therefore (D, D) would be played at every period in any subgame perfect equilibrium.

For $\pi_B^{CD,S} > \pi_B^{CC}$ to hold, we need $\pi_B^{CD,S} - \pi_B^{CC} = \delta_B \left(R - (v_B - p_1 p_2 K - p_1 z_B - p_1 G_B)\right) > 0$. Noting that $G_B > \frac{y_B}{1-\delta_B}$, we get $y_B > c_B + (1 - \delta_B)\left[\frac{v_B - R}{p_1} - p_2 K - z_B\right]$ as a sufficient condition.

Observe that the dynamic game we have defined is symmetric in A and B. Therefore, by symmetry, we also get $y_A > c_A + (1-\delta_A)\left[\frac{v_A - R}{p_1} - p_2 K - z_A\right]$ as a sufficient condition, which yields the desired result.

A.2 Proof of Proposition 2

Our proof is similar to the analysis in Chapter 7.3.2 of [9]. We highlight the important details here.

The firm must decide whether to hire a provider of type δ. Let $x : [0,1] \to [0,1]$ represent the firm's hiring decision i.e., the firm hires type δ with probability $x(\delta)$.

The firm's utility is $u_0 \equiv u_0(x(\delta), \delta) = x \cdot \left(V - \frac{y}{1-\delta_0} + e\right)$ and the provider's utility is $u_A \equiv u_A(x(\delta), \delta) = x \cdot \left(\frac{y-c}{1-\delta} - e\right)$. The firm's goal is to choose $x(\delta)$ to maximize it's expected NPV, $E_\delta[u_0]$, subject to the incentive-compatibility constraint $u_A(x(\delta_A), \delta_A) \geq u_A(x(\delta), \delta_A)\ \forall \delta$. The expectation is taken over the provider's type δ, which is unknown to the firm, and assumed to have a probability density $p(\delta)$ and associated cumulative distribution $P(\delta)$.

Following the analysis of [9], we can rewrite our objective function to be the following:

$$\max_{x(\delta)} \int_0^1 x(\delta) \cdot \left[V - \frac{y}{1-\delta_0} + \frac{y-c}{1-\delta} - \frac{1-P(\delta)}{p(\delta)} \frac{y-c}{(1-\delta)^2}\right] d\delta \tag{1}$$

Define $W(\delta; y) \equiv V - \frac{y}{1-\delta_0} + \frac{y-c}{1-\delta} - \frac{1-P(\delta)}{p(\delta)} \frac{y-c}{(1-\delta)^2}$. In order to maximize the objective function, we simply set $x(\delta) = 1$ whenever $W(\delta; y) \geq 0$, and set $x(\delta) = 0$ whenever $W(\delta; y) < 0$.

We assume that the probability density $p(\delta)$ has a monotone hazard rate (examples of such densities include the uniform, normal, Laplace, and certain beta distributions). Under the monotone hazard rate assumption, $\frac{d}{d\delta}\left(\frac{1-P(\delta)}{p(\delta)}\right) \leq 0$. The monotone hazard rate assumption implies that $W(\delta; y)$ is increasing in δ. Clearly, $W(0; y) < 0$ and $W(1; y) > 0$. Let $\delta^*(y)$ be the value of δ such that $W(\delta; y) = 0$. The optimal hiring decision $x^*(\delta)$ is therefore: $x^*(\delta) = 0 \forall \delta > \delta^*(y)$, and $x^*(\delta) = 0$ otherwise.

Observe that $W(\delta^*(y); y) = 0 = \frac{1}{y-c} W(\delta^*(y); y) = \frac{1}{y-c}\left(V - \frac{y}{1-\delta_0}\right) + \frac{1}{1-\delta^*(y)}\left(1 - \frac{1-P(\delta^*(y))}{p(\delta^*(y))} \frac{1}{(1-\delta^*(y))}\right)$. The last expression is decreasing in y and decreasing in $\delta^*(y)$. Since $\delta^*(y)$ is defined such that $W(\delta^*(y); y) = 0$ for all y, we conclude that if y increases, $\delta^*(y)$ must also increase in order to maintain $W(\delta^*(y); y) = 0$. In other words, $\delta^*(y)$ must be increasing in y.

The per-period payment y must always be greater than the provider's marginal cost c, to ensure that the provider participates. At $y = c$, $W(\delta; c) = V - \frac{c}{1-\delta_0}$. Therefore, if $V < \frac{c}{1-\delta_0}$, the firm will simply set $x(\delta) \equiv 0$, and won't even try to deter collusion.