

Privacy Regulation and Market Structure

James Campbell,^{*} Avi Goldfarb[†] and Catherine Tucker^{‡§}

September 27, 2011

Abstract

This paper models how regulatory attempts to protect the privacy of consumers' data affect the competitive structure of data-intensive industries. Our results suggest that the commonly used consent-based approach may disproportionately benefit firms that offer a larger scope of services. Therefore, though privacy regulation imposes costs on all firms, it is small firms and new firms that are most adversely affected. We then show that this negative effect will be particularly severe for goods where the price mechanism does not mediate the effect, such as the advertising-supported internet.

^{*}Rotman School of Management, University of Toronto, Canada

[†]Rotman School of Management, University of Toronto, Canada

[‡]MIT Sloan School of Management, MIT, Cambridge, MA and National Bureau of Economic Research.

[§]We thank Alex White, Michael Grubb and Torben Stühmeier and seminar participants at the 2011 IIOC and the University of Toronto for useful comments. All errors are our own.

1 Introduction

Firms now automate, parse, and collect customer data at an unprecedented rate. Many firms, from search engines like Google to credit card companies like Capital One, have realized considerable profits on the basis of the ability to analyze massive amounts of customer data in order to improve their offerings. This leads to two concerns from a regulatory perspective. First, data-intensive operations can lead to natural economies of scale and, on many occasions, network effects. This may generate market power and monopoly (Heyer et al., 2009). Second, data-intensive operations can lead to concerns about privacy. In this paper, we build a theoretical model to ask how attempts to protect consumer privacy can affect the competitive structure of such industries.

In a world with no transaction costs, one might expect privacy regulation to favor small firms over large ones - if data generates economies of scale, then reduced access to data might help to overcome such effects. However, this ignores that most privacy regulation requires firms to obtain one-time individual consumer consent to use consumer data (rather than the consent requirements increasing with the amount of data used). Therefore, privacy regulation imposes transaction costs that, our model suggests, will fall disproportionately on smaller firms. Consequently, rather than increasing competition, the nature of transaction costs implied by privacy regulation suggests that privacy regulation may be anti-competitive.

Specifically, we build a model of competition between a generalist firm offering products that appeal to a variety of consumer needs and a specialist firm offering a product that serves fewer consumer needs. The specialist firm offers higher-quality content, but only for their particular niche. A firm's profits depend on how many customers they attract. Customer data helps both generalist and specialist firms to optimize product offerings. The revenue per customer is higher when firms can leverage customer data. Table 1 gives examples of industry settings to which our model could apply. In each case, larger generalist and smaller specialist firms leverage customer data to increase the profits per customer.

In our model, without privacy regulation, both the generalist and the specialist use freely

Table 1: Examples of Generalist and Specialist Firms in Customer Data-Intensive Industries

Generalist Firm	Specialist Firm	Sample Use of Customer Data	Profit Implication of Data Usage
Google News	CelebrityBuzz.com	Use earlier web browsing behavior to target online banner advertising.	More revenues from advertising
HMO	Specialist clinic	Access digital medical histories	Lower record-keeping costs
Visa credit card	Department store credit card	Perform credit checks to determine credit-worthiness	Fewer defaults
Facebook	Pedal Room (social network site for cyclists)	Consumers create personalized content	Increased customer usage and advertising revenue

available data to optimize their product offerings, and under general circumstances, consumers use both services.

We model privacy regulation as meaning that there is now a cost associated with obtaining a consumer’s consent to use their data. This reflects the costs imposed by current consent requirements in the EU’s Data Protection Directive (95/46/EC) and Privacy and Electronic Communications Directive (2002/58/EC) as well as the language in proposed US privacy regulation (Corbin, 2010; FTC, 2010).

We show that such privacy regulation can preclude profitable entry by the specialist firm. Under regulation, the extra costs required to obtain consent mean that in some cases where entry had been profitable without regulation, the specialist firm will choose not to enter. The generalist firm then captures the whole market. This implies that privacy regulation can increase the advantage enjoyed by a large generalist firm. This deprives consumers of the higher-quality niche product offered by a specialist firm, which represents a loss that must be balanced against any gain to consumers due to the increased privacy. This basic model also implies that if the generalist is sufficiently strong relative to the specialist, consumers are willing to tolerate greater exploitation of data by the generalist than the specialist.

We extend this basic model in two ways. First, we show that the impact of regulation is strongest in industries with little price flexibility. This would be more likely to be the case for digital goods such as the advertising-supporting internet where consumers traditionally do not pay a price for the service. Second, we allow for investment in quality and show that investment is lower under regulation if the entrant's initial quality is relatively low. This relationship between pricing, quality, and privacy builds on Acquisti and Varian (2005) and Fudenberg and Villas-Boas (2006), who emphasize behavioral-based price discrimination.

Overall, our model suggests that privacy regulation can alter the competitive market structure of data-intensive industries. The idea that regulation designed to protect consumers can entrench incumbents has a long history in industrial organization. For example, Farr et al. (2001) and Clark (2007) argue that advertising restrictions (for cigarettes and children's breakfast cereals respectively) benefited the existing producers. Similarly, an average price cap regime can act as a powerful source of entry deterrence (Armstrong and Sappington, 2007). However, to our knowledge we are the first to study how attempts to regulate privacy can affect market structure.

The small literature on privacy in economics has focused on questions of allocative efficiency (Taylor, 2004; Hermalin and Katz, 2006). An earlier legal debate has evaluated the interaction of privacy and antitrust concerns, but from a very different perspective. This literature, for example Edwards (2008), has argued that privacy considerations should be part of antitrust deliberations such as the Google and Doubleclick merger. Gray (2010) argues that privacy regulations will encourage competition in the level of privacy-protection offered by firms. Our paper comes to a very different conclusion: privacy protection can lead to antitrust concerns. In this we echo a recent paper by Commissioner Brill of the FTC who warns that self-regulation in privacy could actually stifle market entry (Brill, 2011).

Our model contributes to the ongoing debate into the appropriate level of governmental privacy protection for consumers in an age of digital data. The debate about privacy regulation has focused on the tradeoff between protection of consumer privacy and continued innovation in a vibrant industry (Miller and Tucker, 2009; Goldfarb and Tucker, 2011; FTC, 2010; Department

of Commerce, 2010). For example, when the White House Council launched a Subcommittee on Privacy and Internet Policy on Oct 24, 2010, it gave it the differing objectives “to promote innovation and economic expansion, while also protecting the rule of law and individual privacy” (Kerry and Schroeder, 2010). Therefore, our results suggest that in addition to concerns about consumer protection and data-driven innovation, the impact on market structure should be an important part of the discussion on privacy regulation.

Next, we formalize the model outlined above.

2 The model

The customer data-intensive markets described in Table 1 have many subtle differences. The aim of the model is to capture the key features of multiple markets rather than perfectly matching any single setting. We use the advertising-supported internet as a motivating example because it has been recently the target of regulation.

Consumers have a unit interval of needs. In the case of websites, they satisfy these by visiting (advertising-supported) sites with different content. There is a single general interest firm (W) which supplies content across the whole unit interval. A potential entrant, specialist firm (S), produces content across a fraction $1 - \alpha$ of the unit interval of content, $\alpha \in [0, 1]$. This fraction $1 - \alpha$ represents the breath of the niche covered by the specialist relative to the full spectrum of consumer interests. The quality of content produced by the specialist in its niche is q_H , and the quality of content produced by the generalist across the whole unit interval is $q_L < q_H$. That is, the generalist does everything fairly well, but the specialist is better than the generalist over that part of the unit interval it covers. For example, a consumer may use a large news aggregator to consume news, but he also may be especially interested in celebrity gossip and so visits a specialist site to consume celebrity news.

Consumers choose how to address their interests by picking a basket X of firms whose product they will adopt, $X \subseteq \{W, S\}$. Consumers prefer higher quality: the value v that a representative

consumer C derives from choosing both content providers, only the generalist content, and only the specialist content respectively, are:

$$v_C(W, S) = (1 - \alpha)q_H + \alpha q_L \quad (2.1)$$

$$v_C(W) = q_L \quad (2.2)$$

$$v_C(S) = (1 - \alpha)q_H \quad (2.3)$$

That is, the consumer can obtain one of q_L across the range 1 , q_H across the range $1 - \alpha$, or q_H across $1 - \alpha$ and q_L across the remaining α . Fix $v_C(\emptyset) = 0$.

Firms want revenues, which depend on its share of customers and the extent to which they use prior customer data to streamline and optimize their offering for these consumers. If a firm is not chosen by the consumer, it earns zero revenue. If a firm is chosen, it earns revenue according to (i) the share of the representative consumer's interests it satisfies (as described above) and (ii) whether it has chosen to use customer data to streamline the operations that support its offering. This is captured in a choice by each firm of $A \in \{T, U\}$, where a data-enhanced product (T) gives revenue of p_T per consumer, and a non-data enhanced product (U) gives revenue of $p_U < p_T$. In the advertising case, p_T represents the revenue per impression of targeted advertising T , and p_U represents the revenue per impression of untargeted advertising U .¹

The revenue earned by a firm j depends on its revenue per consumer and on its market share, which is the share of the representative consumer's interests that it satisfies (α for the general interest firm and $(1 - \alpha)$ for the specialist). Specifically, the revenue earned by a firm that satisfies some arbitrary share β of the representative consumer's interests is given by:

$$R_j = p_A \beta \quad (2.4)$$

For each firm that the consumer chooses to allow to use her data, she incurs a cost d . Her total

¹When applied to the advertising-supported internet this higher payoff to targeted advertising is consistent with the models of this industry used by Athey and Gans (2010), Bergemann and Bonatti (2011), and Johnson (2009).

payoff is given by

$$U_C(X) = v_C(X) - nd, \tag{2.5}$$

where n is the number of firms in X that ask for consent in order to use consumer data to streamline their operations. Assume that $q_L > d$, so that $U_C(W) > U_C(\emptyset)$.

This cost d is a fixed one-time cost consumers incur when giving explicit consent for their data to be used. This requirement to give explicit consent is consistent with the EU's Data Protection Directive (95/46/EC) and the European Privacy and Electronic Communications Directive (2002/58/EC) as well as the language in proposed US privacy regulation (Corbin, 2010; FTC, 2010). One notable feature about newly proposed privacy regulation in both Europe and the US is the demand that consent be explicit, and require explicit action on the part of consumers, rather than being passive.

Our assumption that such consent represents a fixed cost reflects a consistent empirical finding in the privacy literature that requiring explicit opt-in consent deters consumers relative to a setting where no explicit consent is required. For example, Junghans et al. (2005) show in a randomized trial that obtaining opt-in consent relative to opt-out substantially reduced recruitment for a patient study. There are several theories about why requiring explicit opt-out consent deters consumers. One explanation is that opt-in consent forces a consumer to spend time familiarizing themselves with the data policies and agreeing to them. These costs are not inconsiderable: McDonald and Cranor (2008) estimate that it would take 201 hours each year for the typical American internet user to read privacy policies prior to consenting. Furthermore, even if consumers do not invest time familiarizing themselves with privacy policies, the interruption in the experience engendered by having to opt-in could impose costs that deter consumers from proceeding. For example, Lambrecht et al. (2011) show that interruptions caused by a need to ensure data privacy and security in the adoption process can deter consumers.

Implicit in this model is the assumption that consumers are indifferent over whether their data

is used to enhance the product, outside of the direct cost of explicit consent. This is most realistic in cases where the benefits and costs of data collection and usage are not directly visible or obvious to consumers. There is evidence that this is often the case in online commercial environments (Turow et al., 2005). Even if consumers were aware of their data being used, it is not clear in which direction, if any, it would affect their behavior. To the extent that the data is useful in providing consumers with improved services, consumers might be eager to facilitate data sharing. Alternatively, regulation may itself reflect existing consumer distaste for data use. In order to focus on competitive effects, we abstract away from these other potential costs and benefits of regulation.

We compare the case where there is privacy regulation to the case where there is no privacy regulation. In our model, privacy regulation mandates that if a firm collects and uses data on its users to improve its operation, it must obtain explicit consent. We model this requirement such that if privacy regulation is in place, when a firm chooses a data-enriched product T , then the consumer incurs the cost d to adopt the firm's product, while absent regulation there would be no such cost.²

We can represent this setting as an extensive form-game G . Because $U_C(W) > U_C(\emptyset)$ when the generalist uses customer data, the outcome when the specialist does not enter is fixed for simplicity at the generalist using customer data to optimize its operations and the consumer choosing $X = \{W\}$ ³. The order of play is as follows:

1. The specialist chooses whether to enter (E) at a fixed cost F or stay out (O).
2. If the specialist enters, both choose whether or not to use customer data to enrich their product offering (T or U).
3. The consumer C chooses a basket X .

²In Section 4 we discuss how the structure of consent-gathering may influence the nature of the cost d and its implications.

³In a trivial variation of the model in which, after the specialist does not enter, the generalist must choose whether to use customer data and then the consumer must choose a basket, this is the unique equilibrium in the subgame following the specialist staying out. Because it does not impact the core results of the model, we implement it as an assumption to simplify the strategy space for the generalist and the consumer.

This is represented by the following game tree in Figure 1, in which payoffs at terminal nodes are for the specialist, the generalist, and the consumer respectively. The objects of interest will




Figure 1: Extensive form of G

be the subgame perfect equilibria (SPE) of this game. Note that the second stage sees the two firms simultaneously choose whether to play T or U , with payoffs determined by the consumer's response in stage 3, which in turn depends on parameters; below we explicitly consider the normal form of this subgame. A strategy for the specialist is $\sigma_S \in \{(O), (E, A \in \{T, U\})\}$. A strategy for the generalist is $\sigma_W = A \in \{T, U\}$. A strategy for the consumer is $\sigma_C = X \subseteq \{W, S\}$. Total payoff for the consumer is $U_C(X)$, for the generalist R_W , and for the specialist $R_S - F$ if it enters and zero otherwise.

2.1 No regulation benchmark

First consider the case when firms are not required to obtain consumers' consent to use prior customer data to enhance their operations. Consumers will therefore not incur consent costs. Consider the product offering choice subgame g that begins after the specialist firm chooses to

enter. If both firms exist, consumers will trivially choose to consume both ($\sigma_C^g = \{W, S\}$) regardless of product offering choice, because no consent cost is payable.

The subgame g can therefore be represented by payoff matrix in Figure 2. Choosing a data-

		Specialist	
		T	U
Generalist	T	$\alpha p_T, (1 - \alpha)p_T$	$\alpha p_T, (1 - \alpha)p_U$
	U	$\alpha p_U, (1 - \alpha)p_T$	$\alpha p_U, (1 - \alpha)p_U$

Figure 2: Product offering subgame g , no regulation case

enriched product by playing T is a dominant strategy for both firms, and so both play T in equilibrium in g . Payoffs in equilibrium in g are then as follows:

$$U_C = (1 - \alpha)q_H + \alpha q_L \quad (2.6)$$

$$R_W = \alpha p_T \quad (2.7)$$

$$R_S = (1 - \alpha)p_T \quad (2.8)$$

We can then characterize equilibria in the full game.

Theorem 1. *There exists a unique SPE in G in which:*

$$\sigma_S^* = \begin{cases} O, & \text{if } F > (1 - \alpha)p_T \\ E, T & \text{if } F < (1 - \alpha)p_T \end{cases} \quad (2.9)$$

$$\sigma_M^* = T \quad (2.10)$$

$$\sigma_C^* = \{W, S\} \quad (2.11)$$

All proofs appear in the appendix. This result says that provided the fixed cost of entry is smaller than the revenue the specialist earns after entry, the specialist enters, both specialist and generalist offer a data-enriched product, and the consumer uses both services.

2.2 Regulation case

Now consider the case in which privacy regulation, as described above, mandates that firms require consumers to sign up before they can use data in their operations. Denote the extensive game G' . Again we proceed by first characterizing the product offering choice subgame g' .




Figure 3: Consumer's optimal basket of sites in the online-advertising subgame, given firm strategies

Figure 3 summarizes the consumer's optimal (U_C -maximizing) choice of products X^* in the

final stage of the game for each possible pair of product-type choices by the firms and the values of parameters. For example, Figure 3a illustrates the case in which both firms choose to offer a data-enriched product. In that case, when the cost to the consumer of giving their consent (d), is sufficiently small, and the quality of the specialist (q_H) is sufficiently high, the consumer chooses to consume both ($X^* = \{W, S\}$). The specialist then earns revenue $(1 - \alpha)p_T$ and the generalist earns αp_T .

These consumer responses thus define payoffs in g' to each firm given a pair of product choices. The payoffs in the matrix for g' analogous to that for g in Figure 2 therefore depend on the parameters q_H , q_L , α and d , because these define the consumer's optimal basket X^* given some pair of choices by the firms in g' .

Given this, the following result characterizes equilibria in the product choice subgame g' :

Lemma 1. a. *When $q_L > (1 - \alpha)q_H$:*

- i. *If $d > \alpha q_L$, there is a unique equilibrium in g' with $(\sigma_W = U, \sigma_S = U)$.*
- ii. *If $d \in ((1 - \alpha)(q_H - q_L), \alpha q_L)$, there is a unique equilibrium with (T, U) .*
- iii. *If $d < (1 - \alpha)(q_H - q_L)$, there is a unique equilibrium with (T, T) .*

b. *When $q_L < (1 - \alpha)q_H$:*

- i. *If $d > (1 - \alpha)(q_H - q_L)$, there is a unique equilibrium with (U, U) .*
- ii. *If $d \in (\alpha q_L, (1 - \alpha)(q_H - q_L))$, there is a unique equilibrium with (U, T) .*
- iii. *If $d < \alpha q_L$, there is a unique equilibrium with (T, T) .*

In all equilibria $\sigma_C = \{W, S\}$.

Equilibrium product offering by each firm depends on the relative strength of the two firms and on the consumer's response. Perhaps most intuitively, if the costs to consumers of giving consent for the use of their data is 'small', the unique equilibrium in the product choice subgame is for both firms to choose a data-enriched product, and if distaste is 'large', the unique equilibrium is for both

to choose a non-data-rich product. For intermediate cases, the unique equilibrium is asymmetric, but in which direction depends on the quality premium and scope of the specialist. If the quality premium and the scope of the specialist are relatively small, then the unique equilibrium has the generalist offering a data-enriched product and the specialist a non-data-rich product.

Figure 4 illustrates these equilibria in the online-advertising choice subgame for various values of the consumer's distaste parameter d in the cases $q_L > (1 - \alpha)q_H$ and $q_L < (1 - \alpha)q_H$.




Figure 4: Lemma 1

2.2.1 Implications for extent of privacy intrusion

Lemma 1 demonstrates that equilibria in the product choice subgame feature the generalist offering a data-enriched product if $d < \alpha q_L$ and the specialist offering a data-enriched product if $d < (1 - \alpha)(q_H - q_L)$. One way to interpret these conditions is as defining the largest tolerable exploitation of data that the consumer will consent to while still using a product in equilibrium. In a slightly different setting in which the consumer's costs from opting in and the firm's revenue from the data-enriched product both increase in the degree to which the firm exploits data on the consumer, these thresholds on d would define the broadest exploitation of data to which the consumer is willing to consent. Under this interpretation, if the generalist is sufficiently strong relative to the specialist, so that $q_L > (1 - \alpha)q_H$, then consumers are willing to tolerate a greater exploitation of data by the

generalist than the specialist. Notably, the use of data by each firm that consumers are willing to tolerate depends not simply on its own quality but also on the strength of its potential competitor.

2.3 The equilibrium effect of regulation on revenue after entry

Because both products are used in equilibria in the product choice subgame, the payoff in the subgame is strictly positive for both firms.

Consider the revenue of the specialist in the product choice subgame. Specialist revenue under regulation is the same as without regulation in those equilibria in which the specialist offers a data-enriched product. These equilibria are characterized by the cost of consent to consumers (d) being ‘small enough’. Further, the threshold below which d is ‘small enough’ is more permissive the larger $(1 - \alpha)$ and better $(q_H - q_L)$ is the specialist.

The converse implies that a specialist that fills a smaller niche and offers a smaller quality premium over the equivalent function of the generalist is more likely to earn lower revenue after entry in the case with regulation than in the case without. For a given fixed cost of entry, lower revenue in the post-entry subgame implies that the conditions for profitable entry become tighter; this foreshadows the result below characterizing those specialists that enter absent regulation and stay out under regulation.

Theorem 2. *There exists a unique SPE in G' in which:*

$$\sigma_S^* = \begin{cases} O, & \text{if } F > (1 - \alpha)p_T \text{ or } F \in [(1 - \alpha)p_U, (1 - \alpha)p_T], d > (1 - \alpha)(q_H - q_L) \\ E, T & \text{if } F < (1 - \alpha)p_T, d < (1 - \alpha)(q_H - q_L) \\ E, U & \text{if } F < (1 - \alpha)p_U, d > (1 - \alpha)(q_H - q_L) \end{cases} \quad (2.12)$$

$$\sigma_M^* = \begin{cases} T & \text{if } d < \alpha q_L \\ U & \text{if } d > \alpha q_L \end{cases} \quad (2.13)$$

$$\sigma_C^* = \arg \max_X U_C(X) \quad (2.14)$$

In equilibrium, this strategy profile results in the consumer choosing the menu $\{W, S\}$ whenever both firms operate in the game with regulation. Intuitively, absent regulation, entrants offer a data-enriched product after entry, and if the content of the firm's product offering has broad enough appeal this generates enough revenue to allow them to profitably enter. With regulation, the consumer's costs of giving consent for their data to be used - the parameter d - defines the equilibrium in the advertising type subgame played between the two firms. Smaller entrants and entrants that offer a smaller quality premium in their niche are less likely to offer a non-data-enriched product in equilibrium after entry. Since a non-data-enriched product generates less revenue, this means that, all else equal, the marginally profitable entrant must be larger than before to overcome the fixed cost of entry.

Therefore we identify which specialists can profitably enter when privacy regulation does not exist but cannot profitably enter in the case with regulation:

Corollary 1. $\sigma_S^*(G) = E, T$ and $\sigma_S^*(G') = O$ if

$$(1 - \alpha)(q_H - q_L) < d, \quad (2.15)$$

$$(1 - \alpha) \in \left[\frac{F}{p_T}, \frac{F}{p_U} \right] \quad (2.16)$$

This follows directly from Theorems 1 and 2. The condition 2.15 is such that the entrant does not offer a data-enriched product in the equilibrium in the post-entry subgame g' , and 2.16 is the condition under which this precludes profitable entry.

The smaller the quality premium, the smaller the potential entrant, the larger the gap between revenues earned by data-rich and non-data-rich products, and the more the consumer incurs costs giving consent, the larger the set of potential entrants precluded from entry by regulation. Without regulation, the entrant need only be better than the incumbent in its niche and sufficiently large to overcome the fixed cost of entry. With regulation, it must be *sufficiently* better relative to the consumer's distaste for signing up, and large enough relative to the premium commanded by the data-enriched product, in order to either freely offer the data-enriched product without being

shunned by the consumer or be profitable despite having to use less lucrative non-data-enriched products. The regulation therefore helps entrench relatively strong generalists.

Figure 5 illustrates equilibrium strategy of the entrant as a function of its quality and the fixed cost it faces. Figure 5a illustrates the case with no regulation and Figure 5b the case with regulation. The area in which the entrant can profitably enter shrinks under regulation, and now includes an area in which the entrant offers a non-data-enriched product. The existence of this latter region suggests the potential for competitive distortion even in the case in which profitable entry is not precluded.

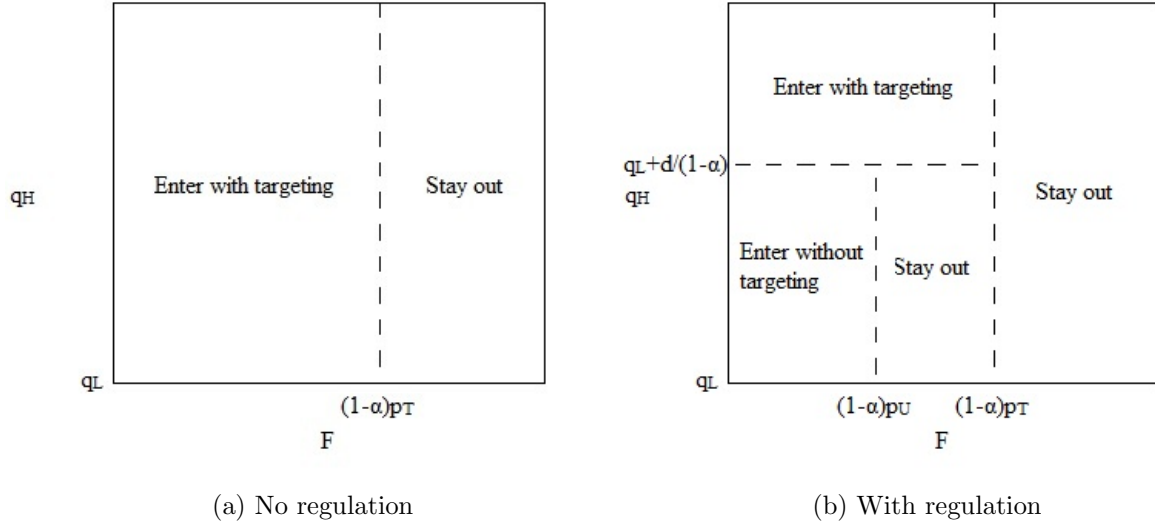


Figure 5: Equilibrium strategy of the entrant for various q_H , F

On the other side of this result is the incumbent. The intuitive notion that the emergence of a large, high quality competitor harms an incumbent is confirmed in the model with or without regulation. It is also the case that if the consumer's distaste for giving consent is sufficiently high then equilibria feature the incumbent choosing non-data-enriched operations in equilibrium and earning a lower payoff than in the case without regulation. It is, however, also true that privacy regulation can shield a large, general incumbent from potential competition because regulation raises the threshold quality and scope for profitable entry by a challenger. In those cases defined in Theorem 1, regulation *benefits* the incumbent, which earns $R_W = p_T$ when the entrant chooses

to stay out, and at best $R_W = \alpha p_T$ when the entrant enters. This is more likely for relatively strong incumbents: the stronger the incumbent, the better the marginal entrant must be.

3 Extensions

3.1 Flexible prices

The results of the previous section demonstrated cases in which the specialist entrant could profitably enter absent regulation but could not profitably enter under regulation. In this section we consider the extent to which this could be mitigated when the firm is able to alter price. For example, the firm may be able to reduce the price of a data-enriched product to compensate consumers for the cost of consent, or raise the price of a non-data-enriched product to make entry profitable.

For simplicity, we focus only on the specialist's problem: fix the generalist's choice to be the data-enriched product T and assume that the consumer will certainly include W in their basket. The order of play is then:

1. The specialist chooses whether to enter (E) at a fixed cost F or stay out (O).
2. If the specialist enters, it chooses whether or not to use customer data to enrich their product offering (T or U).
3. The consumer C chooses between the baskets $\{W\}$ and $\{W, S\}$.

As before we distinguish the case with no regulation, so that the firm need not obtain consent to use T , and the case with regulation in which the firm must obtain consent to use T . Denote the former \hat{G} and the latter \hat{G}' .

Decompose the payoff to using each type of product to $p_T - c_T$ and $p_U - c_U$, where c_T and c_U are the cost of providing each product, $c_T < c_U$, and p_T and p_U now represent a price chosen by

the firm. This means that the revenue earned by the specialist when it uses product $A \in \{T, U\}$ is

$$R_S = (p_A - c_A)(1 - \alpha) \quad (3.1)$$

We can reinterpret the consumer's payoff $U_C(X)$ as maximal willingness to pay for the menu X , and so find the consumer's maximal willingness to pay to add the specialist's product to her basket in each of three cases: under no regulation, under regulation when the specialist uses T , and under regulation when the specialist uses U .

In this modified setting, the analog of Corollary 1 is as follows:

Theorem 3. $\sigma_S^*(\hat{G}) = E, T$ and $\sigma_S^*(\hat{G}') = O$ if

$$c_T + \frac{F}{1 - \alpha} < (1 - \alpha)(q_H - q_L), \quad (3.2)$$

$$c_U + \frac{F}{1 - \alpha} > (1 - \alpha)(q_H - q_L), \quad (3.3)$$

$$c_T + \frac{F}{1 - \alpha} + d > (1 - \alpha)(q_H - q_L). \quad (3.4)$$

The proof appears in the Appendix. Equation 3.2 is the condition for which the specialist can profitably enter absent regulation, and Equations 3.3 and 3.4 are the conditions under which no price that consumers are willing to pay is sufficient to allow profitable entry after regulation with product U and T respectively.

To compare this to the result in Corollary 1 we can rewrite that result's conditions under the decomposition of p_A into $p_A - c_A$. Then Corollary 1 shows that under fixed prices the specialist can no longer profitably enter under privacy regulation if

$$d > (1 - \alpha)(q_H - q_L) \quad (3.5)$$

$$c_U + \frac{F}{1 - \alpha} > p_U. \quad (3.6)$$

Comparing the flexible-prices result to the fixed-prices result suggests that flexible prices make the

region of parameter space for which regulation precludes profitable entry *smaller* but not empty. That is, flexible prices can partly but not fully mitigate the incumbent-favoring effect of the privacy regulation.

This suggests that direct antitrust concerns around privacy regulation may be most acute in an industry with little price flexibility. This is true for the advertising-supported internet (Evans, 2009). Since websites typically offer content to consumers at zero price, their ability to, for example, cut price to compensate a consumer for having to register for the website (and in the process provide informed consent to be tracked for targeted advertisements) is stunted.

3.2 The effect of regulation on investment in innovation

In this section we consider an extension of the framework in which the firms invest in improvements to the quality of the content of the service that they provide to ask how privacy regulation can affect innovation. To do this, we add an investment stage to the beginning of the game G (G' with regulation) so that the quality of the two firms' products now depends on their decisions to invest in quality innovation.

Let the incumbent generalist have a baseline quality q_L . First the specialist entrant will decide whether to pay a fixed cost F , which we now interpret as an investment in innovation yielding some quality of at least q_L , or to stay out. If the entrant invests, next the generalist incumbent will decide whether to invest in its own quality, and then the two firms and the consumer play the game G or G' . That is:

1. The specialist chooses whether to play f and pay F to invest, which yields a random variable $q_S \sim [q_L, \infty)$, or to stay out ($\neg f$).
2. If the specialist chose to invest, the generalist chooses whether to play i and pay I to invest, which yields a random variable $q_W \sim [q_L, \infty)$, or not to invest ($\neg i$) and maintain quality q_L .
3. If the specialist chose to invest, the firms and consumer play G (if no regulation) or G' (if regulation) under the parameters q_S and either q_W or q_L depending on whether the incumbent

chose to invest.

This defines an extensive game as represented by the game tree in Figure 6. As before denote by J and J' the versions without and with regulation respectively.




Figure 6: Extensive form of J

A strategy in J for the specialist is $\sigma_S(J) \in \{f, \neg f\}$, and for the generalist $\sigma_W(J) \in \{i, \neg i\}$, and similarly for J' . Payoffs after f are defined by those from G or G' with quality parameters as resulting from the investment decisions.

Fix $(1 - \alpha) \in [\frac{F}{p_T}, \frac{F}{p_U}]$ so that the entrant can make positive profit only when offering a data-enriched product. The problem for the entrant is whether it will obtain quality in the investment stage sufficiently high for it to be able to offer a data-enriched product and thus to profitably enter.

By assumption $\Pr(q_S > q_L) = 1$. Define the following:

$$\rho \equiv \Pr(q_S > \frac{d}{1 - \alpha} + q_L), \quad (3.7)$$

$$\tau \equiv \Pr(q_S > q_W), \quad (3.8)$$

$$\rho\tau \equiv \Pr(q_S > \frac{d}{1 - \alpha} + q_W), \quad (3.9)$$

where $\rho, \tau \in [0, 1]$. This means that if the entrant invests, with probability 1 it obtains a quality high enough to profitably enter in G if the incumbent does not invest. With probability τ it obtains a quality high enough to profitably enter in G if the incumbent invests. With probability ρ

it obtains a quality high enough to profitably enter in G' if the incumbent does not invest. Finally, with probability $\rho\tau$ it obtains a quality high enough to enter in G' if the incumbent invests.

The object of interest is investment along the equilibrium path. The entrant invests on the equilibrium path when $\sigma_S^*(\cdot) = f$, and the incumbent invests on the equilibrium path if both $\sigma_S^*(\cdot) = f$ and $\sigma_W^*(\cdot) = i$. The following result characterizes the change in investment behavior under the no-regulation and regulation regimes:

Theorem 4. i. $\sigma_W^*(J) = \neg i \Rightarrow \sigma_W^*(J') = \neg i$.

ii. For $\sigma_S^*(J) = \neg f$, $\sigma_S^*(J') = f$, it is necessary but not sufficient that $\rho > \tau$.

iii. $\sigma_W^*(J) = i$, $\sigma_W^*(J') = \neg i$ if

a) $I \in [\rho(1 - \tau)K, (1 - \tau)K]$, or

b) $F \in [\rho\tau K, \tau K]$

iv. $\sigma_S^*(J) = f$, $\sigma_S^*(J') = \neg f$ if

a) $I > (1 - \tau)K$ and $F > \rho K$, or

b) $I < \rho\tau K$ and $F \in [\rho\tau K, \tau K]$

c) $I \in [\rho(1 - \tau)K, (1 - \tau)K]$, $\rho < \tau$ and $F \in [\rho K, \tau K]$

These results demonstrate that the incumbent never invests more under regulation than under no regulation, that the entrant invests more under regulation only if ρ is high, and that the range of parameters for which each firm invests less under regulation is decreasing in ρ . The parameter ρ captures the probability that the entrant's raw quality (absent innovation by either firm) is sufficient to overcome the consumer's costs from sharing their information with a new firm. We can thus view ρ as capturing the quality of the entrant's exogenous 'idea', and so this result mirrors that in Corollary 1: just as privacy regulation raises the threshold quality premium an entrant must offer to profitably enter, it also raises the threshold intrinsic quality such that the entrant,

and in turn the incumbent, will choose to invest in the quality of their product. The regulation therefore increases the difficulty the entrant faces in challenging a relatively strong incumbent.

4 Implications for the design of consent requirements

We have seen that privacy regulation that requires each producer to gain informed consent can disproportionately hinder small and young firms. In the model, the consumer was prompted to consent to sharing her data when choosing whether to adopt a product. In this section we discuss the implications of our results for the design of consent requirements in privacy policy.

One way that regulators could avoid the anti-competitive effects of privacy regulation is to reduce d . This could be achieved via mandated ‘standardized’ privacy agreements that are easy to follow and understand and therefore plausibly reduce the cost to the consumer of consenting to share her data. A smaller d would function in the model to push the marginally profitable entrant in the model under regulation closer to the marginally profitable entrant absent regulation. In this sense standardization of consent requirements would partially mitigate the anti-competitive effect.

Another characteristic of the model is that the consumer incurs a cost d every time she gives consent. An alternative to case-by-case opt-in could be a global opt-in, whereby the consumer consents to sharing certain pieces of data in advance with all producers. If a global opt-in was accepted by a consumer, which depends on the regulator willingness to give up consumer case-by-case sovereignty, it would amount to the consumer paying some fixed cost of consenting to data use before the game modeled above begins. If they did offer global consent, then there would be no further barrier to entry in the regulatory game than the game without regulation. A global opt-in would then have the potential to mitigate the anti-competitive effect.

The challenges to global opt-in are largely practical. One possibility is to give consumers a simple option to opt-in to all tracking. If consumers do not opt-in, then no firm is able to collect data. The difficulty with this is that consumers will have incentives to allow some firms to collect data if the benefits to data are high enough. Therefore, the model largely reflects the

discussion above, with perhaps a higher transaction costs for consumers to opt-in and override their global preferences. Also many existing (and proposed) laws, such as the EU Privacy and Electronic Communications Directive, are impossible to implement via a single ‘opt-in’. The reasons is that such laws state that a sliding scale of consent is required given the privacy risk. For example, ICO (2011) suggests that in order to comply with the ‘EU Privacy and Electronic Communications Directive,’ while some cookies could be acceptable if accepted by a browsers’ settings, other cookies may require pop-up windows requesting consent. This sliding scale means that the regulatory requirements necessarily differ across websites and contexts, forcing a company-specific context. More generally, the issue is that as long as it is legally possible for a consumer to decide to give permission for data use in a firm-specific context, but not globally, the above implications for market structure hold.

5 Conclusion

In this paper we investigate the relationship between offering privacy protection and competitive market structure. Our model brings a new hypothesis to the existing policy discussion surrounding privacy regulation in the United States, Europe, and elsewhere. We show that a potential risk in privacy regulation is the entrenchment of the existing incumbent firms and a consequent reduction in the incentives to invest in quality. These incentives are stronger when firms have little consumer-facing price flexibility, as is the case in online media. We show this result in a setting where large firms have no inherent advantage over small firms in generating trust. If consumers are more likely to trust large firms with data, as suggested by McDonald and Cranor (2008), then consumers might become even less likely to provide consent to small firms, though that depends on the degree to which consumers are aware of tracking absent regulation.

The law literature discusses many possible reasons to favor regulatory protection of consumer privacy (Nissenbaum, 2010; Zittrain, 2008). Thus far, the discussion of the costs of such regulation has focused on data-driven innovation. In this paper, we explore another potential cost of privacy

regulation: it may favor incumbents and large firms over entrants and small firms. This finding goes against current legal debate where the focus has been on the extent to which concerns about privacy should be used as a criterion to reject proposed mergers (Edwards, 2008). Therefore, our model adds a new and potentially important consideration to the debate on privacy regulation.

References

- Acquisti, A. and H. R. Varian (2005). Conditioning prices on purchase history. *Marketing Science* 24(3), 367–381.
- Armstrong, M. and D. E. Sappington (2007). Chapter 27 recent developments in the theory of regulation. Volume 3 of *Handbook of Industrial Organization*, pp. 1557 – 1700. Elsevier.
- Athey, S. and J. S. Gans (2010). The Impact of Targeting Technology on Advertising Markets and Media Competition. *AER Paper and Proceedings* 100(2), 608–613.
- Bergemann, D. and A. Bonatti (2011). Targeting in Advertising Markets: Implications for Offline vs. Online Media. *RAND Journal of Economics Forthcoming*.
- Brill, J. (2011, 7). The intersection of consumer protection and competition in the new world of privacy. *Competition Policy International*.
- Clark, C. R. (2007). Advertising restrictions and competition in the children’s breakfast cereal industry. *Journal of Law and Economics* 50, 757–780.
- Corbin, K. (2010, January 28). Privacy bill nears introduction in house. Technical report, Esecurity, Planet.
- Department of Commerce (2010, December). Commercial data privacy and innovation in the internet economy: A dynamic policy framework. *Report*.
- Edwards, E. (2008). Stepping up to the plate: The Google-DoubleClick merger and the role of the Federal Trade Commission in protecting online data privacy. *Mimeo, Santa Clara University*.
- Evans, D. S. (2009). The online advertising industry: Economics, evolution, and privacy. *The Journal of Economic Perspectives* 23(3), 37–60.
- Farr, S. J., C. Horton Tremblay, and V. J. Tremblay (2001). The welfare effect of advertising restrictions in the U.S. cigarette industry. *Review of Industrial Organization* 18, 147–160. 10.1023/A:1007801608741.
- FTC (2010, December 1). Protecting consumer privacy in an era of rapid change. *Report*.
- Fudenburg, D. and J. M. Villas-Boas (2006). *Volume 1: Handbooks in Information Systems*, Chapter 7: Behavior Based Price Discrimination and Customer Recognition, pp. 377–435. Emerald Group Publishing.
- Goldfarb, A. and C. Tucker (2011). Privacy regulation and online advertising. *Management Science* 57(1), 57–71.
- Gray, E. (2010, September). Ftc to boost competition in privacy. *Global Competition Review*.
- Hermalin, B. and M. Katz (2006, September). Privacy, property rights and efficiency: The economics of privacy as secrecy. *Quantitative Marketing and Economics* 4(3), 209–239.

- Heyer, K., C. Shapiro, and J. Wilder (2009). The year in review: Economics at the antitrust division 2008-2009. *Review of Industrial Organization* 35, 349–367.
- ICO (2011, May). Changes to the rules on using cookies and similar technologies for storing information. *Information Commissioner’s Office*.
- Johnson, J. P. (2009). Targeted advertising and advertising avoidance. *Mimeo, Cornell*.
- Junghans, C., G. Feder, H. Hemingway, A. Timmis, and M. Jones (2005). Recruiting patients to medical research: double-blind randomised trial of opt-in versus opt-out strategies. *BMJ* 331(7522), 940.
- Kerry, C. and C. Schroeder (2010, October 24). White House council launches interagency subcommittee on privacy and internet policy. *Office of Science and Technology Policy Webpage*.
- Lambrecht, A., K. Seim, and C. Tucker (2011). Stuck in the adoption funnel: The effect of interruptions in the adoption process on usage. *Marketing Science* 30(2), 355–367.
- McDonald, A. M. and L. F. Cranor (2008). The cost of reading privacy policies. *Journal of Law and Policy for the Information Society* 4(3).
- Miller, A. R. and C. Tucker (2009, July). Privacy protection and technology adoption: The case of electronic medical records. *Management Science* 55(7), 1077–1093.
- Nissenbaum, H. (2010). *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford University Press.
- Taylor, C. R. (2004). Consumer privacy and the market for customer information. *The RAND Journal of Economics* 35(4), pp. 631–650.
- Turow, J., L. Feldman, and K. Meltzer (2005). Open to exploitation: American shoppers online and offline. *Annenberg Public Policy Center, White Paper*.
- Zittrain, J. (2008). *The Future of the Internet and How To Stop It*. Yale University Press.

A Proofs

A.1 Theorem 1

Proof. $U_i = v_i$ for all σ_W, σ_S since consumers are not required to sign up. Since $v_i(G, S) > v_i(G), v_i(S) > v_i(\emptyset)$, $\sigma_i^* = \{W, S\}$ is strictly dominant for i . The unique equilibrium in g in strictly dominant strategies therefore involves $\sigma^S(g) = T, \sigma^W(g) = T, \sigma^i(g) = \{W, S\}$. $R_W(E, T) = (1-\alpha)p_T$ and so by backward induction $\sigma_S^* = O$ if $(1-\alpha)p_T - F < 0$ and $\sigma_S^* = E, T$ if $(1-\alpha)p_T - F > 0$. \square

A.2 Lemma 1

Proof. Note that $q_L > (1-\alpha)q_H$ iff $\alpha q_L > (1-\alpha)(q_H - q_L)$.

a.iii. and b.iii. Since $d < \min\{\alpha q_L, (1-\alpha)(q_H - q_L)\}$, $X^* = \{W, S\}$ for any $\sigma_W(g'), \sigma_S(g')$ and so $\sigma_W(g') = T, \sigma_S(g') = T$ are dominant strategies for W and S in g' . At $\sigma_W(g') = T, \sigma_S(g') = T, X^* = \{W, S\}$.

a.i and b.i. Since $d > \max\{\alpha q_L, (1-\alpha)(q_H - q_L)\}$, $j \notin X^*$ if $\sigma_j(g') = T$ and so $\sigma_W(g') = U, \sigma_S(g') = U$ are dominant strategies for W and S in g' . At $\sigma_W(g') = U, \sigma_S(g') = U, X^* = \{W, S\}$.

a.ii. Since $d < \alpha q_L, X^* = \{W, S\}$ when $\sigma_W(g') = T, \sigma_S(g') = U$. Since $d > (1-\alpha)(q_H - q_L)$ and $q_L > (1-\alpha)q_H, X^* = \{W\}$ when $\sigma_W(g') = T, \sigma_S(g') = T$. Thus $\sigma_W(g') = T$ is dominant for W in g' , and $BR_S(\sigma_W(g') = T) = U$. At $\sigma_W(g') = T, \sigma_S(g') = U, X^* = \{W, S\}$.

b.ii. Since $d < (1-\alpha)(q_H - q_L), X^* = \{W, S\}$ when $\sigma_W(g') = U, \sigma_S(g') = T$. Since $d > \alpha q_L$ and $q_L < (1-\alpha)q_H, X^* = \{S\}$ when $\sigma_W(g') = T, \sigma_S(g') = T$. Thus $\sigma_S(g') = T$ is dominant for S in g' , and $BR_W(\sigma_S(g') = T) = U$. At $\sigma_W(g') = U, \sigma_S(g') = T, X^* = \{W, S\}$. \square

A.3 Theorem 2

Proof. $R_S \leq (1 - \alpha)p_T$ in all equilibria in g' , so if $F > (1 - \alpha)p_T$ then $\sigma_S(G') = O$ dominates $\sigma_S(G') = (E, \cdot)$. If $d > (1 - \alpha)(q_H - q_L)$, $\sigma_S^*(g') = U$, $X^* = \{W, S\}$ and so $R_S = (1 - \alpha)p_U$ and thus if $F > (1 - \alpha)p_U$ then $\sigma_S(G') = O$ dominates $\sigma_S(G') = (e, \cdot)$.

Conversely if $d < (1 - \alpha)(q_H - q_L)$, $\sigma_S^*(g') = T$, $X^* = \{W, S\}$, $R_S = (1 - \alpha)p_T$ and so if $F < (1 - \alpha)p_T$ then E dominates O and $\sigma_S^*(G') = E, T$. If $d > (1 - \alpha)(q_H - q_L)$, $\sigma_S^*(g') = U$, $X^* = \{W, S\}$, $R_S = (1 - \alpha)p_U$ and so if $F < (1 - \alpha)p_U$ then E dominates O and $\sigma_S^*(G') = E, U$.

$\sigma_W^*(G') = \sigma_W^*(g')$ and $\sigma_C^*(G') = \sigma_W^*(g')$ since W and C have a strategic choice only in subgame g' . □

A.4 Theorem 3

Proof. Under the no-regulation regime, as before S can freely use product T without requiring the consumer to pay the consent cost d to adopt it. The consumer's maximal willingness to pay to add S using T (that is, maximal p_T) is then

$$U_C(W, S) - U_C(W) = [(1 - \alpha)q_H + \alpha q_L] - [q_L] \quad (\text{A.1})$$

$$= (1 - \alpha)(q_H - q_L). \quad (\text{A.2})$$

S can thus profitably enter provided that this p_T is 'large enough':

$$F < R_j = (p_T - c_T)\beta \quad (\text{A.3})$$

$$= ((1 - \alpha)(q_H - q_L) - c_T)(1 - \alpha) \quad (\text{A.4})$$

$$\Rightarrow c_T + \frac{F}{1 - \alpha} < (1 - \alpha)(q_H - q_L) \quad (\text{A.5})$$

Under regulation, again assume that to use T requires the firm to obtain consent from the consumer. Then the consumer's willingness to pay to add S using T to a basket that already

includes W is

$$U_C(W, S) - U_C(W) = [(1 - \alpha)q_H + \alpha q_L - 2d] - [q_L - d] \quad (\text{A.6})$$

$$= (1 - \alpha)(q_H - q_L) - d. \quad (\text{A.7})$$

This, then, is the upper bound on p_T such that the consumer is willing to add the product T from the specialist S . S can thus profitably enter with product T as long as

$$F < R_j = (p_T - c_T)\beta \quad (\text{A.8})$$

$$= ((1 - \alpha)(q_H - q_L) - d - c_T)(1 - \alpha), \quad (\text{A.9})$$

$$\Rightarrow c_T + \frac{F}{1 - \alpha} + d < (1 - \alpha)(q_H - q_L). \quad (\text{A.10})$$

Similarly, the consumer's willingness to pay to add S using U to a basket that already includes W is

$$U_C(W, S) - U_C(W) = [(1 - \alpha)q_H + \alpha q_L - d] - [q_L - d] \quad (\text{A.11})$$

$$= (1 - \alpha)(q_H - q_L). \quad (\text{A.12})$$

This is the upper bound on p_U such that the consumer is willing to add the product U from the specialist S . S can thus profitably enter with product U if

$$F < R_j = (p_U - c_U)\beta \quad (\text{A.13})$$

$$= ((1 - \alpha)(q_H - q_L) - c_U)(1 - \alpha), \quad (\text{A.14})$$

$$\Rightarrow c_U + \frac{F}{1 - \alpha} < (1 - \alpha)(q_H - q_L). \quad (\text{A.15})$$

Together Equations A.5, A.10 and A.15 form the result. \square

A.5 Theorem 4

Proof. The following result defines equilibria in the games J and J' .

Lemma 2. Define $K \equiv (1 - \alpha)p_T$. There exists a unique SPE (σ_S^*, σ_W^*) in J at

$$(f, i) \text{ if } I < (1 - \tau)K, F < \tau K, \quad (\text{A.16})$$

$$(f, \neg i) \text{ if } I > (1 - \tau)K, \quad (\text{A.17})$$

$$(\neg f, i) \text{ if } I < (1 - \tau)K, F > \tau K. \quad (\text{A.18})$$

There exists a unique SPE (σ_S^*, σ_W^*) in J' at

$$(f, i) \text{ if } I < \rho(1 - \tau)K, F < \rho\tau K, \quad (\text{A.19})$$

$$(f, \neg i) \text{ if } I > \rho(1 - \tau)K, F < \rho K, \quad (\text{A.20})$$

$$(\neg f, i) \text{ if } I < \rho(1 - \tau)K, F > \rho\tau K, \quad (\text{A.21})$$

$$(\neg f, \neg i) \text{ if } I > \rho(1 - \tau)K, F > \rho K. \quad (\text{A.22})$$

Proof. First, consider the incumbent. In J , if W is called upon to make a decision and W plays $\neg i$, $q_W \equiv q_L < q_S$ for sure and so $R_W(G) = \alpha p_T$. If W plays i , $q_W > q_S$ with probability $(1 - \tau)$ and so $R_W(G) = \alpha p_T$ with probability τ and $R_W(G) = p_T$ with probability $(1 - \tau)$. The expected payoff to playing i , $\tau(\alpha p_T) + (1 - \tau)p_T$, exceeds the payoff to $\neg i$, αp_T , if $I < (1 - \tau)K$. Thus if $I < (1 - \tau)K$, $\sigma_W^*(J) = i$, else $\sigma_W^*(J) = \neg i$.

For the incumbent in J' , similarly, but now if W plays $\neg i$, $q_W \equiv q_L + \frac{d}{1 - \alpha} < q_S$ with probability ρ ; $R_W(G) = \alpha p_T$ with probability ρ and $R_W(G) = p_T$ with probability $(1 - \rho)$. If W plays i , $q_W + \frac{d}{1 - \alpha} < q_S$ with probability $\rho\tau$; $R_W(G) = \alpha p_T$ with probability $\rho\tau$ and $R_W(G) = p_T$ with probability $(1 - \rho\tau)$. The expected payoff to playing i , $\rho\tau(\alpha p_T) + (1 - \rho\tau)p_T$, exceeds the payoff to $\neg i$, $\rho(\alpha p_T) + (1 - \rho)p_T$, if $I < \rho(1 - \tau)K$. Thus if $I < \rho(1 - \tau)K$, $\sigma_W^*(J') = i$, else $\sigma_W^*(J') = \neg i$.

Next, consider the entrant in J . By backward induction, since $F < K$ by assumption, $\sigma_S^*(J) = f$

when $I > (1 - \tau)K$, since $\sigma_W^*(J) = \neg i$, $q_S > q_W$ for sure and $R_S(G) = K$. If $I < (1 - \tau)K$ so that $\sigma_W^*(J) = i$, $q_S > q_W$ with probability τ . Thus the expected payoff to playing f , $\tau K - F$, exceeds the payoff to playing $\neg f$, zero, if $F < \tau K$, and so $\sigma_W^*(J) = i$, $\sigma_S^*(J) = f$ if $I < (1 - \tau)K$, $F < \tau K$, and $\sigma_W^*(J) = i$, $\sigma_S^*(J) = \neg f$ if $I < (1 - \tau)K$, $F > \tau K$.

Finally, consider the entrant in J' . If $I < \rho(1 - \tau)K$, $\sigma_W^*(J') = i$. If the entrant plays f , $q_W + \frac{d}{1-\alpha} < q_S$ with probability $\rho\tau$, and so in expectation $R_S(G) = \rho\tau K$. If the entrant plays $\neg f$, its payoff is zero. Thus if $I < \rho(1 - \tau)K$ and $F < \rho\tau K$, then $\sigma_W^*(J') = i$, $\sigma_S^*(J') = f$ and if $I < \rho(1 - \tau)K$ and $F > \rho\tau K$, then $\sigma_W^*(J') = i$, $\sigma_S^*(J') = \neg f$.

If $I > \rho(1 - \tau)K$, $\sigma_W^*(J') = \neg i$. If the entrant plays f , $q_W + \frac{d}{1-\alpha} < q_S$ with probability ρ , and so in expectation $R_S(G) = \rho K$. If the entrant plays $\neg f$, its payoff is zero. Thus if $I > \rho(1 - \tau)K$ and $F < \rho K$, then $\sigma_W^*(J') = \neg i$, $\sigma_S^*(J') = f$ and if $I > \rho(1 - \tau)K$ and $F > \rho K$, then $\sigma_W^*(J') = \neg i$, $\sigma_S^*(J') = \neg f$. □

Theorem 4 follows directly. □