Conditional Privacy Rights*

Murat C. Mungan†

April 16, 2016

Abstract

People have subjective valuations of privacy. Thus, absent further considerations, efficiency requires that a person be afforded privacy if, and only if, his subjective valuation of privacy exceeds the social value of the information that would be disclosed through a violation of that person's privacy. Absolute regimes that either always allow privacy, or never allow privacy, cannot achieve this result. This article shows that a conditional privacy regime can lead to efficient separation among people based on their subjective valuations of privacy. Moreover, this regime does not inefficiently distort information collection incentives or incentives to refrain from various acts that may generate collectible information.

1. Introduction

Rapid technological developments are enabling cheap and effective means of information collection and dissemination. While this is generally a positive development that reduces the cost of learning, it generates some negative consequences by making available private, embarrassing, damaging and even humiliating information. Thus, an important question is whether these negative consequences can be eliminated or mitigated, without interfering much with the free flow of information.

A solution that was previously proposed is the creation of a "National Information Market" (NIM).[1] A well functioning NIM has the same appealing features as any other market: people engage in transactions voluntarily, thus, every transaction implies a Pareto improvement, and a NIM allows more voluntary transactions. Stated differently, a person would sell private information in a NIM only if he is offered a price that is greater than the inconvenience to him of this information being disseminated (henceforth, privacy cost). Moreover, a buyer would be willing to offer a price greater than the privacy cost to the seller, only if the social value of the information is greater than this privacy cost. Thus, all transactions would be wealth enhancing.

*I am grateful for the financial support provided by Law and Economics Center, George Mason University School of Law, as part of the 2015-16 Privacy Fellowship.

†D'Alemberte Professor of Law, Florida State University College of Law; Courtesy Professor of Economics, Florida State University Department of Economics. E-mail: mmungan@law.fsu.edu

[1] Laudon (1996).

There is, of course, an important problem that a NIM may not be able to solve. The buyer of the information will often be interested in purchasing the information to disseminate it to third parties, and collect proceeds from doing so. In many cases, the buyer will not be able to collect the entirety of the reservation prices of third parties from accessing this information (among other reasons, because he cannot perfectly price discriminate). Thus, even assuming that a NIM were feasible, it would not be able to facilitate all transactions that are wealth enhancing, but, every transaction that takes place in a NIM would be wealth enhancing. That the buyer of the information cannot internalize the entirety of the positive externalities he generates by disseminating the information is only one of the reasons for this, but, other considerations may imply the same result.

Based on the discussion above, one may claim that a NIM would unambiguously increase social welfare: they only allow welfare enhancing transactions. This view is correct, but only given the initial allocation of rights. If, for instance, people did not have any legal ownership interest in information about them, then more information would be disseminated, including those whose dissemination is welfare enhancing, but, due to problems discussed above, end up not being disseminated in a NIM. Thus, the normative ranking of a NIM, as described above, and a legal regime where people have no property rights (call this the 'publicity regime') attached to information about them, depends, among other things, on (i) our beliefs regarding the proportion of instances where a NIM versus a publicity regime prevents welfare enhancing transactions, and, (ii) the average value of the welfare enhancing transactions that each regime promotes and prohibits.

While this is inherently an empirical question whose answer depends on the particular circumstance, the normative values that are reflected through various legal and social norms, as well as the First Amendment, indicates that we, as a society, place a very high value on allowing free information flow. This can be interpreted as a result of an implicit belief that, in an overwhelming majority of cases, the social value of the information being shared is greater than the privacy cost to the individual. By saying this, I am not at all attempting to undervalue the interest of any individual to keep certain information to himself and to have it shared only with individuals of his own choosing. Rather, I am providing a simple economic rationale for our default position that we should not be interfering with the free flow of information in an overwhelming majority of circumstances.

These observations have important implications regarding the optimal allocation of property rights over private information. This claim is contrary to early comments on the issue, which claim that an application of the Coase theorem implies that the initial allocation of property rights is irrelevant.[2] This is because, there are important transaction costs associated with enforcing any conceivable protection of privacy rights (in addition to the costs of even defining these rights), which implies that allocating rights to the wrong party may

[2] Noam (1997).

cause at least two sources of significant inefficiencies. First, in instances where transactions costs are prohibitively high there will be losses from unrealized trades, i.e. allocative inefficiencies. Second, even in instances in which transaction costs are small enough to make trade between parties profitable, these small costs will nevertheless need to be incurred if rights are held by the wrong party. These costs, which are small per transaction, can sum up to very large aggregate costs, if the number of transactions that need to take place is large. Thus, if we believe that in most cases the social value of sharing the information is greater than cost of breaching privacy, it is presumably best to not allocate privacy unconditional rights (as suggested in the NIM proposal).

The obvious alternative, is to have a publicity regime that unconditionally allows information dissemination. The initial concern that comes to mind with this type of regime is that it may lead to the dissemination of undesirable (e.g. embarrassing or humiliating) information, which may generate net-losses. A counter-claim to this argument may be that people who are really significantly adversely affected by the presence of this type of information may approach the party disseminating it (henceforth, publishers), and pay for the non-dissemination of this information. Although this response may have some appeal, it does not consider real and important externalities.

First, this type of regime may have the effect of over-incentivizing the capturing of negative value information.[3] This type of behavior, motivated by the prospect of rent-seeking, causes at best distributive effects that do not enhance welfare ex-post, but requires investments ex-ante.[4] Hence, allowing no privacy protection may cause losses, too, by incentivizing wasteful investments in rent-seeking activity. The opposite problem is also possible: if publishers are able to internalize only a small proportion of the social gains from information dissemination, they may have sub-optimal incentives to collect socially valuable information. Which of these problems is generated hinges on a comparisons between the proportion of social value that the publisher is unable to capture (this can be called consumer surplus in the market for information) through information dissemination and the average privacy cost of individuals about whom information is disseminated. If the former [latter] is larger, then the publisher is under-incentivized [over-incentivized] to collect information.

Second, as Daughety and Reinganum (2010) point out, the prospect of information collection and dissemination may affect actors' ex-ante behavior. By not engaging in activity about which information can be collected, the actor can minimize the likelihood of embarrassing information being leaked out. Wether this is a good or a bad thing depends on the nature of the act that generates the information. If the act does not cause any significant harm onto others (and perhaps confers benefits to some parties), then, one may refer to this type of ex-ante impact as *chilling effects*. On the other hand, if the act is socially harmful, one may appropriately refer to them as deterrence effects. A publicity

[3] Hirshleifer (1971) makes a similar observation regarding investments in private information.

[4] These costs are quite simlar to those that exist in the context of duress in contract law and are analyzed in Seidenfeld and Mungan (2015).

regime can *over-chill* socially desirable behavior, and may over- or under-deter socially harmful behavior.

To summarize, giving people property rights over their private information may generate large transaction costs and may also lead to the over-removing of information from the public sphere. On the other hand, the publicity regime may over-incentivize publishers to invest in information acquisition activities, and may also have undesirable effects on the activities that generate the information that is disseminated.

In this article, I consider a hybrid-regime of 'conditional privacy rights' where the rule is that a publisher may disseminate any information that it captures, but an individual may remove this information from the public sphere conditional on paying an amount determined by a court. This enables a sorting of the information to be removed based on the person's subjective valuation of his privacy. This regime also has the potential of aligning the publisher's incentives to collect information with the social objective. In particular, if the type of information is one that, on average, generates privacy costs that are smaller than the fraction of the social value of information dissemination that the publisher is unable to reap, than it is optimal for part of the proceeds from the information removal to be redirected to the publisher. Alternatively, if the opposite condition holds, it may be optimal to impose a sanction on the publisher for each information removal transaction.

Thus, the advantages of this regime are that it enhances allocative efficiency by allowing individuals to remove information from the public sphere when the privacy cost exceeds the social value of the information. Moreover, the regime mitigates costs associated due to misalignments between the publishers incentives to collect information and the social objective. Finally, the regime can improve social welfare through its effects on actors' ex-ante behavior: chilling behavior is mitigated due to the possibility of removing embarrassing information, and when the act is socially harmful, gains from privacy can be optimally traded-off against deterrence benefits through adjustments in the price that one has to pay to exercise privacy rights.

To formalize these potential benefits in a simple framework, I consider a setting where individuals can remove information from the public sphere only through court orders, and that they cannot directly negotiate with publishers. These assumptions are justified if there are large transaction costs, perhaps due to the reasons described above, that prevent Coasean bargaining between the publisher and individuals. Thus, the model presented cannot formalize any transaction-cost-reducing-effects. However, an additional plauisble benefit from conditional privacy rights -relative to a NIM- is presumably obtained by allowing the parties to avoid transactions costs in an overwhelming majority of cases (where the social value of the information is greater than the private value to the individual). I briefly elaborate on transaction costs related issues in section 4.

It is important to note that the analysis in this article relates to only a sub-set of the concerns raised in the prior literature on privacy. Acquisti et al. (2016) provide an extensive survey of this literature, and conclude that there are

three main themes in this literature. Some previous articles focus on problems associated with behavioral biases which make it hard for people to correctly interpret the trade-offs associated with privacy related decisions. Another strand of the literature focuses on the lack of information regarding how, when, and why information may be collected. Finally, there articles focusing on the theoretical welfare effects of privacy protection. The instant article is related to this last strand of the literature. In particular, it focuses on the *simple meaning* of privacy costs which are generated by concrete pieces of information that are embarrassing or reveal negative traits about an individual.

It is also worth noting that the analysis presented here has a close resemblence to the seemingly urelated literature on contractual hold-up and duress. In particular, Shavell (2007), Bar-Gill and Ben-Shahar (2004), and Mungan and Seidenfeld (2015) provide analyses of cases where a party may expend resources to place itself in a position to threaten a second party unless that party complies with the threatenor's demands. The analyses in these articles differ from the analysis presented here, because they consider the relationship between two specific parties (as opposed to publishers and the public generally). Moreover, Shavell (2007), which, as in this article, considers court induced price caps in these contractual hold-up situations, considers a setting where, contrary to the setting considered here, all costs and benefits are incured by the two parties interacting, i.e. there are no externalities to third parties.

This article proceeds as follows. In section 2, I propose a simple model where a publisher chooses how much effort to exert in information collection, and, this decision does not affect people's ex-ante behavior (i.e. no chilling or deterrence effects are captured). This section highlights how sanctions and/or rewards can be used to align the publisher's incentives with the social objective. Section 3, extends this model to incorporate potential ex-ante effects, and shows that the same mechanism proposed in section 2 can be used when ex-ante effects are possible. This section also points out that one can always construct a conditional privacy regime that is superior to a publicity regime. Section 4 concludes, and an appendix in the end contains proofs of various propositions.

2. Simple Model of Information Gathering and Conditional Privacy Rights

In this section, I consider a simple model, where people's ex-ante behavior are not taken into consideration (or, alternatively, unaffected by the information gathering behavior of publishers). The objective is to highlight the effect of conditional privacy rights in two primary dimensions. First, these rights can be used to separate between people based on their subjective valuations of privacy. Second, allowing people to exercise conditional privacy rights can reduce the incentives of publishers to obtain information. This reduction contributes to welfare, if the publisher has incentives to over-invest in information gathering in the first place. The opposite conclusion is true if the publisher is under-incentivized to collect information. If privacy rights cause any perverse incentives for the publisher, they can be fixed by re-directing some of the proceeds from the collection of fees towards publishers.

To formalize these ideas, I consider a continuum of individuals, about whom

publishers may collect information. Because these individuals differ from each other with respect to their privacy costs, f denotes the density function describing the measure of individuals who have costs c, with $f(c) > 0$ for all $c > 0$, and F denotes the cumulative distribution function associated with f.

Publishers attempt to collect information about individuals, by expending costs of ψ, which leads to successful information collection, per individual, with probability $q(\psi)$ where $q' > 0$ and $q'' < 0$. The social value of disclosing information is g, and by disclosing this information the publisher reaps α fraction of these benefits.

Finally, the government allows individuals to exercise privacy rights, conditional on paying a fee of x. Thus, if publishers obtain information regarding a person and disclose it, the individual has two options: pay x and not incur the privacy cost of c, or do nothing and suffer the privacy cost of c. Thus, only individuals with $c > x$ exercise their conditional privacy rights.[5]

Given this notation, the publisher's expected benefit, as a function of ψ, is given by:

$$q(\psi)F(x)\alpha g - \psi \tag{1}$$

Therefore, the publisher's choice of effort, denoted ψ^P satisfies:

$$\frac{\partial q(\psi^P)}{\partial \psi} = \frac{1}{F(x)\alpha g} \tag{2}$$

Social welfare, which consists of the social value and the privacy costs associated with the information published minus the information gathering expenses, is given by:

$$q(\psi)\int_0^x (g-c)f(c)dc - \psi \tag{3}$$

Thus, the optimal investment in information collection, given any x, satisfies:

$$\frac{\partial q(\psi^*)}{\partial \psi} = \frac{1}{\int_0^x (g-c)f(c)dc} \tag{4}$$

A comparison of (2) and (4) reveals the conditions under which the publisher over- or under-invests in information collection, given any x. This comparison hinges on the average privacy cost incurred by individuals about whom private information is published. This value, which is a function of x, can be defined as:

$$\overline{c}(x) = \frac{\int_0^x cf(c)dc}{F(x)} \tag{5}$$

Lemma 1: *The publisher is over-incentivized [under-incentivized] to collect information, if the average privacy costs to individuals from information publication is greater [smaller] than the social value of the information that the*

[5] I assume that indifferent individuals do not exercise their conditional privacy rights.

publisher is unable to internalize. Stated differently, $\psi^* \leq \psi^P$ *(which implies* $q^* \leq q^P$*) iff* $(1-\alpha)g \leq \overline{c}(x)$.

Proof:

$$\frac{\partial q(\psi^*)}{\partial \psi} = \frac{1}{\int_0^x (g-c)f(c)dc} \geq \frac{1}{F(x)\alpha g} = \frac{\partial q(\psi^P)}{\partial \psi} \text{ iff} \tag{6}$$

$$F(x)\alpha g \geq \int_0^x (g-c)f(c)dc = F(x)(g-\overline{c}(x)) \tag{7}$$

Thus, $\frac{\partial q(\psi^*)}{\partial \psi} \geq \frac{\partial q(\psi^P)}{\partial \psi}$ iff $(1-\alpha)g \leq \overline{c}(x)$. Since $q'' < 0$, this implies that $\psi^* \leq \psi^P$ iff $(1-\alpha)g \leq \overline{c}(x)$.■

The rationale behind lemma is intuitive. The publisher's and the social objective would be aligned with each other if the publisher were to internalize all relevant costs and benefits. However, information dissemination causes positive as well as negative externalities. In particular, the publisher is unable to internalize $(1-\alpha)g$ of the social value associated with each piece of information disclosed, because he is unable to collect the entire value of the information from third parties who use the information, i.e. because he is unable to price discriminate. The publisher's inability to reap the entirety of the benefits associated with his investments leads him to under-invest in collecting information. On the other hand, the publisher imposes an average privacy cost of $\overline{c}(x)$ per person about whom information is published, and he is not forced to internalize any portion of these costs. The publisher's disregard of these costs cause him to over-invest in information collection. Thus, in aggregate, whether the publisher over- or under-invests in information collection depends on the magnitude of the positive versus negative externalities associated with information dissemination.

An immediate corollary of these observations is that social welfare can be increased, relative to the publicity regime, by using conditional privacy rights and setting $x = g$ when $(1-\alpha)g \leq \overline{c}(g)$. This is because $\overline{c}_x > 0$, and thus, absent conditional privacy rights, the publisher is over-incentivized to collect information. His deviation from the optimal investment in information collection can be mitigated by reducing the profitability of information collection, which is what happens when conditional privacy rights are introduced.

On the other hand, if $(1-\alpha)g \geq \overline{c}(g)$, social welfare can be increased by setting $x = g$, and diverting αg of the proceeds, per conditional privacy right exercised, to the publisher. This regime simply keeps the publisher's incentives unchanged relative to the publicity regime, but, causes the publication only of information whose social value is smaller than the privacy cost to the individual.

Finally, the government can align the publisher's incentives with the social objective, if it can impose a sanction on the publisher or provide him with a reward for each person who exercises his privacy rights. Sanctions would be optimal if the publisher is over-incentivized when $x = g$, and rewards would be optimal if he is under-incentivized when $x = g$. To formalize these results, let s

denote the amount that the government pays the publisher, per individual who exercises his conditional privacy right.

Proposition 1: *(i) It is optimal to set the price of privacy equal to the social benefit from the information (i.e. $x^* = g$), and it is optimal to reward [sanction] the publisher when the average privacy cost to individuals from publication (i.e. $\overline{c}(g)$) is smaller [greater] than the social gains from information dissemination that the publisher is unable to capture (i.e. $(1-\alpha)g$). The optimal reward [or sanction] is $s^* = \frac{F(g)}{1-F(g)}[(1-\alpha)g - \overline{c}(g)]$. (ii) A regime where $x = g$ and $s = \alpha g$ dominates the publicity regime. (iii) When $(1-\alpha)g < \overline{c}(g)$, a regime where $x = g$ and $s = 0$ dominates the publicity regime as well as the regime where $x = g$ and $s = \alpha g$.*

Proof: (i) Let $\widehat{s}(x) = \frac{F(x)}{1-F(x)}[(1-\alpha)g - \overline{c}(x)]$. When $s = \widehat{s}(x)$, it follows that the publisher's pay-off as a function of ψ is given by: $q(\psi)(F(x)\alpha g + (1 - F(x))\widehat{s}(x)) - \psi = q(\psi)(F(x)\alpha g - F(x)\overline{c}(x)) - \psi = q(\psi)\int_0^x (g-c)f(c)dc - \psi$. This implies that $\widehat{s}(x)$ alligns the publisher's incentives with the social objective, and therefore, given any x, maximizes social welfare. Given $s = \widehat{s}(x)$, social welfare becomes:

$$W = q(\psi^*)\int_0^x (g-c)f(c)dc - \psi^* \tag{8}$$

Thus, $\frac{dW}{dx} = (g-x)f(x)$, and therefore $x^* = g$.

(ii) Setting $s = \alpha g$ causes the publisher to have the same information collection incentives that he has in the publicity regime. Thus, both regimes produce the same ψ, and, therefore, the same q. Thus, the welfare in the two regimes, and their difference, are given by:

$$\begin{aligned} W^P - W^C &= q^P\int_0^\infty (g-c)f(c)dc - \psi^P - q^C\int_0^g (g-c)f(c)dc + \psi^C \qquad (9)\\ &= q^P\left(\int_0^\infty (g-c)f(c)dc - \int_0^g (g-c)f(c)dc\right) = q^P\left(\int_g^\infty (g-c)f(c)dc\right) < 0 \end{aligned}$$

where the P superscript denotes publicity and the C superscript denotes the regime where $s = \alpha g$ and $x = g$.

(iii) Let q^C, q^P, and q^0 denote the resulting probabilities in the regime where $s = \alpha g$ and $x = g$, the publicity regime, and the regime where $s = 0$ and $x = g$, respectively. It follows that $q^C = q^P > q^0 > q^*$ (the equality follows from part (ii) of this proposition; $q^C > q^0$, because setting $s > 0$ increases information collection incentives; and $q^0 > q^*$ because $(1-\alpha)g < \overline{c}(g)$). Thus, the first two regimes generate greater costs than the regime where $s = 0$, due to deviations from the optimal investment.■

Part (i) of proposition 1 identifies optimal policies assuming that s can be set at any level. Parts (ii) and (iii), on the other hand, focus on policies that

are simpler to implement: they require s to be set at either 0, or αg. The latter reward of αg has an intuitive meaning: it corresponds to the lost revenue of the publisher as a result of information removal. These policies may be more attractive when specific determinations of s^* would require large administrative costs, which is an issue that is not explicitly incorporated in the models presented in this article.

This section presents results pertaining to the social desirability of using conditional privacy rights when information collection causes no ex-ante effects. The next section incorporates these effects through a more complex model.

3. Conditional Privacy Regimes when there are Ex-ante Effects

To incorporate potential ex-ante effects, I consider a setting where the information of interest is generated by an act that confers benefits to the actor and may generate externalities. The benefit to the individual is denoted b, and is distributed with $k(b)$ across individuals, and the act causes an externality of $-h$ to third parties (there are harms to third parties iff $h > 0$, and there are benefits to third parties iff $h < 0$). The act generates information, which, as in the previous section, may be captured by a publisher with probability $q(\psi)$. If the information is captured and disclosed, the actor suffers a cost of c as in the previous section. To simplify the analysis, I assume that b and c are independently distributed. The government, as in the previous section, chooses x and s.

Thus, the sequence of events can be described as follows:

Period 0: The government sets x and s

Period 1: Actors decide whether or not to commit the act

Period 2: The publisher (henceforth J) decides how much effort to exert

Period 3: Actors, about whom information is captured, decide whether or not to pay x to exercise their conditional privacy rights.

Next, I proceed by backward induction to determine the subgame perfect Nash equilibria (SPNE) of this game.

3.1. Actors' and Publishers' Behavior

Period 3:

Actors about whom information is collected decide whether or not to exercise their privacy rights. It immediately follows that these individuals exercise their rights if $c > x \equiv c^*$.

Period 2:

The publisher has to decide how much effort to exert. It's pay-off is given by:

$$q(\psi)[(1-\gamma)\alpha g + \beta s] - \psi \tag{10}$$

where $1-\gamma$ and β are the measures of individuals with $c \leq x$ and $c > x$, respectively, who commit the act in period 1. This implies that a measure of $\gamma - \beta$ individuals do not commit the act. Thus, J chooses effort $q^o = q(\psi^o)$ that satisfies:

$$q'(\psi^o) = \frac{1}{(1-\gamma)\alpha g + \beta s} \tag{11}$$

Period 1:

In this period, actors engage in a simultaneous game where each chooses whether or not to commit the act. Each individual takes β and $1-\gamma$, as given, since no person has positive measure. These measures, i.e. β and $1-\gamma$, imply an expected probability of information collection in the second period, which I denote as q^e. Thus individuals commit the act if

$$\begin{aligned} b - q^e c &> 0 \text{ and } c \leq x; \text{ or} \\ b - q^e x &> 0 \text{ and } c > x. \end{aligned} \tag{12}$$

on the other hand individuals with

$$\begin{aligned} b - q^e c &\leq 0 \text{ and } c \leq x\,; \text{ or} \\ b - q^e x &\leq 0 \text{ and } c > x. \end{aligned} \tag{13}$$

do not commit the act. Thus, the actual first period measures of individuals committing the act in the first period, as a function of q^e, are given by:

$$1 - \gamma(q^e) = \int_0^x \int_{q^e c}^{\infty} k(b)dbf(c)dc \tag{14}$$

and

$$\beta(q^e) = \int_x^{\infty} \int_{q^e x}^{\infty} k(b)dbf(c)dc \tag{15}$$

Thus, a SPNE exists when the expected probability of information collection, q^e, equals the observed probability of information collection, $q^o = q(\psi^o(q^e))$, which is defined in (11) and the accompanying text.[6] Thus, the equilibrium condition can be defined as:

$$q(\psi^o(q^e)) = q^e \tag{16}$$

The next lemma formalizes that an equilibrium exists, and provides a sufficient condition for its uniqueness. Figure 1, below, illustrates the behavior of individuals in a SPNE.

[Insert Figure I]

Lemma 2: *Given any x, there is a an interior SPNE q^* such that $q(\psi^o(q^*)) = q^*$. Moreover, the equilibrium is unique when $s \geq 0$.*

Proof: First, note that the right hand side of (11) is positive when $q^e = 0$, which, in turn, implies that $\psi^o > 0$ and $q(\psi^o(0)) > 0$. Next, note that $q(\psi^o(1)) < 1$ since, $1-\gamma(1) > 0$, which implies, per (11), that ψ^o is finite. Thus,

[6] Otherwise, some actors have profitable deviations from their strategies in the first period. In particular, if $q^e > q^o$ an actor with $b \in (q^o, q^e)$ and $c < x$, can increase his expected pay-off by committing the act, and if $q^e < q^o$ an actor with $b \in (q^e, q^o)$ and $c < x$, can increase his expected pay-off by not committing the act.

if we express the equilibrium condition in (16) as $E(q^e) = q(\psi^o(q^e)) - q^e$, it follows that $E(0) > 0$ and $E(1) < 0$. Therefore, an equilibrium exists. Moreover, implicitly differentiating (11) reveals that:

$$\frac{\partial \psi^o}{\partial q^e} = -\frac{1}{[(1-\gamma(q^e))\alpha g + \beta s]^2 q''}(-\alpha g \frac{\partial \gamma}{\partial q^e} + s\frac{\partial \beta}{\partial q^e}) \tag{17}$$

as (14) and (15) illustrate $\frac{\partial \gamma}{\partial q^e} > 0$ and $\frac{\partial \beta}{\partial q^e} < 0$. Thus, $\frac{\partial \psi^o}{\partial q^e} < 0$ when $s > 0$. This implies that $\frac{\partial q(\psi^o(q^e))}{\partial q^e} = q' \frac{\partial \psi^o}{\partial q^e} < 0$, which, in turn implies that $\frac{\partial E(q^e)}{\partial q^e} < 0$, and therefore, the equilibrium is unique.■

In addition to showing that an equilibrium exists, lemma 1 also illustrates that unless s is negative and large, there is a single equilibrium. As illustrated in section 2, in a simpler framework it is not necessary to resort to a negative s to increase welfare relative to the publicity regime, and this reasoning extends to the present analysis. Proposition 2, below, formalizes this result, by constructing a regime that generates the same incentives for the publisher as in the publicity regime, but allows welfare to be maximized by choosing x to trade off other costs and benefits.

Proposition 2: *A regime where* $s = \frac{\alpha g}{\beta(q^P)}\gamma(q^P) \equiv s^I < \alpha g$ *produces the same information collection probability as a publicity regime for all* x*, where* q^P *is the information collection probability in a publicity regime.*

Proof: The publisher's pay-off in the proposed regime is given by

$$\begin{aligned} & q(\psi)[(1-\gamma)\alpha g + \beta s^I] - \psi \\ = \quad & q[(1-\gamma(q^e))\alpha g + \alpha g \frac{\beta(q^e)\gamma(q^P)}{\beta(q^P)}] - \psi \end{aligned} \tag{18}$$

Thus, it follows that when $q^e = q^P$, J chooses an effor that produces a collection probability of q^P, i.e. $q(\psi^o(q^P)) = q^P$. This is because, when $q^e = q^P$ (18) becomes $q(\psi)\alpha g - \psi$, and q^P, by definition, maximizes this expression.■

In the proceeding parts, I use the abbreviation *IPCP* to refer to regimes where $s = \frac{\alpha g}{\beta(q^P)}\gamma(q^P)$. This is short for **I**ncentive **P**reserving **C**onditional **P**rivacy regimes. Since $s > 0$ in these regimes, it follows that an IPCP generates a unique equilibrium. The next proposition illustrates the effect of varying x among IPCPs. These effects are illustrated in figure 2, below.

[Insert Figure 2]

Proposition 3: *Increasing* x *among IPCPs (i) reduces the measure of individuals who commit the act in the first period (chilling or deterrence effect) and (ii) increases the amount of information available (information effect).*

Proof: As proposition 2 demonstrates, all IPCPs generate the same probability of information collection, q^*. Thus, as (15) illustrates

$$\partial\gamma/\partial x = -f(x)(1 - K(q^* x)) < 0 \tag{19}$$

and

$$\partial\beta/\partial x = -f(x)(1-K(q^*x)) - q^*k(q^*x)(1-F(x)) < 0 \qquad (20)$$

(19) implies that $q^*(1-\gamma)$, i.e. the measure of individuals about whom information is disclosed (and who do not exercise their conditional privacy rights), is decreasing in x, which is the claim made in part (ii).

(i) On the other hand, the measure of individuals who do not engage in the act in the first period is given by $\gamma - \beta$. (19) and (20) together imply that the change in this measure in response to a change in x is given by $\gamma_x - \beta_x = q^*k(q^*x)(1-F(x)) > 0$.■

Proposition 3 illustrates how varying the cost of conditional privacy rights, among IPCPs, affects the behavior of individuals in the first period, and how it affects the amount of information publicized. These observations play a key role in illustrating the social welfare effects of varying x.

3.2. Social Welfare Analysis

Given individuals' behavior, as described in the preceding section, one can formulate social welfare as a function of the government's choice of x. The utilitarian social welfare that I consider consists of the sum of all benefits (the relevant bs and gs) net of costs (the relevant cs and ψs). Thus, social welfare is given by:

$$W = \int_0^x \int_{q^*c}^{\infty} (b - h + q^*(g-c))k(b)dbf(c)dc + \int_x^{\infty} \int_{q^*x}^{\infty} (b-h)k(b)dbf(c)dc - \psi^* \quad (21)$$

All components of social welfare, except for $-\psi^*$ are illustrated in figure 2. The same figure also illustrates two effects which can be described with reference to figure 2 as Δ_1 and Δ_2 as follows:

Δ_1 = changes in welfare due to chilling or deterrence effects ($-(b-h)$ per individual moving from area I to area III)

Δ_2 = changes in welfare as a result of more information publication

It is straight forward to note that the contribution of Δ_1 to social welfare is non-negative iff $q^*x \leq h$ and that the contribution of Δ_2 is non-negative iff $x \leq g$. Thus, an immediate result is that if $h > 0$, the welfare maximizing x among all IPCPs, denoted as x^I is such that $x^I \in [\min\{\frac{h}{q^*}, g\}, \max\{\frac{h}{q^*}, g\}]$. Equally straightforward is the observation that $x^I < g$ when $h \leq 0$. More generally, the welfare maximizing IPCP regime is obtained when x is chosen to optimally trade-off chilling/deterrence effects against gains from more information. However, the observations that $x^I < \max\{\frac{h}{q^*}, g\}$, regardless of the value of h, has a simple implication: there exists an IPCP that generates greater welfare than the publicity regime. This follows, because the publicity regime produces the same welfare as the limiting IPCP where $x \to \infty$. The fact that there exists a finite x^I that maximizes welfare among all IPCPs implies that the publicity regime is dominated. Proposition 4 summarizes these results.

Proposition 4: *(i) The welfare maximizing IPCP regime dominates the publicity regime. (ii)* $x^I \in [\min\{\frac{h}{q^*}, g, 0\}, \max\{\frac{h}{q^*}, g\}]$. *(iii) There exists a threshold*

externality $\underline{h} < 0$*, such that* $x^I > 0$ *if* $h > \underline{h}$*. (iv) When* $x^I > 0$*, it follows that* $x^I_g > 0$ *and* $x^I_h > 0$*.*

Proof: (i) Given that $s = s^I$ it follows that

$$\frac{\partial W}{\partial x} = \int_x^\infty k(q^*x)(h - q^*x)f(c)dc + \int_{q^*x}^\infty f(x)q^*(g-x)k(b)db \tag{22}$$

$$= k(q^*x)(h - q^*x)(1 - F(x)) + f(x)q^*(g-x)(1 - K(q^*x))$$

Thus, $\partial W/\partial x < 0$ for all $x > \max\{\frac{h}{q^*}, g\}$.

(iv) $\frac{\partial x^I}{\partial g} = -\frac{W_{xg}}{W_{xx}} = -\frac{f(x)q^*(1-K(q^*x))}{W_{xx}} > 0$, since $W_{xx}|_{x=x^I} < 0$, because x^I maximizes W. Similarly, $\frac{\partial x^I}{\partial h} = -\frac{W_{xh}}{W_{xx}} = -\frac{k(q^*x)(1-F(x))}{W_{xx}} > 0$.

(ii) If $\min\{\frac{h}{q^*}, g\} > 0$, it follows from (22) that $\frac{\partial W(0)}{\partial x} > 0$ and $\frac{\partial W(0)}{\partial x} < 0$ for all $x < \max\{\frac{h}{q^*}, g\}$, thus $x^I \in [\min\{\frac{h}{q^*}, g\}, \max\{\frac{h}{q^*}, g\}]$. However, when $\min\{\frac{h}{q^*}, g\} \leq 0$ it is possible that $x^I = 0$, since $\frac{\partial W(0)}{\partial x} < 0$ for sufficiently small h.

(iii) Denote $\frac{\partial W}{\partial x}$ as $\frac{\partial W(h,x)}{\partial x}$ and $x^I = x^I(h)$. It follows that $\frac{\partial W(0,0)}{\partial x} > 0$ which implies that there exists $\underline{h} < 0$ such that $\frac{\partial W(h,0)}{\partial x} > 0$ for all $h > \underline{h}$. Hence, $x^I(h) > 0$ for all $h > \underline{h}$. ■

Proposition 4 illustrates the simple trade-offs that emerge when the analysis is restricted to IPCPs. When the act that generates the information is socially desirable, the welfare maximizing x among IPCPs is smaller than the social value of the information. This means that some information, whose disclosure would, but for other effects, be socially desirable, is not disclosed. This is because the welfare maximizing x trades-off some informational benefits for benefits in the form of reduced chilling effects. On the other hand, when the act that generates the information is sufficiently harmful (i.e. $h > q^*g$), the welfare maximizing x is greater than g: some information whose disclosure causes significant privacy costs are nevertheless disclosed to increase deterrence. It is technically possible for $x^I = 0$. But, this is possible only if information disclosure is likely to cause very large chilling effects. Finally, x^I is increasing in both g and h, because increases in these variables imply increased gains from information disclosure in the form of increased deterrence (or reduced chilling of behavior) and greater informational benefits.

If s is not constrained to keep the publisher's incentives unchanged relative to the publicity regime, welfare can be increased even further. These regimes do not necessarily keep q^* constant, and thus, generate an additional trade-off that is tied to changes in this probability. These complications do not affect any of the qualitative results presented in proposition 4. In other words, the optimal conditional privacy rights regime responds to h and g in the same manner that the best IPCP regime does. The primary difference is that the optimal s may be a penalty on the publisher rather than a reward, i.e. $s^* < 0$ is possible. These claims are formalized in proposition 5.

Proposition 5: *(i) x^* is finite and there exists a threshold externality $\underline{h} < 0$, such that $x^I > 0$ iff $h > \underline{h}$. (ii) $s^* \geq 0$ iff $\overline{c}(x^*) \geq (1-\alpha)g$. (iii) When $x^I > 0$, it follows that $x_g^I > 0$ and $x_h^I > 0$.*

Proof: See appendix.

4. Transaction Costs and Strategic Interactions

The primary purpose of the analysis in section 3 is to highlight the separating function of conditional privacy rights. Only those individuals who have high valuations of privacy are willing to incur the cost of information removal. Thus, one can set x to make sure that only socially detrimental information is removed from the public sphere. This analysis abstracts from a couple of interrelated issues pertaining to Coasean bargaining, which can potentially be addressed in future research that extends the analysis contained herein. In particular, it is assumed that i-) legal actions are costless to take, ii-) a person who would otherwise have information removed through court order cannot circumvent the court by offering the publisher some compensation that is lower than x, iii-) a person with $c \in (\alpha g, x)$ cannot approach the publisher and offer him an amount larger than αg to stop him from publishing information, and iv-) the publisher cannot induce a person by paying him an amount greater than $c - x$ to publish the information, which the person could otherwise remove from the public space through court order.

Some of these assumptions are justified when there are large transaction costs (t) associated with negotiations and communications between the publisher and the person about whom information is collected. For instance, if $t > g$, in the regime described in section 2, the publisher and individuals (with any c) cannot generate a large enough surplus to make it worthwhile for them to negotiate. Even if this assumption were reasonable, it would be an interesting approach to consider transactions costs and simultaneously ignore litigation costs. Thus, an interesting question to ask is whether there are good reasons to think that negotiations between publishers and individuals may break down even if litigation is relatively inexpensive. One potential reason, briefly mentioned in the introduction, is the possibility of strategic interactions between publishers and individuals. If for instance, publishers may refuse information removal to establish a reputation to increase their over-all bargaining positions, fewer transactions than expected may take place.

Studying these and related issues can highlight the hidden costs and benefits associated with conditional privacy rights, and may reveal that the proposed optimal conditional privacy right regimes must be modified in light of these costs.

5. Conclusion

The economics literature has made multiple observations regarding the potential social welfare effects of privacy protection, and has pointed out both benefits and costs associated with regimes that protect privacy. This literature generally focuses on absolute privacy rights (as in the NIM example) as well as absolute publicity regimes (where there are no privacy rights). In this article, I have studied the properties of a hybrid regime where privacy rights are conditional, and have highlighted the benefits that can be generated through

this hybrid regime. The analysis reveals that a conditional privacy right regime can generate welfare increases by allowing people to signal their subjective valuation of privacy. Moreover, this regime can be designed in a way that does not interfere with information collection incentives, and without generating any additional costs in the form of inefficient chilling (or deterrence) effects.

Appendix

Notes on the Proof of Proposition 5:

(ii) s^* is calculated by solving for the s that equates the social objective function to the publisher's objective function.

(iii) This is obtained through the same method as in proposition 4 with the additional Δ_3.

(i) Same as in proposition 4.

References

[1] Acquisti, A., Taylor, C., and L. Wagman (Forthcoming). The Economics of Privacy.

[2] Bar-Gill, Oren, and Omri Ben-Shahar, "The Law of Duress and the Economics of Credible Threats", Journal of Legal Studies 33 (2004): 391-430.

[3] Daughety, A. and J. Reinganum (2010). Public goods, social pressure, and the choice between privacy and publicity. American Economic Journal: Microeconomics 2 (2), 191-221.

[4] Hirshleifer, J. (1971). The private and social value of information and the reward to inventive activity. The American Economic Review 61 (4), 561-574.

[5] Laudon, K. C. (1996). Markets and privacy. Communications of the ACM 39 (9), 92{104.

[6] Mungan, M, and M. Seidenfeld (2016). Investments to Make Threats Credible, Rent-Seeking, and Duress. Supreme Court Economic Review

[7] Noam, E. M. (1997). Privacy and self-regulation: Markets for electronic privacy. In Privacy and Self-regulation in the Information Age. US Department of Commerce.

[8] Seidenfeld, M & M. Mungan, "Duress as Rent Seeking", Minnesota Law Review 99 (2015): 1423-1457.

[9] Shavell, S., "Contractual Holdup and Legal Intervention", Journal of Legal Studies 36 (2007): 325-354.