

PRIVACY, PROPERTY RIGHTS & EFFICIENCY: The Economics of Privacy as Secrecy

9 November 2005

**Benjamin E. Hermalin
and
Michael L. Katz***

Abstract

There is a long history of governmental efforts to protect personal privacy and strong debates about the merits of such policies. A central element of privacy is the ability to control the dissemination of personally identifiable data to private parties. Posner, Stigler, and others have argued that privacy comes at the expense of allocative efficiency. Others have argued that privacy issues are readily resolved by proper allocation of property rights to control information. Our principal findings challenge both views. We find: (a) privacy can be efficient even when there is no “taste” for privacy *per se*, and (b) to be effective, a privacy policy may need to ban information transmission or use rather than simply assign individuals control rights to their personally identifiable data.

* The authors would like to thank Thomas Davidoff, Giancarlo Spagnolo, Jean Tirole, Hal Varian, Michael Waldman, and Luc Wathieu, as well as participants in the 2004 “Quantitative Marketing and Economics Conference” and 2004 “Conference on the Economics of Electronics Communications Markets,” for helpful discussions of these issues. We are particularly grateful to an anonymous referee. An earlier version of this paper was titled “Is Privacy Efficient? The Economics of Privacy as Secrecy.”

I. INTRODUCTION

There is a long history of contentious policy debates and governmental efforts to protect personal privacy. A central element of privacy is the ability to maintain control over the dissemination of personally identifiable data—privacy as secrecy.¹ Recent technological developments in information collection and processing have heightened privacy concerns.² Today, online bookstores know what you like to read, your TiVo reports your viewing habits to the company’s central database,³ and airlines keep a record of where you travel. Each year sees a number of privacy bills introduced in state legislatures and the U.S. Congress in response to privacy concerns, but there is little consensus on the appropriate approach.⁴ There are many calls for strong governmental intervention to restrict the use of personally identifiable data. However, there are also calls simply to establish appropriate property rights to information on the grounds that market forces will then lead to efficient privacy levels.

Broadly speaking, there are two reasons that an individual might wish to withhold personal information from others. First, an individual may wish to conceal information from a potential trading partner because revelation of the information would lead the partner to take actions that would have adverse consequences for the person about whom the information was

¹ A distinct conception of privacy is autonomy, both from the state (*e.g.*, the right to choose to have an abortion) and from annoyance by other private parties (*e.g.*, the ability to be free of telemarketing calls). For an early discussion by an economist of privacy as autonomy, see Hirshleifer (1980).

² Even this trend is not new. Concerns about increasing surveillance and data processing led to the amendment of the California State Constitution in the early 1970s to include an explicit right to privacy.

³ See, *e.g.*, “Greeting Big Brother with Open Arms,” *New York Times*, January 17, 2004, B9.

⁴ See Smith (2003) for a recent summary of federal legislation.

revealed.⁵ For example, revelation that a policyholder smokes might lead to less favorable insurance rates, or someone who revealed he was HIV positive would, to many, be a less attractive sexual partner. We refer to this motivation as concern for market-mediated consequences. Second, an individual may have a taste for privacy *per se* even if there are no market-mediated consequences (*e.g.*, someone simply does not like the thought of neighbors reviewing her online pornography purchases). Our focus is on the first reason.

We are concerned with two questions. First, is privacy efficient? Second, is there an assignment of property rights to personally identifiable information that leads to an optimal level of privacy or disclosure? Clearly, the provision of privacy can be efficient when individuals have a taste for privacy. In what follows, we assume that there are no such tastes and, instead, focus on the first motive for privacy. We assume that both parties to the trading relationship are private agents and neither party has the power to compel the other to take action involuntarily. We also assume that both private parties are economically rational, expected-utility maximizers.

At least since the work of Hirshleifer (1971), it has been recognized that economically rational parties may invest inefficient amounts in collecting (or concealing) information in order to affect the distribution of rents.⁶ For example, in order to increase its revenues from a given volume of sales, a seller may expend resources to collect information about potential customers' willingness to pay for its product that is privately valuable but socially worthless. This

⁵ Another situation in which privacy concerns arise is one in which the individual wishes to prevent a trading partner from intentionally or unintentionally sharing information with a third party (*e.g.*, the sale of mailing lists or the failure to take adequate measures to secure a database of credit card numbers). Of course, the two cases are linked when the third party obtaining the information is also one of the individual's trading partners. For a recent analysis of third-party sharing, see Kahn *et al.* (2000). See also Calzolari and Pavan (2004), who establish conditions under which privacy is an equilibrium outcome but do not examine whether such outcomes are efficient, and work cited therein.

⁶ See Taylor (2004) for a recent analysis along these lines.

observation suggests that there can be an efficiency-enhancing role for privacy regulations.⁷

However, exactly the same logic suggests that there can be an efficiency-enhancing rationale for making illegal any effort to keep information private! The reason is that the parties who possess information that is unfavorable to them (*e.g.*, workers with low expected marginal revenue products) have incentives to expend resources to conceal that fact even though there may be no social value from doing so.

In many situations, the administrative or transactions costs associated with disseminating information to—or concealing it from—a trading partner are low. A more fundamental question is how revelation of information to a potential trading partner affects the efficiency of the equilibrium actions taken by the parties. The Chicago School, most notably Posner (1981) and Stigler (1980), asserts that privacy is harmful to efficiency because it stops information flows that would otherwise lead to improved levels of economic exchange. The literature has identified at least three mechanisms through which these adverse effects may arise. First, a lack of information can prevent the realization of matching benefits.⁸ For example, it might be efficient for an employer to provide the most extensive training to those employees with the best long-term health prospects. A privacy policy that limited the disclosure of health information would be an obstacle to such matching. Second, privacy can lead to informational asymmetries that destroy markets and prevent efficient exchange.⁹ This would be true, for example, when individuals have significant private knowledge about their likelihood of suffering a particular

⁷ This does not, however, imply that protecting privacy will promote efficiency. Depending on the elasticity of demand for information, implementing a privacy policy that raises the cost of collecting information might actually worsen the inefficiency by leading to higher levels of socially unproductive expenditures.

⁸ Posner (1981) at 405, and Stigler (1980) at 629.

harm and insurance companies consequently face severe adverse selection problems.¹⁰ Lastly, privacy can discourage productive investments.¹¹ If one cannot reveal one's productivity, there can be less incentive to invest in increased productivity. An example would be policies that prohibit business school students from revealing their grades to potential employers.

Under the Chicago School view, more information is better, at least if obtained without cost, and thus privacy is generally inefficient unless some parties have a demand for privacy for its own sake. One author summarized the situation as follows:

In grossly oversimplified terms, the consensus of the law and economics literature is this: more information is better, and restrictions on the flow of information in the name of privacy are generally not social wealth maximizing, because they inhibit decisionmaking, increase transactions costs, and encourage fraud.¹²

A superficial analysis appears to support the Chicago position: Rational-actor models predict that, in the absence of transactions costs, perfectly informed parties will undertake all efficient trades, while it is well-established that imperfectly informed agents may fail to trade efficiently. This argument is, however, incomplete in three important respects.

First, welfare depends on more than *ex post* trade efficiency. Specifically, there can be *ex ante* efficiency effects on the provision of insurance and on investment incentives.¹³ Consider the effects on insurance. Privacy protection can create insurance that would otherwise be

⁹ This effect of privacy is implicit in the lemons model of Akerlof (1970) lemons model of asymmetric information. It can also be viewed as an extreme form of the first mechanism identified by Posner and Stigler.

¹⁰ We observe that, in order to understand the full effects of privacy policies, one must also examine other potential market responses to privacy, such as insurance suppliers' relying on employer-purchased plans to reduce self-selection. Wathieu (2002) examines the roles of intermediaries (*e.g.*, employers) who possess finer information about customers (*e.g.*, insurees) than product or service producers (*e.g.*, insurance companies).

¹¹ Stigler (1980) at 630-631.

¹² Murphy (1996) at 2382.

¹³ For a seminal analysis of the effects of *ex post* contracting on *ex ante* incentives, see Williamson (1975).

destroyed. For instance, if a potential policyholder can be tested for the likelihood of developing a fatal health condition, life insurance companies might demand to test potential policy holders and adjust prices according to the test results. The competitive equilibrium would be *ex post* efficient: each risk-averse person would purchase full insurance at an actuarially fair rate based on his or her test results. However, from an *ex ante* perspective, individuals would bear the risks associated with the outcomes of their test results. If testing—by either individuals or insurance companies—were banned, then the competitive equilibrium would entail all risk-averse individuals’ buying full insurance at a common rate. Welfare would be greater than under the testing equilibrium both because the (socially wasteful) costs of testing would be avoided and because risk-averse individuals would bear less risk.

With respect to investment, privacy policies affect not only the dissemination of information, but also its acquisition or investments in complementary assets (*e.g.*, developing a statistical model of consumer behavior to target marketing efforts based on household characteristics). Specifically, if a party has to disclose any information that it has collected, then it has reduced incentives to collect the information.¹⁴ For example, absent the ability to keep information confidential, people may not collect information about themselves (*e.g.*, individuals might forgo AIDS testing if disclosure were mandatory), resulting in unintended adverse consequences. When the information that would otherwise be collected has social value, this is a

¹⁴ The structure of the argument is isomorphic to the logic of granting patents and other intellectual property rights.

bad thing.¹⁵ Similarly, policies that influence the cost of obtaining information will also influence the incentives to make complementary investments.¹⁶

A second consideration, absent from the Chicago School argument, follows from the fact that there may be other market imperfections that interact with privacy. Specifically, in the presence of price rigidities (say due to regulations or social norms), markets can fail to adjust efficiently to additional information. For example, if a would-be employer cannot lower an employee's wage offer in response to unfavorable information, she may simply refuse to hire at all. Consequently, greater information dissemination can reduce the efficiency of the resulting trading equilibrium.¹⁷

In a sense, each of the two issues just summarized concerns an extension of the model covered by the Chicago School argument. The next point, however, goes to the heart of the argument itself. The supporting logic given above considers only the limiting case of full information and does *not* establish that intermediate increases in information are efficiency enhancing. In fact, we show below that additional intermediate levels information can, in general, raise or lower ex post trade efficiency and total surplus. We also establish specific conditions under which the welfare effects of additional intermediate levels of information are unambiguous. Roughly speaking, we establish conditions under which allowing households to reveal personally identifiable information increases total surplus because it allows firms to make tailored offers to households that facilitate efficient transactions that would otherwise fail to

¹⁵ Curiously, despite reaching his overall conclusion that privacy is harmful, Stigler (1980) also observes that disclosure can discourage efficient investment in obtaining information. He apparently failed to notice that this fact can be construed as an argument that privacy protection can be efficiency enhancing.

¹⁶ For a recent analysis in a related context, see Kahn *et al.* (2000).

¹⁷ This argument is an application of the general theory of the second best (Lipsey and Lancaster, 1956).

occur. However, we also establish conditions under which intermediate increases in information break what would have been efficient broad pooling of households types under privacy and lead to narrower pools that entail greater levels of inefficient exclusion.

Although we identify conditions under which the efficiency effects of privacy are unambiguous, there are many situations to which these conditions do not apply. In principle, one way around this problem would be to rely on the incentives of the private parties. For example, some people have argued that consumer surplus or total surplus could be maximized by giving households property rights to their personally identifiable information.¹⁸ Our analysis, however, demonstrates that such rights would be worthless in many cases. Specifically, we establish conditions under which the equilibrium outcome is independent of the assignment of rights to the personally identifiable data. Although this latter result is reminiscent of the Coase (1960) Theorem, in that the assignment of property rights is irrelevant to the determination of total welfare, there are three important differences. First, in most applications, the Coase Theorem implies that the assignment of property rights affects the distribution of surplus but not the total. Here, the assignment of property rights has no effects on either the distribution or the total level of surplus. Second, the Coase Theorem often fails when, as here, the parties bargain under asymmetric information.¹⁹ Third, the resulting outcomes in our applications are inefficient.

¹⁸ For example, Shapiro and Varian (1997, pages 29 and 30) argue that:

The right way to think about privacy, in our opinion, is that it is an externality problem. I may be adversely affected by the way people use information about me and there may be no way that I can easily convey my preferences to these parties. The solution to this externality problem is to assign property rights in information about individuals to those individuals. They can then contract with other parties, such as direct mail distributors, about how they might use the information.

¹⁹ See, *e.g.*, Hermalin and Katz (2005) for a discussion.

The remainder of the paper is organized as follows. Section II presents the general model and two results suggesting that privacy property rights are irrelevant. Section III examines a specific model of price discrimination by a monopolist or monopsonist that seeks personally identifiable data concerning potential trading partners. This example is of interest, in part, because many privacy advocates have asserted that e-commerce and other technologies (*e.g.*, supermarket frequent-buyer cards) are going to lead to pervasive, inefficient price discrimination and that state intervention is necessary to prevent this outcome.²⁰ Section IV examines a competitive market, which we describe in terms of an employment example. Privacy advocates strongly argue that workers should be protected from invasive questioning. But proponents of the Chicago School argue that efficiency is harmed by governmental limitations on employers' abilities to seek and act on information from or about potential employees.²¹ Our analysis demonstrates that there are complicated tradeoffs missed by both sides of the debate. The paper closes with a brief conclusion in Section V. Proofs not provided in the text can be found in the Appendix.

II. A GENERAL MODEL AND RESULTS

In this section, we describe the general model. We are interested in situations in which one side of the market, "firms," would like to learn the value of certain individually identifiable data concerning the other, "households." The labels firms and households are purely for expositional convenience. The analysis would apply equally well to any situation in which one party seeks information about another.

²⁰ For discussions of the Internet and price discrimination, see Acquisti and Varian (2001) and Odlyzko (2003).

²¹ See, *e.g.*, Posner (1981) at 405.

A type- θ household earns utility $u(\theta, x, t)$ from a transaction with a firm, where x is the amount of trade between the household and the firm, and t is the monetary transfer *from* the household *to* the firm. It is common knowledge that each household knows its type. A household's outside opportunity (its payoff in the absence of trade with a firm) is $\underline{u}(\theta)$. A firm earns $\pi(\theta, x) + t$ from a transaction with a type- θ household.²² In the case of a firm's selling a unit of some commodity at price p with marginal cost c that is independent of the buyer's type, we have $x \in \{0,1\}$, $t = px$ and $\pi(\theta) = -cx$. In the case of insurance, the cost could be a function of the buyer's type: $\pi(\theta, 1) < 0$ would be the expected payable claim of a type- θ household that purchased insurance. For an employer paying wage w to a risk-neutral worker with marginal revenue product θ , we have $t = -wx$, $\pi(\theta, x) = x\theta$, and $u(\theta, x, t) = w$, where x equals 1 if the worker is employed by that firm and 0 otherwise.

The value of θ is *soft* information, in that a household can choose to misreport its value because neither a firm nor a court can ever observe its value directly. For example, θ may be a measure of a household's willingness to pay for some product.

There is an indicator variable, σ , that is informative with respect to a household's type, θ . We assume that σ is drawn from a finite set, Σ . We also assume each household knows the value of its indicator variable but firms cannot observe a household's indicator variable unless the information is released to them. In contrast to the household's type, we assume that the indicator variable is *hard* information (*i.e.*, it can be concealed in some circumstances, but its value cannot be distorted or lied about if it is revealed). For example, a potential employee might be asked to release his or her medical records, or an e-tailer might seek to track a

²² Observe that we are assuming that the value of transacting with a given household is independent of a

consumer's purchase history. In a slight abuse of notation, we denote situations in which a household's indicator variable is concealed by $\sigma = 0$.

We now describe the structure of the game. We consider two versions of the game, which differ in terms of whether firms or households move first. We begin with the version in which firms move first. Firms simultaneously make offers. An offer by firm i , denoted M^i , is a menu of options—a mechanism—that yields a level of trade and transfer, (x, t) , as a function of a household's report, r , chosen from message space R and the household's revelation (or not) of its indicator value, $s \in \Sigma \cup \{0\}$. The value of r could be the household's representation of its value of θ , but in principle it could be an action (*e.g.*, as in a purchase decision under price discrimination) or even a multidimensional variable. Because it is hard information, $s = \sigma$ unless the household conceals it, in which case $s = 0$.

Households then simultaneously decide which, if any, offers to accept. Accepting an offer means that the household agrees to play according to the mechanism proposed by the corresponding firm. Let $i = 0$ denote the refusal to accept any offer. Our equilibrium concept is perfect Bayesian equilibrium.

Many people have argued that the assignment of property rights to personally identifiable information is an important element of privacy policy.²³ Thus, we are interested in comparing: equilibrium when firms have the right to compel revelation of the indicator variable with equilibrium when households have the right to conceal this information. Observe that, when firms possess the information property rights, a firm can—if it wishes—order households not to choose $s = 0$.

firm's transactions with other households.

Proposition 1: *If firms commit to contract offers before households reveal their indicator variables, then the set of equilibrium outcomes is independent of whether the property rights to personally identifiable data are granted to the firms or to the households.*

Proof of Proposition 1: Regardless of whether firms or households possess the information property rights, the set of strategies available to a firm comprises mappings of the form

$$M^i(r, s) = (x, t).$$

Consider, first, the situation in which firms possess the right to compel households to reveal σ . Observe that there is no loss of generality in assuming that firms always compel revelation of the indicator variable when they have the right to do so because: (a) a firm could compel revelation but commit to ignoring the value of the indicator variable (*i.e.*, it could offer a mechanism such that $M^i(r, s) = M^i(r, s')$ for all r, s , and s'); and (b) there is no value to non-revelation as a means of transmitting other information because the mechanism could instead rely on r .²⁴ Consider an equilibrium of this game and let $\tilde{M}^i(\cdot, \cdot)$ denote the mechanism offered by firm i with compulsory revelation. Let $(\tilde{i}(\theta, \sigma), \tilde{r}(\theta, \sigma), \tilde{s}(\theta, \sigma))$ denote the equilibrium strategy of an arbitrary household. For notational convenience, define $\tilde{M}^0(\cdot, \cdot) \equiv (0, 0)$; that is, if a household rejects all offers, it gets the no-trade outcome $(0, 0)$.

We next show that there is an equilibrium with the same outcome when the firms cannot compel households to reveal σ . Suppose that firm i offers the following contract to households:

²³ See, for instance, Varian (1997).

²⁴ In what follows, we could invoke the revelation principle, but the discussion in terms of general mechanisms better shows the logic of the argument.

(a) reveal σ and play mechanism $\tilde{M}^i(\cdot, \cdot)$, or (b) refuse to reveal σ and be refused trade. That is, firm i offers

$$\hat{M}^i(r, s) = \tilde{M}^i(r, s) \text{ if } s \in \Sigma \text{ and } (0, 0) \text{ if } s = 0.$$

We will now demonstrate that $\hat{M}^i(\cdot, \cdot)$ with voluntary revelation is equivalent to $\tilde{M}_i(\cdot, \cdot)$ with compulsory revelation from the perspective of households and the resulting profit levels for the firm.

Because trade was not compulsory in the first equilibrium, it follows that, if the firms offer $\{\hat{M}^i(\cdot, \cdot)\}$, there exists a continuation equilibrium in which households play $(\tilde{i}(\theta, \sigma), \tilde{r}(\theta, \sigma), \tilde{s}(\theta, \sigma))$. By construction, these strategy profiles yield the same outcome as under the first equilibrium.

The remaining question is whether the offerings $\{\hat{M}^i(\cdot, \cdot)\}$ constitute an equilibrium. As shown above, any feasible deviation when revelation of the indicator variable is not compulsory would have been feasible when it was. Thus, the fact that no deviation was profitable in the first equilibrium means that no deviation is profitable in the second.

One can readily make the argument in the reverse direction because a firm that possesses the information property rights can always choose a mechanism that does not make use of some or all of the available information. ***Q.E.D.***

Next, consider situations in which households make their revelation decisions before firms make offers. Let $\{\tilde{M}^i(r, s)\}$ denote the set of offers in an equilibrium that arises when firms have the information property rights. Define $V(M, s, \theta) \equiv \max_{r \in R} u(\theta, x(r, s), t(r, s))$ to be the utility of a type- θ household under mechanism M , when the household plays its optimal r

given s . We say that this equilibrium has a least-favored indicator value if there exists a $\underline{\sigma}$ such that $V(\tilde{M}^i, \underline{\sigma}, \theta) \leq V(\tilde{M}^i, \sigma, \theta)$ for all $\sigma \neq \underline{\sigma}$, for all i , and for all θ .

Proposition 2: *Suppose that households make their information revelation decisions before firms make offers. Consider any equilibrium that arises when firms are granted the information property rights. If there is a least-favored indicator value, then there exists an equilibrium when households hold the right to conceal their personally identifiable data that supports the identical outcome as the equilibrium that arises when firms possess the information property rights.*

Proof: Let $\{\tilde{M}^i(r, s)\}$ denote the set of offers in an equilibrium under compulsory disclosure.

Let $\underline{\sigma}$ be such that $V(\tilde{M}^i, \underline{\sigma}, \theta) \leq V(\tilde{M}^i, \sigma, \theta)$ for all $\sigma \neq \underline{\sigma}$, for all i , and for all θ .

Consider the game in which disclosure is voluntary. Faced with a set of disclosure decisions by households, firms must form beliefs about those households that choose not to disclose σ .

Consider the following strategies and beliefs when revelation cannot be compelled. All households choose to reveal their indicator variable. All firms believe that any household that chooses to conceal its indicator variable has an indicator variable equal to $\underline{\sigma}$. Firms also believe that the distribution of θ given non-revelation is the distribution of θ given $\underline{\sigma}$. If a household reveals σ , then firm i makes offer $\tilde{M}_i(\cdot, \sigma)$. Otherwise the firm offers $\tilde{M}_i(\cdot, \underline{\sigma})$.

Given that the compel equilibrium is an equilibrium, it follows that no firm would wish to deviate from its strategy when revelation cannot be compelled. Households for whom $\sigma = \underline{\sigma}$ would be indifferent between concealing and not. By construction, no household for whom

$\sigma \neq \underline{\sigma}$ would wish to deviate by concealing.²⁵ Thus, these strategies and beliefs constitute an equilibrium. Moreover, this equilibrium yields the same outcomes as the compel equilibrium.

Q.E.D.

Observe that Proposition 2, unlike Proposition 1, does not establish the equivalence of equilibrium sets. First, Proposition 2 depends on the existence of a least-favored indicator value. Although, this assumption is satisfied in many cases of interest, there could exist situations in which it is not. Second, even when a least-favored indicator value exists, the set of equilibria when households possess the information property rights may be larger than the set of equilibria when firms possess those rights. We illustrate both of these points in Section IV below.

Propositions 1 and 2 suggest that granting households privacy property rights will often fail to be a meaningful policy. Observe, however, that neither result establishes that the personally identifiable information will, in fact, be used, whether or not firms can compel its revelation. In the remainder of the paper, we examine two specific models that place greater structure on the problem and establish conditions under which the assignment of privacy property rights is indeed irrelevant.²⁶

The two specific models we will consider differ along several dimensions. One, as suggested above, is whether firms make offers before or after households make revelation decisions. A second is the nature of the households' types. Logically, there are two types of situation to examine: (a) pure price discrimination (*e.g.*, a seller learns something about the

²⁵ For an early application of this type of unraveling argument in a monopoly context, see, *e.g.*, Grossman (1981).

²⁶ Kahn et al. (2000) appear to obtain a conflicting result. However, as they themselves note, the assignment of property rights affects the equilibrium outcome in their model because of restrictions they place on the scope of contracts that can be written.

buyer's valuation, but the information tells the seller nothing about its costs), and (b) benefit-relevant information (*e.g.*, an insurance company obtains information about a potential buyer's health status or a firm learns about a potential employee's work record). The key distinction between these two is that, in the first but not the second, the transactions price is a sufficient statistic for calculating the benefits of trading with the specific partner. Pure price discrimination effects will arise only if the firms have some degree of market power. Below, we consider a monopolistic price discriminator but consider a perfectly competitive example of benefit-relevant information.

Given the possible irrelevance of whether households have the right to conceal information or firms have the right to compel its revelation, we are also interested in equilibrium when firms are legally prevented from making use of σ whether or not households would otherwise consent to its release and use. Below, we refer to this third case as "privacy." We also refer to situations in which households are legally compelled to reveal the value of their indicator variable whether or not firms desire it as the "absence of privacy."

III. MARKET POWER AND PRICE DISCRIMINATION

We begin by considering a market in which there is a single firm. Although the analysis applies equally well to a monopsonist, we describe the analysis in terms of a monopoly seller seeking information to serve as the basis of price discrimination. A privacy issue arises in this model because the monopolist potentially has available data on household characteristics that would allow the seller to divide buyers into sub-populations. For example, an e-commerce site might be able to divide its customers by certain socio-economic data (*e.g.*, home ownership) or by their past purchases.

We assume quasi-linear utility over the monopolist's output, x , and downward sloping demand curves. Letting $b_\theta(x)$ denote a type- θ household's inverse demand curve, a household's total dollar benefits from consumption of the monopolist's good are $\int_0^x b_\theta(t)dt \equiv B_\theta(x)$. In terms of our general formulation, $u(\theta, x, t) = B_\theta(x) - t$.

We assume demands are finite: for each type θ , $b_\theta(0) < \infty$ and there exists an $x_\theta^w < \infty$ such that $b_\theta(x) = 0$ for all $x \geq x_\theta^w$ and $b_\theta(x) > 0$ for all $x < x_\theta^w$. As is standard in the screening literature, we assume the demand curves of different household types do not cross: $\hat{\theta} > \theta$ implies $b_{\hat{\theta}}(x) > b_\theta(x)$ for all x such that $b_\theta(x) > 0$.

We assume that the marginal cost of production is constant and, for convenience, we take it to be 0. Hence, x_θ^w is the total-surplus-maximizing level of trade for a type- θ household.

We consider a model with a finite number of household types, $\theta \in \{1, 2, \dots, K\}$. Let N_σ denote the number of households whose indicator variable equals σ (we refer to such households as being in sub-population σ). Let $N_\sigma f_\sigma(\theta)$ denote the number of households of type θ whose indicator variable equals σ . The function $f_\sigma(\cdot)$ is the probability mass function for household types conditional on having indicator value σ . Recall that we use $\sigma = 0$ to denote the situation in which the firm cannot condition its offer on σ (privacy), so that $f_0(\cdot)$ is the probability mass function for the entire population of households. There is no loss in generality from assuming that $f_0(\theta) > 0$ for all θ .

As shown by Milgrom (1981), $\sigma = j$ is favorable information (is "good news") about the demand of members of group j if $\frac{f_j(\theta)}{f_0(\theta)}$ is strictly increasing. This case corresponds to a

situation in which households with $\sigma = j$ tend to have higher θ than households drawn from the overall population. It is immediate that, if the indicator variable takes only two values, say 1 and 2, then, $\sigma = 2$ being good news implies $\sigma = 1$ is bad news.²⁷

The general properties of the solution to the monopolist's pricing problem are well known, so we only summarize those properties here.²⁸ By the revelation principle, the seller can do no better than to offer a set of options (*e.g.*, package sizes) and let the buyers self-select among them. The seller's optimal mechanism presents the households with K options designed so that a household of type θ purchases an x_θ -unit package of output for a total payment of t_θ . It is readily shown that a higher-demand buyer must purchase at least as much as a lower-demand buyer; that is, $\hat{\theta} > \theta$ implies $x_{\hat{\theta}} \geq x_\theta$ for all $\theta, \hat{\theta} \in \{1, 2, \dots, K\}$. Without loss of generality, we can ensure all types are offered an option in equilibrium if we allow the firm to "sell" a package with 0 units for a price of 0.

As is well known, the seller will induce the first-best level of consumption for the highest-value consumers (*i.e.*, $x_K = x_K^w$). The reason is that, at any lower level of consumption, the monopolist could offer to sell additional output to its customers for an incremental amount equal to the relevant area under the demand curve of a household of the highest type. Doing so would not attract any lower types to mimic the highest type (the highest type has the highest marginal willingness to pay), and it would increase the seller's profits as long as the new quantity did not exceed the point at which marginal willingness to pay is equal to marginal cost.

²⁷ That is, $\frac{f_2(\theta)}{f_0(\theta)}$ increasing implies $\frac{f_0(\theta)}{f_1(\theta)}$ is increasing as well.

For all other household types, consumption is distorted downward. The reason is most easily seen by considering an illustrative special case. Suppose the monopolist faces two types of households with demand curves as drawn in Figure 1. As shown in Figure 1, the most that the monopolist could charge a low-demand buyer for x_1 units is given by the area under the lower demand curve up to the purchase quantity, area E in the figure. The most that the seller can charge a high-demand buyer for x_2 units is equal to the sum of areas E , G , and J . In particular, the seller must allow each high-demand household to earn an information rent equal to area I , which is the surplus a high-demand household would enjoy from purchasing x_1 for a total outlay equal to area E . Figure 2 illustrates the incremental revenue effects of increasing x_1 . The revenues collected from type-1 households rises by area Y times the number of type-1 households. But raising x_1 also imposes a cost on the seller: the revenues collected from type-2 consumers fall by area Z —the increased information rents—time the number of type-2 households. This cost induces the seller to reduce the consumption of each low-value customer below the first-best level (β in this example).

If the ratio of type-1 to type-2 households is sufficiently high, there exists an interior solution, which balances the marginal revenue gains and losses. Otherwise, it is privately optimal for the monopolist to exclude low-demand households from the market (*i.e.*, the monopolist sets $x_1 = 0$) in order to extract greater surplus from high-demand households. For future reference, observe that such exclusion of the low type results in the maximal equilibrium level of deadweight loss in the two-type case.

²⁸ See Katz (1983) for an example providing details of the analysis. The fundamental approach to this class of problems was pioneered by Mirrlees (1971) and applied by Spence (1980) to a multi-product monopolist facing discrete types.

We are now ready to consider the effects of information revelation. For the moment, assume there are only two household types, $\theta \in \{1,2\}$, and two values of the indicator variable, $\sigma \in \{1,2\}$. As long as the two sub-populations do not have the same ratio of high- and low-demand household types, one value of σ (say, $\sigma = 2$) must be good news and the other bad news. First, suppose that each sub-population contains both high and low demanders. Then, in each case, high types will continue to consume the efficient quantity. The story for low demanders is different. Because group 2 has proportionately fewer low types relative to high types, the seller will find it profitable to decrease the quantity offered to the low types in order to reduce the information rents enjoyed by the high types. Similarly, because group 1 has proportionately more low types relative to high types, the seller will find it profitable to increase the quantity offered to the low types because the revenue gains from the additional sales to low demanders will more than offset the increase in the information rents enjoyed by the high consumer types. Because the consumption level of low demanders is inefficiently low, the increase in consumption level for group 1 is efficient and the decrease for group 2 is inefficient.

This logic generalizes to arbitrary finite numbers of household types and values of the indicator variable:

Proposition 3: *Regardless of the privacy regime, a type- K household purchases an efficient amount. Suppose that, for all $x < x_\theta^w$, $b_{\theta+1}(x) - b_\theta(x)$ is non-decreasing in x for all $\theta \in \{1,2,\dots, K-1\}$.²⁹ Then for a household of type $\theta < K$ with indicator value $\sigma = q$:*

²⁹ This assumption rules out the possibility that marginal information rents fall as x_θ increases. If the marginal rents fell sufficiently fast, the monopolist's problem would not be concave and multiple equilibria might exist, making comparative statics difficult.

- (a) *If q is good news about the household's type, then privacy raises the household's consumption toward its efficient level if that household has positive consumption under privacy, and privacy has no effect otherwise (i.e., the household consumes 0 under either regime).*
- (b) *If q is bad news about the household's type, then privacy lowers the household's consumption away from its efficient level if the household would have positive consumption absent privacy, and privacy has no effect otherwise (i.e., the household consumes 0 under either regime).*

Proposition 3 demonstrates that privacy (i.e., prohibition of pricing contingent on σ) lowers efficiency for the bad-news realizations of σ , but raises it for the good-news realizations, which leaves open the question of the net effect on total surplus. It is well known that moving from simple monopoly pricing to third-degree price discrimination has ambiguous welfare effects in general (see, e.g., Varian, 1989). One can interpret this finding as a result about privacy: allowing the monopolist to segment the market into two or more groups results in an increase in information. Thus, the ambiguous welfare effects of third-degree price discrimination can be interpreted as ambiguous effects of increased information. However, except when each consumer has a 0–1 demand for the product, uniform pricing within each segment does not make full use of the information available to the monopolist. Hence, one cannot be sure that the ambiguous welfare effects aren't simply due to *ad hoc* restrictions imposed on the seller or on the form of consumer demands.³⁰ As our next result shows, however, this ambiguity persists even when both restrictions are dropped.

³⁰ The restriction to third-degree price discrimination when individual consumers have downward sloping demand curves is as an example of the price rigidities discussed in the introduction.

Proposition 4: *Suppose that there are two types of household, with linear demand functions having a common slope, and the indicator variable can take two values, one of which is good news (implying the other is bad news).*

- (a) If low-type households would be excluded from the market under privacy, then total surplus is weakly lower under privacy than in the absence of privacy.*
- (b) If, absent privacy, low-type households in both sub-populations would have positive consumption levels, then privacy leads to strictly higher total surplus than the absence of privacy.*
- (c) In other cases, privacy may raise or lower total surplus relative to the absence of privacy depending on the parameter values.*

The intuition for part (a) is straightforward. As noted earlier, exclusion of the low-type households maximizes deadweight loss, so no regime can do worse than a regime that excludes the low type. Moreover, if the proportion of low-type households in the bad-news sub-population (*e.g.*, $\sigma = 1$) is high enough, the seller will not exclude these households when conditioning on $\sigma = 1$, which results in welfare strictly greater than under privacy. The intuition for part (b) is somewhat more involved. If, in the absence of privacy, both sub-populations have sufficiently many low-type households that they are not excluded, then the low type is not excluded with privacy either. Moreover, it can be shown that, when the low type is not excluded, the deadweight loss from selling inefficiently little to the low type falls at a diminishing rate as the probability of the low type, $f_\sigma(1)$, rises. In other words, the deadweight loss for a population or sub-population is convex in $f_\sigma(1)$. Hence, the sum of the deadweight losses of the sub-populations is greater than the deadweight loss when the population is not divided and $f_\sigma(1)$ is at its average value. Part (c) follows from the first two because the relevant deadweight losses are continuous functions of the relevant parameters.

This analysis demonstrates that both sides of the e-commerce privacy debate have overstated their cases. Proposition 4 shows that those who assert that additional information gathering necessarily will improve efficiency are incorrect. Thus, there might appear to be scope for efficient governmental intervention. However, claims by privacy advocates that consumers necessarily are harmed by the loss of privacy are also incorrect. Depending on the specifics of the market under consideration, privacy may benefit or harm consumers, as we now demonstrate.

Regardless of the seller's information, low-demand households derive no surplus in equilibrium, so our interest is in high-demand households. Consider two scenarios. First, suppose that there are sufficiently few low-demand households in the overall population that, absent access to the indicator variable, the seller sets $x_1 = 0$. Then high-demand households enjoy no surplus (area I in Figure 1 collapses). Suppose that an improvement in information leads to at least one sub-population with both household types in which the proportion of low-demanders is sufficiently high that the seller sets $x_1 > 0$. Now, high-demand households in that sub-population earn positive surplus. Because the surplus of the high-demand households in the other sub-population can't be less than zero, total household surplus has increased.

Next, suppose that, if revealed, the indicator variable perfectly sorts the two types but absent access to the indicator variable, the seller chooses $x_1 > 0$. In this case, improved information lowers equilibrium household surplus: high-demand households enjoy positive surplus equal to area I in Figure 1 under privacy and but neither type gets surplus absent privacy.

Summarizing this discussion, we have proved:

Corollary 1: *Under monopoly screening, a switch to privacy raises equilibrium consumer (household) surplus for some parameter values and lowers it for others.*

As seen, a simple ban on the seller's collecting and using personally identifiable information can raise or lower equilibrium total surplus. A natural question is whether there are conditions under which one can unambiguously state whether privacy is efficient. Parts (a) and (b) of Proposition 4 provide an answer, but only for a limited set of circumstances. We turn next to deriving more broadly applicable conditions.

One such condition arises when the indicator variable partitions household types into one *block* (subset) of high types and one block of low types. Specifically, let there exist an $S \in \{1, 2, \dots, K - 1\}$ such that

$$f_1(\theta) > 0 = f_2(\theta) \text{ for all } \theta \in \{1, 2, \dots, S\}$$

and

$$f_1(\theta) = 0 < f_2(\theta) \text{ for all } \theta \in \{S + 1, S + 2, \dots, K\}.$$

To develop intuition, suppose there were just two consumer types. Then this partition results in perfect sorting. Hence, the monopolist engages in first-degree price discrimination, and total surplus is maximized. High-type households' consumption is unchanged; it is efficient with or without sorting. Low-type households' consumption rises because increasing their consumption level no longer forces the seller to offer greater information rents to the high-type households.

Under certain assumptions, this logic can be extended to an arbitrary number of household types.

Proposition 5: *Suppose that: (a) the personally identifiable information partitions household types; (b) the inverse demand curves of the different household types do not get farther apart as the households' type rise (i.e., $b_{\theta+1}(x) - b_{\theta}(x)$ is non-increasing in θ for all $1 \leq \theta < K$); and (c) the population hazard rate, $\frac{f_0(\theta)}{1 - F_0(\theta)}$, is non-decreasing. Then privacy is less efficient than information dissemination if:*

(i) *given any pair of blocks, all of the households in one block are of a higher type than any household in the other block; or*

(ii) $\frac{f_0(2)}{1-F_0(2)} > 1$;

or both.

Assumptions (b) and (c), the latter of which is standard in the screening literature, are sufficient to rule out bunching. Were certain household types bunched, it could be possible that partitioning into ordered blocks would lead to some types consuming less than under privacy.³¹

The intuition behind condition (i) can be understood as follows. Consider a partition of ordered blocks. The consumption levels of all household types at the top of their blocks are efficient, which, except for type K , is a strict efficiency improvement over privacy. For types in the top block, the tradeoff between efficiency and information rent is unchanged (given that none of these types was bunched with types in the next lower block) because the information rent is influenced only by higher types. Hence, the consumption levels for types in the top block remain unchanged. For types in other blocks, the tradeoff between efficiency and information rent tilts towards efficiency because the output allocated to any one of these types influences the information rent of fewer higher types than under privacy because some higher types are now in different blocks (again given no bunching with types in lower blocks). Hence the consumption of types in blocks below the top block weakly rise toward the efficient level. Because the consumption of at least some type is strictly more efficient and the consumption of no types is

³¹ For instance, suppose there are three types, with types 1 and 3 plentiful (especially type 1), while type-2 households are relatively rare and have demands close to those of type-1 households. Under privacy, it could be profit-maximizing to bunch types 1 and 2 at the same $x > 0$. Now consider the partition in which 1 is alone and 2 and 3 are together. Now the monopolist could prefer to shut out type 2. That is, partitioning in this way leads to type 2's consumption going down.

less efficient, the use of personally identifiable data that partition the type space in this fashion raises total surplus.

The logic behind the sufficiency of condition (ii) is similar to that of condition (i). For any type, the number of types higher than it in its block is no greater than in the overall population and typically less. If there were no change to any type's marginal contribution to higher types' information rent, then reducing the relevant number of higher types would improve welfare for the reasons given above. With an arbitrary partition, however, there is no guarantee that the next higher type within a block is the same as the next higher type within the overall population. When they are not the same, a type's contribution to the information rent of higher types goes up on a per-capita basis (*i.e.*, $b_{\theta+j}(x) - b_{\theta}(x)$ is increasing in j), which is a countervailing effect to the reduction in the relevant number of higher types. Combined with assumptions (b) and (c), condition (ii) is sufficient to prevent this countervailing effect from swamping the reduction-in-higher-types effect.³²

Lastly, we consider whether the allocation of property rights to personally identifiable information matters. Generically, there is a unique solution to the monopolist's profit-maximization program. Hence, the game has a single equilibrium and the following result is an immediate application of Proposition 1.

Proposition 6: *Suppose that there is single firm, which can commit to offering mechanisms on a take-it-or-leave-it basis. The equilibrium outcome is the same whether the property rights to personally identifiable data are given to the firm or to households.*

³² Assumption (c) and condition (ii) are satisfied, for example, by a Poisson distribution with mean less than $1 + \sqrt{3}$.

Recall that, when households possess the information property rights, the monopolist can (1) refuse to sell to any household that does not reveal its indicator variable,³³ and (2) offer the menus described above, conditional on indicator variable, to each household that reveals its indicator variable. The resulting outcome will be identical to the outcome without privacy. In other words, assigning information property rights is not enough to protect households or promote efficiency.

What is surprising is not that revelation is induced, but that the monopolist incurs no cost to do so. In other words, the assignment of property rights generates no information rents for households. This reflects a critical difference between hard and soft information—with hard information there is no way for one type to mimic another’s information, and hence no scope for it to capture an information rent.

IV. COMPETITIVE MARKETS AND ADVERSE SELECTION

We next examine the welfare effects of incremental information when households move before firms in a competitive market where the personally identifiable information is directly payoff relevant to a trading partner. Specifically, we consider a competitive employment market in which there are N_0 workers of varying abilities indexed by θ . In the following, the type space is either an interval, $[\underline{\theta}, \bar{\theta}]$, or there are just two types, θ_1 and θ_2 . To ensure gains from trade, we assume the minimum ability is positive.

³³ For example, an e-merchant could require consumers sign up and provide personal information before being allowed to shop.

There are more than N_0 potential employers, each of whom hires at most one employee.³⁴

Employers simultaneously bid for workers by announcing wage offers, where their offers can be made contingent on personally identifiable data if available to them. For example, if the data are available, wages can be made contingent on past employment or the potential employee's health status. Once wage offers have been made, workers decide which, if any, to accept.

A worker's utility is $\theta(1-x) + wx$, where $x \in \{0,1\}$ denotes whether she is employed in this sector ($x = 1$) or uses her time in some other way ($x = 0$), and w is her wage. Each worker knows her own ability and is willing to accept a job if the wage offered is θ or higher. The correlation of the worker's reservation wage and ability reflects that the value of her outside option (*e.g.*, becoming self-employed or working in some other sector) is likely an increasing function of her ability.

An employer's profit is $(v\theta - w)x$, where θ is the realized ability of the worker hired and $v > 1$ is the marginal revenue product of ability. Because $v > 1$ and $\min \theta > 0$, the first-best outcome requires that all workers be employed in this sector.

An employer cannot observe θ directly and must, instead, form a prediction of its value. A rational employer will pay a worker up to $v\theta^e$, where θ^e is the employer's expected value of the worker's ability conditional on the available information. To make the problem nontrivial, we assume that the amount rational employers bid for workers known to be of the lowest productivity is less than the reservation wage of the highest-ability worker.

Our first result concerns the continuous case.

³⁴ The adverse selection structure is similar to a used car market in which workers are the sellers of used cars and employers are potential buyers of used cars as studied in Akerlof's (1970) seminal article.

Proposition 7: *If $\sigma = q$ is good news about a household's ability, then privacy lowers the equilibrium wage paid to households for which $\sigma = q$. If $\sigma = q$ is bad news about a household's ability, then privacy raises the equilibrium wage paid to households for which $\sigma = q$.*

The efficiency effects of privacy are complex, as can be seen by considering what happens when the indicator variable takes on only two values, 1 and 2, and a value of 1 is bad news. Use of the indicator variable leads to a fall in the equilibrium wage paid to workers for whom $\sigma = 1$, which suggests that high-ability workers in this group could become inefficiently unemployed. But one must also account for the change in the mix of household abilities within each group. Suppose, for example, that all households for whom $\sigma = 1$ had abilities below θ^* and all households for whom $\sigma = 2$ had abilities above θ^* for some $\theta^* \in (\underline{\theta}, \bar{\theta})$. Even though the equilibrium wage offered to households with $\sigma = 1$ falls, total employment could rise. The reason is that those households with $\sigma = 1$ have lower reservation wages, and the households with higher reservation wages (*i.e.*, those with $\sigma = 2$) see their wage rise.

In order to keep track of these mix effects, we use a straightforward, two-type example to demonstrate that privacy can raise or lower welfare.³⁵ Denote the average worker ability when the proportion of high-ability workers is g as $\theta_A(g) \equiv (1-g)\theta_1 + g\theta_2$. A rational employer will pay a worker up to $v\theta_A(g^e)$, where g^e is the employer's belief about the proportion of high-ability job applicants conditional on available information.

³⁵ Levin (2001) has shown that a particular form of improved information (discussed below) can raise or lower efficiency when θ is continuously distributed.

We now characterize the competitive equilibrium for a population where employers know $f_\sigma(\theta_2) = g_\sigma$ and have no additional information. We consider two cases, which depend on whether g_σ is greater or less than $g^* \equiv \frac{\theta_2 - v\theta_1}{v(\theta_2 - \theta_1)}$. Define the *high-productivity case* as $g_\sigma > g^*$. In the high-productivity case, average productivity exceeds the high type's reservation wage (i.e., $v\theta_A(g_\sigma) > \theta_2$). Given this relationship, there exists an equilibrium in which all workers accept a wage of $v\theta_A(g_\sigma)$.³⁶ Because all workers are employed, the outcome is efficient. Thus, the asymmetric information about worker ability does not adversely affect the market outcome.

Define the *low-productivity case* as $g_\sigma < g^*$ (i.e., average productivity is less than high-ability households' reservation wage). In this case, the asymmetry of information leads to an adverse-selection problem—if employers thought all workers would be in the market, the most employers would bid is $v\theta_A(g_\sigma)$, which is less than θ_2 . Consequently, high-ability workers would exit the labor force. In equilibrium, employers anticipate this exit, so the equilibrium wage offer is $v\theta_1$, and high-quality workers are unemployed. The resulting deadweight loss is $N_\sigma g_\sigma (v-1)\theta_2$.

Our interest is in the role of information. Suppose that, initially, $g_1 = g_0 = g_2$. Then the indicator variable is uninformative, and revelation of the indicator variable has no effect on efficiency. Now consider the following thought experiment. Exchange a low-ability worker

³⁶ Because $v\theta_1 < \theta_2$, there is also a perverse Nash equilibrium in which employers expect workers of ability θ_2 not to seek employment and, thus, employers never bid above $v\theta_1$. However, this is not Bayesian perfect under the market structure assumed here. If an employer deviated and offered a wage of $\theta_2 + \varepsilon$, ε an arbitrarily small positive number, then all workers would be willing to be employed by that firm and the deviating employer would earn positive expected profit for sufficiently small ε because $v\theta_A(g) - \theta_2 > 0$.

from sub-population 2 with a high-ability worker from sub-population 1. That is, begin to sort low-ability workers into sub-population 1 and high-ability workers into sub-population 2. g_2 rises and g_1 falls. As we will now show, although complete employer information about worker types (full sorting) induces the efficient outcome, total surplus is a non-monotonic function of intermediate levels of information (partial sorting).

First, suppose that, with privacy, the overall population falls into the high-productivity case. Then the equilibrium with privacy is efficient and eliminating privacy (*i.e.*, revealing additional information) is weakly harmful. Provided both sub-populations are high-productivity, the additional information has no effect on efficiency. But once the workers are sufficiently sorted that sub-population 1 becomes low productivity (*i.e.*, g_1 falls below g^*), eliminating privacy destroys what would otherwise have been efficient pooling. Observe, however, that once efficient pooling has been destroyed, *further* sorting raises total surplus back towards the efficient-pooling level—such sorting has no effect on the employment of low-ability workers but it moves high-ability workers from inefficient unemployment as members of sub-population 1 to efficient employment as members of sub-population 2 ($g_1 < g^* < g_0$ implies $g_2 > g^*$). Nonetheless, until the sorting is perfect, the level of total surplus is still lower than under privacy.

Next, suppose that the overall population falls into the low-productivity case with privacy. In this case, the privacy equilibrium is inefficient—no high-ability worker is employed. Hence, improved information weakly increases total surplus. As long as both sub-populations are low-productivity, the additional information has no effect on efficiency. However, once the workers are sufficiently sorted that $g_2 > g^*$, high-ability workers in sub-population 2 are efficiently employed. Moreover, additional sorting increases total surplus further because the

number of (employed) high-ability workers in sub-population 2 rises and the number of (unemployed) high-ability workers in sub-population 1 falls.

This discussion establishes:

Proposition 8: *Suppose there are two levels of ability.*

- (a) *If the overall population is in the high-productivity case, then privacy is weakly more efficient than information dissemination and is strictly so if one of the sub-populations falls in the low-productivity case.*
- (b) *If the overall population is in the low-productivity case, then privacy is weakly less efficient than information dissemination and is strictly so if one of the sub-populations is in the high-productivity case.*

In addition to characterizing total surplus effects, we can examine the distributional consequences of information flows from workers to employers. Employers always make zero expected profits conditional on their information. In the high-productivity case, low-ability workers are made increasingly worse off by more information when $g_1 < g^* < g_0$. Weighted by the relative numbers, the increased income of those workers in sub-population 2 is more than offset by increased number of those workers in sub-population 1. In other circumstances, low-ability workers can gain from increased information. In particular, they gain when release of the indicator variable leads some of them to be pooled with high-ability workers and such pooling would not arise absent revelation of the indicator variable (*i.e.*, when $g_0 < g^* < g_2 < 1$). Moreover, continued increases in g_2 reduce the number of low-ability workers benefiting from pooling, but increases the wage, $v\theta_A(g_2)$, of those who are pooled. It can be shown that the net effect is ambiguous.³⁷

³⁷ Straightforward calculations reveal that the sign of the change is equal to the sign of $1 - 2g_2$.

High-ability workers gain from increased information when $g^* < g_1 < g_0 < g_2$. In this case, everyone is employed with or without privacy, but increased information means high-ability workers are less co-mingled with low-ability workers and, thus, appropriate a higher percentage of their revenue product. In this case, the interests of high- and low-ability workers are opposed. In other cases, however, both types can gain. Specifically, when $g_1 < g_0 < g^* < g_2 < 1$, both high- and low-ability households gain from being pooled in sub-population 2 compared to the original situation in which the equilibrium wage was $v\theta_1$ and only low-ability workers were employed. Households in sub-population 1 are unaffected. In yet other cases, both types can be harmed by increased information. This happens, for example, when $g_1 < g^* < g_0 < g_2 = g_1 + \varepsilon$ and ε is a small positive number. In this situation, the dominant effect is the collapse of pooling in sub-population 1 and both types of household see their expected wages fall.

Summarizing the analysis of expected wages, we see again that there is no dominant policy even if the goal is simply raising workers' welfare:

Corollary 2: *Suppose that there are two levels of ability. Depending on the parameter values, privacy can: (a) benefit both types of worker; (b) harm both types of worker; or (c) benefit low-ability workers at the expense of high-ability workers.*

The ambiguous welfare effects of additional information are not an artifact of our simple two-type setting. Proposition 8 can readily be generalized.

Corollary 3: *For a population with any number of ability levels, privacy is strictly more efficient than information dissemination whenever: (i) the equilibrium wage under privacy is higher than the reservation wage of the highest-ability worker, and (ii) there is a realization of the indicator*

variable such that the resulting equilibrium wage is lower than the reservation wage of the highest-ability worker who can experience that realization.

It is useful to ask whether there are other conditions under which the welfare effects of privacy are unambiguous. Levin (2001) examined information structures which, in our notation, correspond to situations in which $\sigma \in \{1,2\}$ and there exists $\hat{\theta}$ such that

$$f_1(\theta) > 0 = f_2(\theta) \text{ for all } \theta \in [\underline{\theta}, \hat{\theta}]$$

and

$$f_1(\theta) = 0 < f_2(\theta) \text{ for all } \theta \in (\hat{\theta}, \bar{\theta}].$$

In our context, his result can be viewed as follows. Suppose that, under the privacy equilibrium, an interval of the highest ability workers choose not to participate in this sector (*i.e.*, the equilibrium wage, w_0 , is below $\bar{\theta}$). Then, if $\hat{\theta} > w_0$, allowing employers to condition wage offers on σ cannot harm efficiency—there is no effect on the pool of workers already employed in this sector—but the use of σ may improve efficiency by allowing high-ability households to be hired at a wage they are willing to accept. Of course, for the reasons discussed in the example above, this result does not extend to situations in which the some of these high types would participate under the privacy equilibrium (*i.e.*, when $w_0 > \hat{\theta}$). Formally,

Proposition 9: *Suppose that under privacy the equilibrium wage is w_0 . Allowing households with reservation wages greater than w_0 to reveal themselves as such weakly raises welfare.*

Allowing households with reservation wages less than w_0 to reveal themselves can raise or lower welfare, depending on parameter values.

Now, consider a public policy of assigning to workers the property rights to their personally identifiable information, as many writers have advocated. If employers commit to

offers before workers choose whether to reveal their personally identifiable information, then Proposition 1 implies that the set of equilibrium outcomes is independent of the assignment of information property rights.

Now suppose that households move first. Consider an example in which workers have the right to keep their health status secret from potential employers, and consider the high-productivity case in our example. As shown above, if no worker were expected to reveal her health status, the equilibrium wage would be $v\theta_A(g_0)$. Observe that a worker with good health (*i.e.*, in sub-population 2) would have an incentive to reveal her health status if that would lead employers to hold the belief $\tilde{g} > g_0$ that she were a high-ability worker, because competition would then yield her an offer of $v\theta_A(\tilde{g})$, which is greater than $v\theta_A(g_0)$. Note the importance of employer beliefs. In particular, depending on employers' beliefs, there can be an equilibrium with full revelation (employers expect all workers to reveal their health status and they believe any worker who doesn't reveal has poor health) and an equilibrium with no revelation (employers expect all workers to conceal their health status and they believe only low-ability workers reveal their health status).³⁸

Because—even under the perfect Bayesian equilibrium concept—there is considerable latitude as to what beliefs firms can hold when households move first, we cannot establish the general result that the set of equilibrium outcomes is the same whether the property rights to personally identifiable data are given to the households or the firms. However, we can show,

Proposition 10: *Consider any equilibrium of the competitive employment model when firms are granted the property rights to personally identifiable information. The resulting outcome can*

³⁸ If health status is a perfect indicator of ability, then the no-revelation equilibrium doesn't exist.

also be supported as an equilibrium when the information property rights are granted to households.

V. CONCLUSION

With so many people making extreme claims in discussions of privacy and related public policy, and with so little understanding of the underlying economics, it is important to identify the fundamental forces clearly. A central fact is that, contrary to the Chicago School argument, the flow of information from one trading partner to the other can reduce *ex post* trade efficiency when the increase in information does not lead to symmetrically or fully informed parties. This finding gives rise to two broad, related questions for public policy. First, can one identify conditions under which public policy should or should not promote privacy? Second, in those circumstances where privacy is desirable, how can public policy promote it?

Unfortunately, we have not been able to find a simple answer to the first question. Although Levin (2001) and we have identified certain situations in which privacy is inefficient, such cases are limited. We are not optimistic that general conditions will be found. For decades, it has been known that third-degree price discrimination can be efficient or inefficient, and policy makers still lack readily applicable tests for when it falls into one category or another.³⁹ The response of American antitrust authorities has been to adopt a policy of allowing price discrimination except where it can be shown to harm the competitive process.⁴⁰ We suspect that privacy concerns are more important in situations where households are sellers because the

³⁹ A well-known necessary condition for efficiency is that total output rises, but this is not easily computed a priori.

impact of a wage change on household welfare is likely to be much more significant than that of a change in the price of a typical good or service purchased by the household. Casual observation suggests that, to date, distributional considerations have played a much more prominent role in policy deliberations over privacy than have efficiency concerns. Yet, as our analysis shows, discerning a privacy policy's distributional consequences is not always clear either; certainly in the case of employment, changes in privacy policy can make some households winners and others losers.

Turning to the question of how to promote privacy, economic analysis provides a strong indication of what *not* to do. Whether the uninformed side of the market is monopolized or is competitive, the assignment of privacy rights to personally identifiable information may have no effect on agents' equilibrium welfare levels and need not lead to an efficient equilibrium privacy level.

In some situations, the only effective policy would be explicitly to block the dissemination or use of such information. Public policy could block dissemination in several ways. One is to make it illegal to reveal personally identifiable data. Another is to destroy employment or prison records or other forms of tangible evidence, which would prevent households from credibly revealing the information even if they chose to do so. A related policy would be to refuse to enforce sanctions against people who lie about their protected characteristics.

⁴⁰ One rationale for this approach is that allowing price discrimination weakly raises a producer's profits, and this may generate increased investments in plant or R&D that benefit consumers in the long run. Of course, letting all firms in an industry engage in price discrimination might lower their profits. And there is no general theorem stating that the additional R&D is always worth more to consumers than its cost.

Policy makers might allow dissemination and instead attempt to block the use of the protected information. There are, however, potentially serious shortcomings with this approach. As is well known from the study of discrimination, if each employer were required to pay workers of all types the same wage, employers would be induced to specialize in terms of the abilities of workers they hired, so that the market as a whole could pay contingent wages that no one employer would be allowed to do. And, even if it were possible to mandate a uniform wage across the industry, employers could be expected to attempt to avoid hiring workers with unfavorable values of the protected characteristic.

APPENDIX

We begin by describing the general form of the problem faced by the profit-maximizing monopolist considered in the text. The incentives compatibility and individual rationality constraints associated with the seller's program are

$$B_\theta(x_\theta) - t_\theta \geq B_\theta(x_{\hat{\theta}}) - t_{\hat{\theta}} \quad \forall \theta, \hat{\theta}$$

and

$$B_\theta(x_\theta) - t_\theta \geq 0 \quad \forall \theta,$$

respectively. It is well known, that the seller's profit maximization program reduces to:

$$\max_{\{x_1, \dots, x_K\}} \sum_{\theta=1}^K f_\sigma(\theta) B_\theta(x_\theta) - \sum_{\theta=1}^{K-1} (1 - F_\sigma(\theta)) \{B_{\theta+1}(x_\theta) - B_\theta(x_\theta)\} \quad (\text{A1})$$

$$\text{s.t. } x_{\theta+1} \geq x_\theta \geq 0,$$

where $F_\sigma(\theta) \equiv \sum_{i=1}^{\theta} f_\sigma(i)$.⁴¹ The marginal benefit to the monopolist of increasing x_θ is thus

proportional to

$$b_\theta(x_\theta) - \frac{(1 - F_\sigma(\theta))}{f_\sigma(\theta)} \{b_{\theta+1}(x_\theta) - b_\theta(x_\theta)\} \quad (\text{A2})$$

(we need not worry about the definition of $b_{K+1}(\cdot)$ because $1 - F_\sigma(K) = 0$). The firm increases x_θ as long as (A2) is positive, unless the $x_{\theta+1} \geq x_\theta$ constraint binds.

Define $f_\lambda(\cdot) \equiv (1 - \lambda)f_m(\cdot) + \lambda f_n(\cdot)$, and define $F_\lambda(\cdot)$ analogously. Observe that, if

$f_m(\theta)/f_n(\theta)$ is monotonic in θ , then so too is $f_\lambda(\theta)/f_n(\theta)$. Define

$$V(\mathbf{x}, \lambda) \equiv \sum_{\theta=1}^K B_\theta(x_\theta) - \sum_{\theta=1}^{K-1} \frac{1 - F_\lambda(\theta)}{f_\lambda(\theta)} \{B_{\theta+1}(x_\theta) - B_\theta(x_\theta)\},$$

where $\mathbf{x} \equiv (x_1, \dots, x_K)$. Let X be the subset of \mathfrak{R}^K such that $x_{\theta+1} \geq x_\theta \geq 0$ for all $\theta \in \{1, \dots, K-1\}$.

⁴¹ See, for example, Katz (1983) or Spence (1980) for details.

Lemma A.1: Suppose that, for all $x < x_\theta^w$, $b_{\theta+1}(x) - b_\theta(x)$ is non-decreasing in x for all

$\theta \in \{1, 2, \dots, K-1\}$. Then $V(\mathbf{x}, \lambda)$ and X have the following properties:

- (a) $V(\cdot, \lambda)$ is supermodular for all λ .
- (b) $V(\cdot, \lambda)$ has a unique maximizer in X for all λ .
- (c) $V(\cdot, \cdot)$ has decreasing first differences if $f_n(\theta)/f_m(\theta)$ is increasing in θ .
- (d) X is a lattice.

Proof: (a) By Topkis's Characterization Theorem (Milgrom and Roberts, 1990), property (a) is

implied by the fact that $\frac{\partial^2 V}{\partial x_\theta \partial x_{\hat{\theta}}} \geq 0$ for $\theta \neq \hat{\theta}$.

(b) $V(\cdot, \lambda)$ attains a maximum over X because $V(\cdot, \lambda)$ is continuous and bounded and there would be no loss of generality in limiting its domain to $X \cap [0, x_K^w]^K$, which is closed and bounded. $b_\theta(x)$ is decreasing and $b_{\theta+1}(x) - b_\theta(x)$ non-decreasing in x for all $\theta \in \{1, 2, \dots, K-1\}$. Thus, $V(\cdot, \lambda)$ is concave. Given that X is convex, the maximizer of $V(\cdot, \lambda)$ is unique.

(c) The additive separability of $V(\mathbf{x}, \lambda)$ with respect to the components of \mathbf{x} implies it is sufficient to show decreasing differences with respect to a single x_θ , which can be shown by demonstrating that the cross-partial derivative of V with respect to λ and x_θ ,

$$\frac{(F_n - F_m)f_\lambda + (1 - F_\lambda)(f_n - f_m)}{f_\lambda^2} (b_{\theta+1} - b_\theta), \quad (\text{A3})$$

is negative, where arguments have been suppressed for readability. Straightforward algebra reveals that the sign of (A3) equals the sign of $F_n(\theta) - F_m(\theta)$. As is well known, $f_n(\theta)/f_m(\theta)$ increasing in θ implies that $F_m(\theta) > F_n(\theta)$ for all $\theta < K$, which establishes (c).

(d) X is a lattice if the meet (pointwise minimum) and join (pointwise maximum) of any two elements of X are in X . Suppose the join of \mathbf{x}^1 and \mathbf{x}^2 was not in X for two elements \mathbf{x}^1 and \mathbf{x}^2 of X . Then there must exist a θ such that $\max\{x_\theta^1, x_\theta^2\} > \max\{x_{\theta+1}^1, x_{\theta+1}^2\}$, which is impossible given that $x_\theta^i < x_{\theta+1}^i$ for $i = 1, 2$. The proof that the meet is in X is similar. **Q.E.D.**

Lemma A.2: Let \mathbf{z} be the value of $\mathbf{x} \in X$ that maximizes $V(\mathbf{x}, \lambda)$. Then \mathbf{z} uniquely maximizes $\Pi(\mathbf{x}, \lambda)$ subject to $\mathbf{x} \in X$, where

$$\Pi(x, \lambda) = \sum_{\theta=1}^K f_\lambda(\theta) B_\theta(x_\theta) - \sum_{\theta=1}^{K-1} \{1 - F_\lambda(\theta)\} \{B_{\theta+1}(x_\theta) - B_\theta(x_\theta)\}.$$

Proof: In maximizing either V or Π , there are $K - 1$ constraints of the form $x_{\theta+1} - x_\theta \geq 0$. Form the Lagrangian

$$\begin{aligned} L_\Pi &= \Pi(\mathbf{x}, \lambda) + \sum_{\theta=1}^{K-1} \mu_\theta (x_{\theta+1} - x_\theta) + \mu_0 x_1 \\ &= \Pi(\mathbf{x}, \lambda) + \sum_{\theta=1}^K (\mu_{\theta-1} - \mu_\theta) x_\theta, \end{aligned}$$

where μ_0 is the Lagrange multiplier on the restriction that sales be non-negative and $\mu_K \equiv 0$.

Making the change of variables $\alpha_\theta = \mu_\theta - \mu_{\theta-1}$, the Lagrangean is

$$L_\Pi = \Pi(\mathbf{x}, \lambda) - \sum_{\theta=1}^K \alpha_\theta x_\theta.$$

We have a solution if there exist \mathbf{x} and $\{\alpha_\theta\}$ such that

$$\sum_{k=\theta+1}^K \alpha_k \leq 0 \text{ for all } \theta \in \{0, \dots, K-1\}$$

$$\frac{\partial L_\Pi}{\partial x_\theta} = f_\lambda(\theta) b_\theta(x_\theta) - (1 - F_\lambda(\theta)) (b_{\theta+1}(x_\theta) - b_\theta(x_\theta)) - \alpha_\theta = 0, \quad (\text{A4})$$

and

$$(x_{\theta+1} - x_\theta) \sum_{k=\theta+1}^K \alpha_k = 0.$$

We can similarly write the Lagrangean for the V problem as

$$L_V = V(\mathbf{x}, \lambda) - \sum_{\theta=1}^K \tilde{\alpha}_\theta x_\theta.$$

The problem is unchanged if we scale each Lagrange multiplier by defining $\hat{\alpha}_\theta \equiv \tilde{\alpha}_\theta f_\lambda(\theta)$. We then have a solution for the V problem if there exist \mathbf{x} and $\{\hat{\alpha}_\theta\}$ such that

$$\sum_{k=\theta+1}^K \hat{\alpha}_k \leq 0 \text{ for all } \theta \in \{0, \dots, K-1\},$$

$$\frac{\partial L_V}{\partial x_\theta} = b_\theta(x_\theta) - \frac{1 - F_\lambda(\theta)}{f_\lambda(\theta)} (b_{\theta+1}(x_\theta) - b_\theta(x_\theta)) - \frac{\hat{\alpha}_\theta}{f_\lambda(\theta)} = 0. \quad (\text{A5})$$

and

$$(x_{\theta+1} - x_\theta) \sum_{k=\theta+1}^K \hat{\alpha}_k = 0.$$

Multiplying (A5) by $f_\lambda(\theta)$, we see that, if \mathbf{x} and $\{\hat{\alpha}_\theta\}$ are a solution to (A5), then they also satisfy (A4). V and Π are both concave (the proof of the latter parallels the proof that V is concave) and X is convex. Hence, both V and Π have unique maximizers in X . Therefore the \mathbf{x} and $\{\hat{\alpha}_\theta\}$ that solve the V problem are the unique solution to the Π problem. ***Q.E.D.***

Proof of Proposition 3: (a) When $\sigma = q$ is good news, $f_q(\theta)/f_0(\theta)$ is strictly increasing. Define $f_\lambda(\cdot) \equiv (1-\lambda)f_0(\cdot) + \lambda f_q(\cdot)$. By Lemma A.2, we need only compare the consumption allocation that maximizes $V(\mathbf{x}, 0)$ to the one that maximizes $V(\mathbf{x}, 1)$. Denote the former by \mathbf{x}^0 and the latter by \mathbf{x}^q . By Lemma A.1, Topkis's Monotonicity Theorem (Milgrom and Roberts, 1990) implies that

the pointwise maximum of \mathbf{x}^0 and \mathbf{x}^q also maximizes $V(\mathbf{x},0)$.⁴² Because $V(\mathbf{x},0)$ has a unique maximizer (Lemma A.1), the pointwise maximum of \mathbf{x}^0 and \mathbf{x}^q equals \mathbf{x}^0 ; that is, $x_\theta^0 \geq x_\theta^q$ for all θ . As discussed in the text, there is no distortion at the top: $x_K^0 = x_K^q = x_K^w$. Because there is no negative consumption, $x_\theta^0 = 0$ implies $x_\theta^q = 0$. To conclude our proof of part (a) we need to show that $x_\theta^0 > 0$ implies $x_\theta^0 > x_\theta^q$ for $\theta < K$. Suppose, counterfactually, that there exists at least one $\theta < K$ such that $x_\theta^0 = x_\theta^q > 0$. Consider the smallest such θ , $\tilde{\theta}$. Because $\tilde{\theta}$ is the smallest such θ , it must be that $x_{\tilde{\theta}}^q > x_{\tilde{\theta}-1}^q$ (without loss of generality, we can use the convention $x_0^\sigma = 0$ should $\tilde{\theta} = 1$). The fact that the $x_{\tilde{\theta}}^q \geq x_{\tilde{\theta}-1}^q$ constraint is not binding implies that

$$0 \leq \frac{\partial V(\mathbf{x}^q, 1)}{\partial x_{\tilde{\theta}}} = b_{\tilde{\theta}}(x_{\tilde{\theta}}) - \frac{1 - F_q(\tilde{\theta})}{f_q(\tilde{\theta})} (b_{\tilde{\theta}+1}(x_{\tilde{\theta}}) - b_{\tilde{\theta}}(x_{\tilde{\theta}})).$$

As is well known, if $f_\sigma(\theta)/f_{\sigma'}(\theta)$ is increasing, then

$$\frac{f_\sigma(\theta)}{1 - F_\sigma(\theta)} < \frac{f_{\sigma'}(\theta)}{1 - F_{\sigma'}(\theta)}. \quad (\text{A6})$$

Hence,

$$\frac{\partial V(\mathbf{x}^0, 0)}{\partial x_{\tilde{\theta}}} = b_{\tilde{\theta}}(x_{\tilde{\theta}}) - \frac{1 - F_0(\tilde{\theta})}{f_0(\tilde{\theta})} (b_{\tilde{\theta}+1}(x_{\tilde{\theta}}) - b_{\tilde{\theta}}(x_{\tilde{\theta}})) > b_{\tilde{\theta}}(x_{\tilde{\theta}}) - \frac{1 - F_q(\tilde{\theta})}{f_q(\tilde{\theta})} (b_{\tilde{\theta}+1}(x_{\tilde{\theta}}) - b_{\tilde{\theta}}(x_{\tilde{\theta}})) \geq 0.$$

This expression implies that the constraint that $x_\theta^0 \leq x_{\theta+1}^0$ is binding when $\sigma = 0$ (if not, then it would be profitable to increase $x_{\tilde{\theta}}^0$, contradicting the optimality of \mathbf{x}^0). Given $\mathbf{x}^0 \geq \mathbf{x}^q$,

⁴² The relevant aspects of Topkis's Monotonicity Theorem can be summarized as follows: Let X be a lattice and Y be a partially ordered set (a property clearly satisfied by the interval $[0,1]$). Let $\phi(x, y) : X \times Y \rightarrow \Re$ be supermodular in x for any given y and let that function exhibit decreasing

$x_{\tilde{\theta}}^q = x_{\tilde{\theta}}^0 = x_{\tilde{\theta}+1}^0$ implies $x_{\tilde{\theta}}^q = x_{\tilde{\theta}+1}^q$. Moreover, this argument can be repeated inductively so that if $x_{\tilde{\theta}}^0 = x_{\tilde{\theta}+i}^0$ for $i = 1, \dots, I$, then $x_{\tilde{\theta}}^q = x_{\tilde{\theta}+i}^q (= x_{\tilde{\theta}}^0)$ for $i = 1, \dots, I$. Because it is never optimal to set $x_{\tilde{\theta}} > x_{\tilde{\theta}}^w$, we know $\tilde{\theta} + I < K$; that is, there is a type $\tilde{\theta} + I + 1$ such that $x_{\tilde{\theta}}^0 = \dots = x_{\tilde{\theta}+I}^0 < x_{\tilde{\theta}+I+1}^0$.

Substituting the constraint $x_{\tilde{\theta}}^q = \dots = x_{\tilde{\theta}+I}^q$ in the maximization of $V(\mathbf{x}, 1)$ (recall $\lambda = 1$)

corresponds to $\sigma = q$), the first-order condition with respect to $x_{\tilde{\theta}}^q$ implies

$$\begin{aligned} 0 &\leq \sum_{\theta=\tilde{\theta}}^{\tilde{\theta}+I} \left(b_{\theta}(x_{\tilde{\theta}}) - \frac{1-F_q(\theta)}{f_q(\theta)} (b_{\theta+1}(x_{\tilde{\theta}}) - b_{\theta}(x_{\tilde{\theta}})) \right) \\ &< \sum_{\theta=\tilde{\theta}}^{\tilde{\theta}+I} \left(b_{\theta}(x_{\tilde{\theta}}) - \frac{1-F_0(\theta)}{f_0(\theta)} (b_{\theta+1}(x_{\tilde{\theta}}) - b_{\theta}(x_{\tilde{\theta}})) \right) = \sum_{\theta=\tilde{\theta}}^{\tilde{\theta}+I} \frac{\partial V(\mathbf{x}^q, 0)}{\partial x_{\theta}}, \end{aligned}$$

where the second line follows from (A6) and the fact that $x_{\tilde{\theta}}^0 < x_{\tilde{\theta}+I+1}^0$. But this implies that it would be feasible to increase profits when $\sigma = 0$ by raising $x_{\tilde{\theta}}^0 = \dots = x_{\tilde{\theta}+I}^0$, which contradicts the optimality of \mathbf{x}^0 . Hence, by contradiction, we've established the rest of part (a).

Part (b) has the identical proof, except that, now, $f_{\lambda}(\cdot) \equiv \lambda f_0(\cdot) + (1-\lambda)f_q(\cdot)$. **Q.E.D.**

Proof of Proposition 4: Let $-s$ be the common slope. Let $b_2(x) - b_1(x) = \delta$, where $\delta > 0$. If x_1 is an interior solution, then it follows from (A2) that

$$x_1 = \frac{\beta - \delta H_{\sigma}}{s},$$

where $H_{\sigma} = f_{\sigma}(2)/f_{\sigma}(1)$. An interior solution exists only if and only if $\beta > \delta H_{\sigma}$. The resulting

deadweight loss triangle per low-type is $\frac{s}{2}(x_1^w - x_1)^2$.

differences in x and y . Then, if $y \leq y'$, the join (pointwise maximum) of the x that maximizes $\phi(x, y)$ and the x' that maximizes $\phi(x, y')$ maximizes $\phi(x, y)$.

Consider each case in turn:

- (a) If $\beta \leq \delta H_0$, then $x_1 = 0$ under privacy and the deadweight loss is maximized. If $\beta > \min\{\delta H_1, \delta H_2\}$, then $x_1 > 0$ for one sub-population absent privacy, and privacy is less efficient.
- (b) When $\beta > \max\{\delta H_1, \delta H_2\}$, direct calculations show that privacy reduces deadweight loss.⁴³
- (c) The per-capita deadweight loss under dissemination tends to 0 as the sorting of household types becomes perfect and thus dissemination is then more efficient than privacy. Instances in which the information improvement lowers total surplus can be constructed by making use of (b) and the continuity of average deadweight loss with respect to the proportion of low types, and considering values of $\max\{\delta H_1, \delta H_2\}$ that are just above β . **Q.E.D.**

Proof of Proposition 5: As a preliminary, observe that, if (A2) is strictly increasing in θ for all x , then the associated order constraint (*i.e.*, $x_{\theta+1} \geq x_\theta$) is not binding and x_θ is found by setting (A2) to zero and solving. Therefore, if (A2) is strictly increasing in θ for two distinct values of σ , say m and n , and if

$$\frac{1 - F_m(\theta)}{f_m(\theta)} < \frac{1 - F_n(\theta)}{f_n(\theta)},$$

then x_θ is greater when $\sigma = m$ than when $\sigma = n$.

By assumptions (b) and (c), (A2) is strictly increasing in θ for all x when $\sigma = 0$ (recall $b_\theta(x)$ is strictly increasing in θ).

⁴³ Calculations for parts (b) and (c) available from the authors upon request.

Suppose types are partitioned into J blocks, and let S_j denote the maximal element in block j . Index the blocks so that $1 \leq S_1 < \dots < S_J = K$. Let $\sigma = j$ be that value of the indicator variable that indicates that a household with that realization of σ is in the j^{th} block.

Under condition (i), any type in block i is lower than any type in block j if $i < j$. Observe that $f_\sigma(\theta)$ is equal to $f_0(\theta) \frac{N_0}{N_\sigma}$ if θ is in the σ^{th} block and 0 otherwise. Hence,

$$\frac{1 - F_\sigma(\theta)}{f_\sigma(\theta)} = \frac{\sum_{t>\theta} f_\sigma(t)}{f_\sigma(\theta)} = \frac{\sum_{\substack{t>\theta \\ t \in \sigma}} f_0(t) N_0 / N_\sigma}{f_0(\theta) N_0 / N_\sigma} = \frac{\sum_{\substack{t>\theta \\ t \in \sigma}} f_0(t)}{f_0(\theta)} \leq \frac{\sum_{t>\theta} f_0(t)}{f_0(\theta)} = \frac{1 - F_0(\theta)}{f_0(\theta)} \quad (\text{A7})$$

for θ in the σ^{th} block. The inequality is strict for all blocks except block J . As noted before, there is no distortion for the top type within any population or sub-population; hence, each type S enjoys efficient consumption. Except for S_J , this represents a strict increase in efficiency. For S_J (*i.e.*, K), there is no change in efficiency. For the other types, recall that the marginal profit function is proportional to

$$b_\theta(x) - \frac{1 - F_\sigma(\theta)}{f_\sigma(\theta)} \{b_{\theta+1}(x) - b_\theta(x)\}.$$

From (A7), a switch from $\sigma = 0$ (*i.e.*, privacy) to $\sigma \in \{1, \dots, J\}$, cannot lower the marginal profit function and, for $\sigma \in \{1, \dots, J-1\}$, it strictly increases it. Hence, the consumption of these types increases, at least weakly. Because at least some types consume an amount strictly more efficient under the partition than under privacy, while no type consumes an amount that is less efficient, the result follows when condition (i) is satisfied.

Now assume condition (ii) is met. As above, the consumption of types $\{S_1, \dots, S_{J-1}\}$ is efficient and, thus, more efficient than under privacy. The consumption level of $S_J = K$ is equally efficient with or without privacy. Consider a type θ that is not the maximal type within its block.

Let type $\theta + k$ be the next higher type within that block. If $k = 1$ for all such types, then we have an ordered partition and the previous analysis applies. Restrict attention to the case $k \geq 2$. Now the marginal profit function is proportional to

$$b_\theta(x) - \frac{1 - F_\sigma(\theta)}{f_\sigma(\theta)} \{b_{\theta+k}(x) - b_\theta(x)\},$$

which, by (A7) and the fact that $k \geq 2$, is weakly greater than

$$b_\theta(x) - \frac{1 - F_0(\theta + k - 1)}{f_0(\theta)} \{b_{\theta+k}(x) - b_\theta(x)\}.$$

If we can show that term is greater than

$$b_\theta(x) - \frac{1 - F_0(\theta)}{f_0(\theta)} \{b_{\theta+1}(x) - b_\theta(x)\},$$

then we will have shown that marginal profit weakly increases for all non-maximal types, which completes the proof. A sufficient condition for that relation to hold is that

$$Z(n) \equiv (1 - F_0(\theta + n - 1)) \{b_{\theta+n}(x) - b_\theta(x)\}$$

be decreasing in n . Observe that

$$\begin{aligned} Z(n+1) - Z(n) &= (1 - F_0(\theta + n)) \{b_{\theta+n+1}(x) - b_{\theta+n}(x)\} - f_0(\theta + n) \{b_{\theta+n}(x) - b_\theta(x)\} \\ &\leq f_0(\theta + n) \Delta(x) \left(\frac{1 - F_0(\theta + n)}{f_0(\theta + n)} - n \right), \end{aligned}$$

where $\Delta(x) \equiv b_{\theta+n+1}(x) - b_{\theta+n}(x) > 0$ (because $x < x_{\theta+n+1}^w$) and the inequality follows from assumption (b). Assumption (c) and condition (ii) imply that the term in large parentheses is negative. ***Q.E.D.***

Turn now to our labor market example under the assumption that there is a continuum of household abilities. A necessary condition for a wage, w_σ , to be an equilibrium is that it equal the average productivity of the households willing to work at that wage, or

$$w_\sigma = v \int_{\underline{\theta}}^{w_\sigma} \frac{\theta f_\sigma(\theta)}{F_\sigma(w)} d\theta. \quad (\text{A8})$$

A necessary and sufficient condition for equilibrium is that the wage be the highest such wage satisfying equation (A8)—otherwise there would exist a wage such that the average productivity of job applicants would be higher than that wage and an employer would find it profitable to deviate by offering that wage. We know that such a value exists because both sides of the equation are continuous and the left-hand side ranges from 0 to ∞ while the right-hand side is bounded below by $v\underline{\theta}$ and above by $v\bar{\theta}$.⁴⁴

Proof of Proposition 7: It suffices to show that $\frac{f_i(\theta)}{f_j(\theta)}$ monotonically decreasing implies

$w_i < w_j$. Define $g_\sigma(\theta, w) = \frac{f_\sigma(\theta)}{F_\sigma(w)}$. By the definitions of $f_\sigma(\theta)$ and $F_\sigma(\theta)$,

$$\int_{\underline{\theta}}^w \left[\frac{g_i(\theta, w)}{g_j(\theta, w)} - 1 \right] g_j(\theta, w) d\theta = 0.$$

Moreover, $\frac{g_i(\theta, w)}{g_j(\theta, w)}$ is monotonically decreasing, so it and θ covary negatively. Hence, their covariance,

⁴⁴ We ignore the pathological cases in which the largest solution is a tangency or there is no largest solution because there are infinitely many solutions. $v < 2$ and $F_\sigma(\theta)$ weakly concave are sufficient to rule out such cases.

$$\int_{\underline{\theta}}^w \theta \left[\frac{g_i(\theta, w)}{g_j(\theta, w)} - 1 \right] g_j(\theta, w) d\theta, \quad (\text{A9})$$

is negative for all $w > \underline{\theta}$. Therefore,

$$v \int_{\underline{\theta}}^{w_i} \frac{\theta f_j(\theta)}{F_j(w_i)} d\theta > v \int_{\underline{\theta}}^{w_i} \frac{\theta f_i(\theta)}{F_i(w_i)} d\theta = w_i .$$

Both the integral and w are continuous in w , and $v \int_{\underline{\theta}}^w \frac{\theta f_j(\theta)}{F_j(w)} d\theta - w$ goes to $-\infty$ as $w \rightarrow \infty$.

Hence, there exists $w_j > w_i$ such that $v \int_{\underline{\theta}}^{w_j} \frac{\theta f_j(\theta)}{F_j(w_j)} d\theta = w_j$. **Q.E.D.**

Proof of Proposition 10: Consider an equilibrium that arises when firms can compel households to reveal the values of their indicator variables. Let $t_n(\sigma)$ be the equilibrium wage offered by firm n to any household that has value σ . Because there are no matching benefits, the industry is competitive, and no firm would, in equilibrium, offer a wage that leads it to lose money it must be that $t_n(\sigma) = t_m(\sigma)$ for all n and m . Without loss of generality, order the σ s so that $t(\sigma_i) \geq t(\sigma_j)$ if $i < j$. Observe that σ_1 is the least-favored indicator variable because $u(\theta, x(\theta, t(\sigma_1)), t(\sigma_1)) \leq u(\theta, x(\theta, t(\sigma_j)), t(\sigma_j))$ for all θ and all $j > 1$, where $x(\theta, t)$ is the best-response labor supply decision of a type- θ household when offered t . The result then follows from Proposition 2. **Q.E.D.**

REFERENCES

- Acquisti, Alessandro and Hal R. Varian (2005) "Conditioning Prices on Purchase History," *Marketing Science*, in press.
- Akerlof, George A. (1970) "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism," *The Quarterly Journal of Economics*, **84(3)**: 488-500.
- Calzolari, Giacomo and Alessandro Pavan (2004) "On the Optimality of Privacy in Sequential Contracting," unpublished manuscript, Northwestern University.
- Coase, Ronald (1960) "The Problem of Social Cost," *Journal of Law and Economics*, **3(1)**: 1-44.
- Grossman, Sanford J. (1981) "The Informational Role of Warranties and Private Disclosure about Product Quality," *Journal of Law & Economics*, **24(3)**: 461-83.
- Hermalin, Benjamin E. and Avery W. Katz (2005) "The Law & Economics of Contracts," *The Handbook of Law and Economics* (A.M. Polinsky & S. Shavell, eds.), Amsterdam: North-Holland.
- Hirshleifer, Jack (1971) "The Private and Social Value of Information and the Reward to Inventive Activity," *The American Economic Review*, **61(4)**: 561-574.
- Hirshleifer, Jack (1980) "Privacy: Its Origin, Function, and Future" *The Journal of Legal Studies*, **9(4)**: 649-664.
- Kahn, Charles M., James McAndrews, and William Roberds (2000) "A Theory of Transactions Privacy," Working Paper 2000-22, Federal Reserve Bank of Atlanta.
- Katz, Michael L. (1983) "Non-uniform Pricing, Output and Welfare under Monopoly," *Review of Economic Studies*, **L**: 37-56.
- Levin, Jonathan (2001) "Information and the market for lemons," *RAND Journal of Economics*, **32**: 657-666.
- Lipsey, Richard G., and Kelvin Lancaster (1956) "The General Theory of Second Best," *Review of Economic Studies*, **24**: 11-32.
- Milgrom, Paul R. (1981) "Good News and Bad News: Representation Theorems and Applications," *The Bell Journal of Economics*, **12**: 380-391.
- Milgrom, Paul R. and John Roberts (1990) "Rationalizability, Learning, and Equilibrium in Games with Strategic Complementarities," *Econometrica*, **58**: 1255-1277.
- Mirrlees, James. (1971) "An Exploration in the Theory of Optimal Income Taxation," *Review of Economic Studies*, **38**: 175-208.

- Murphy, Richard S. (1996) "Property Rights in Personal Information: An Economic Defense of Privacy," *Georgetown Law Journal*, **84**:2381-2417.
- Odlyzko, Andrew (2003) "Privacy, economics, and price discrimination on the internet," ACM, *Fifth International Conference on Electronic Commerce*.
- Posner, Richard A. (1981) "The Economics of Privacy," *American Economic Review*, **71(2)**: 405-409.
- Shapiro, Carl, and Hal R. Varian (1997) "US Government Information Policy," unpublished manuscript, University of California, Berkeley.
- Spence, A. Michael (1980) "Multi-product quantity-dependent prices and profitability constraints," *Review of Economic Studies*, **47**: 821-841.
- Stigler, George J. (1980) "An Introduction to Privacy in Economics and Politics," *The Journal of Legal Studies*, **9(4)**:623-644.
- Smith, Marcia S. (2003) "Internet Privacy: Overview and Pending Legislation," Report for Congress, Congressional Research Service, The Library of Congress, updated February 6, 2003.
- Taylor, Curtis R. (2004) "Privacy and Information Acquisition in Competitive Markets," unpublished manuscript, Duke University.
- Varian, Hal R. (1989) "Price Discrimination" in *Handbook of Industrial Organization*, Vol. 1, R. Schmalensee and R. Willig, eds., Amsterdam: North-Holland.
- _____ (1997) "Economic Aspects of Personal Privacy," in *Privacy and Self-regulation in the Information Age*, Washington, D.C.: U.S. Department of Commerce.
- Wathieu, Luc. (2002) "Privacy, Exposure and Price Discrimination," Harvard Business School Marketing Research Papers, No. 02-03.
- Williamson, Oliver (1975) *Markets and Hierarchies: Analysis and Antitrust Implications*, New York: The Free Press.

Figure 1




Figure 2

