Privacy Regulation and Market Structure

Citation	Campbell, James, Avi Goldfarb, and Catherine Tucker. “Privacy Regulation and Market Structure.” Journal of Economics & Management Strategy 24, no. 1 (February 10, 2015): 47–73.
As Published	http://dx.doi.org/10.1111/jems.12079
Publisher	Wiley Blackwell
Version	Original manuscript
Accessed	Tue Nov 20 18:38:25 EST 2018
Citable Link	http://hdl.handle.net/1721.1/99173
Terms of Use	Creative Commons Attribution-Noncommercial-Share Alike
Detailed Terms	http://creativecommons.org/licenses/by-nc-sa/4.0/

Privacy Regulation and Market Structure

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

August 15, 2013

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.

[*] University of Toronto, Toronto ON, Canada

[†] University of Toronto, Toronto ON, 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. We gratefully acknowledge financial support from the NBER Digitization program and the Sloan Foundation. 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 requests increasing with the amount of data used). Therefore, privacy regulation imposes transaction costs whose effects, 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. We focus our discussion on the advertising-supported internet, where larger generalist and smaller specialist firms leverage customer data to increase the profits per customer. In our conclusion, we discuss the application of the ideas to other industry contexts.

In our model, without privacy regulation, consumers suffer harm when consuming any product since rules on data exploitation are ill-defined. Both the generalist and the specialist use freely available data to optimize their product offerings, and under general circumstances, a consumer uses both services if she is not too concerned about privacy and uses no product if she is sufficiently concerned.

We model privacy regulation as meaning that consumers now incur a cost when prompted to give consent to use their data. This reflects the frictions imposed by current consent requirements in the EU's Data Protection Directive (95/46/EC) and Privacy and Electronic Communications Directive (2002/58/EC) and its amendment (2009/135/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 three 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 show that the model's conclusions are robust to there being several potential specialist firms each serving a different niche. Third, we allow for investment in quality and show that the entrant never invests more in quality under regulation than without regulation, and in some cases invests less.

Overall, our model suggests that privacy regulation can alter the competitive market structure of data-intensive industries. The relationship in the model between pricing, quality, and privacy builds on Acquisti and Varian (2005) and Fudenberg and Villas-Boas (2006), who emphasize behavioral-based price discrimination. 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, Armstrong and Sappington (2007) show that an average price cap regime can act as a powerful source of entry deterrence and Armstrong et al. (2009) show that allowing customers to opt-out of marketing can soften price competition. This is echoed in recent work too by Bouckaert and Degryse (2006) who explicitly tie this softening of price competition to opt-in regimes.

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 Julie 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; 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 discuss the history and nature of privacy regulation.

2 Privacy Regulation

Echoing the early literature in economics (Posner, 1981), it is natural to think of privacy regulation as simply being an enforcement of secrecy consequently ensuring that all firms receive less data.

However, the development of privacy regulation both in the US and internationally has evolved into something more complex. Representative of this shift are the privacy principles adopted by the OECD Council in September, 1980. These principles were adopted in response to the nascent advent of the information society and the collection of widespread data by firms about their consumers. The aim of these guidelines was to form a basis for legislation across member states. They are still important today, both because they are commonly used outside of the OECD as the foundation for privacy regulation, but also because they form the basis of the European Commission (EC) Data Protection Directive (Directive 95/46/EC), and other 'EU-style' national privacy legislation which is the basis of Europe-wide privacy regulation.

The guidelines summarize eight key principles:[1]

1. Collection Limitation Principle: There should be limits to the collection of personal data and any such data should be obtained by lawful and fair means and, where appropriate, with the knowledge or consent of the data subject.

2. Data Quality Principle: Personal data should be relevant to the purposes for which they are to be used, and, to the extent necessary for those purposes, should be accurate, complete and kept up-to-date.

3. Purpose Specification Principle: The purposes for which personal data are collected should be specified not later than at the time of data collection and the subsequent use limited to the fulfilment of those purposes or such others as are not incompatible with those purposes and as are specified on each occasion of change of purpose.

4. Use Limitation Principle: Personal data should not be disclosed, made available or otherwise used for purposes other than those specified except: a) with the consent of the data subject; or b) by the authority of law.

5. Security Safeguards Principle: Personal data should be protected by reasonable security safeguards against such risks as loss or unauthorised access, destruction, use, modification or disclosure of data.

[1]Part Two, paragraphs 7 though 14 of Annex to the Recommendation of the Council of 23rd September 1980: Guidelines Governing The Protection of Privacy and Transborder Flows of Personal Data.

6. Openness Principle: There should be a general policy of openness about developments, practices and policies with respect to personal data. Means should be readily available of establishing the existence and nature of personal data, and the main purposes of their use, as well as the identity and usual residence of the data controller.

7. Individual Participation Principle: An individual should have the right: a) to obtain from a data controller, or otherwise, confirmation of whether or not the data controller has data relating to him; b) to have communicated to him, data relating to him within a reasonable time; at a charge, if any, that is not excessive; in a reasonable manner; and in a form that is readily intelligible to him; c) to be given reasons if a request is denied, and to be able to challenge such denial; and d) to challenge data relating to him and, if the challenge is successful to have the data erased, rectified, completed or amended.

8. Accountability Principle: A data controller should be accountable for complying with measures which give effect to the principles stated above.

It is clear that rather than simply limiting the use of data, what these regulations attempt to do is add transparency to the process by which firms use data. In particular, the focus of this paper is the general principles of consent, knowledge, and control that are evoked by these principles.

Nowhere has this emphasis been more visible than in the recent enactment of Cookie tracking regulation (2009/135/EC). In 2009, the EU updated the original 2002 E-privacy directive, (2002/58/EC), that was designed to reflect the more general (1995) EU privacy directive (and by implication the OECD principles). Much of this update reflected the fact that in the previous decade there were many technologies that enhanced online advertising taking it from a medium that largely displayed electronic billboards on websites to an industry that tracked users browsing behavior and used this information to both target ads to precisely defined groups in real-time and personalize the content of ads.

One of the most noticeable features of the revised cookie directive was the passage

> Member States shall ensure that the storing of information, or the gaining of access to information already stored, in the terminal equipment of a subscriber or user is only

allowed on condition that the subscriber or user concerned has given his or her consent, having been provided with clear and comprehensive information.

In other words if a website that was supported by the internet wished to place cookies in order to better serve and target ads it would need to obtain consent. This directive has received attention for the fact it appears to make the use of Cookies an opt-in, rather than an opt-out process.

The implications of this Cookie directive are probably best illustrated by an example. For this, we use a simulation provided by TrustE, a well-known provider of software solutions that make websites compliant with European Privacy regulations that require express consent for the user of cookies. Figures 1, 2, and 3 provide an example of how the JEMS website would have to alter in order to comply with the kind of expressed consent mechanism that is required by the directive. In Figure 1 the user has to click on the top banner in order to agree to the use of Cookies. If the user clicks to allow cookies, Figure 2 shows that the default is for the user to agree to all cookies put there by the website. If they disagree and wish to 'opt-out', as shown in Figure 3 the way the software works is to allow them to limit their use of cookies so that they are purely functional - such as those cookies used for tracking the contents of a shopping basket within a single session.

According to a December 2012 estimate, the cost for a typical website of automating a pre-packaged cookie management solution is $3300 for 15 months.[2] Of course whether this is a large amount is an open question and we do not explicitly model these costs in our approach - we will note though that it is reasonably straightforward to show that, if all fi rms have to pay a fee of $3300 as a result of regulation, this would hit smaller firms disproportionately in a way that would reinforce our basic findings.

An obvious question is the extent to which this evolution of European law will spread to North America. Currently the harmonization of laws works through the Safe Harbor self-certifying framework which allows firms in the US who do business in Europe to signal their compliance. However, a key question that is debated in the US currently is whether the US should follow EU law regarding the regulation of consent for online surveillance. This paper aims to help shed light on that debate.

[2] http://www.socitm.net/info/165/services/31/website_services/6

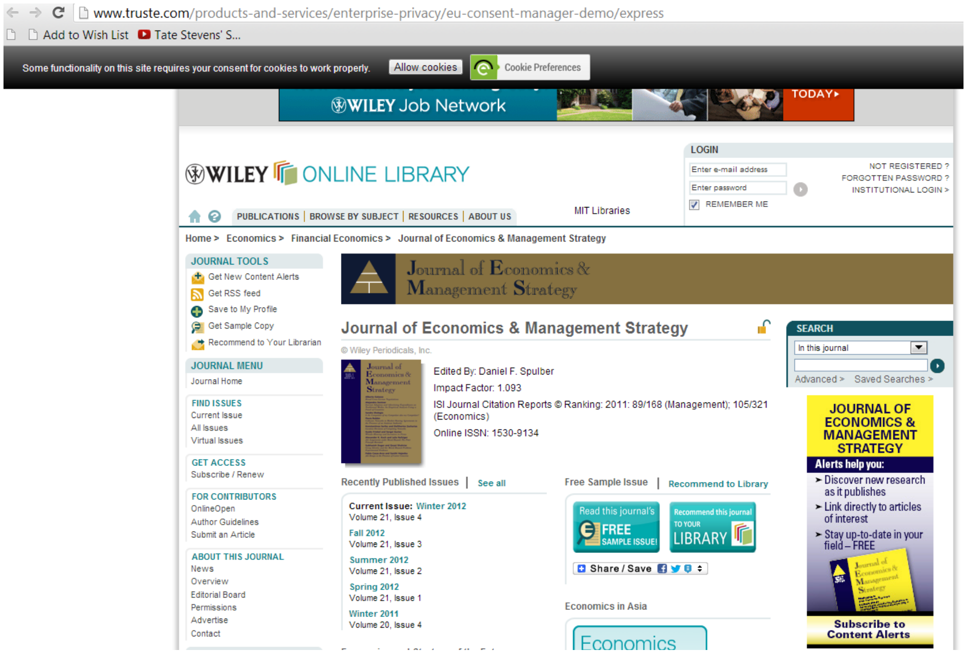

Figure 1: Express Consent

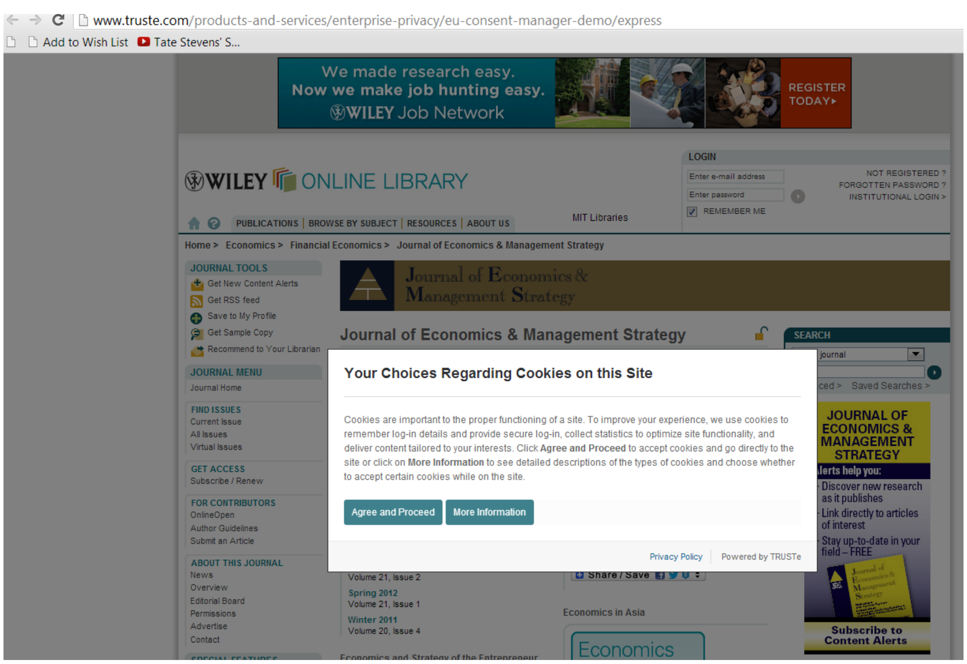

Figure 2: Express Consent

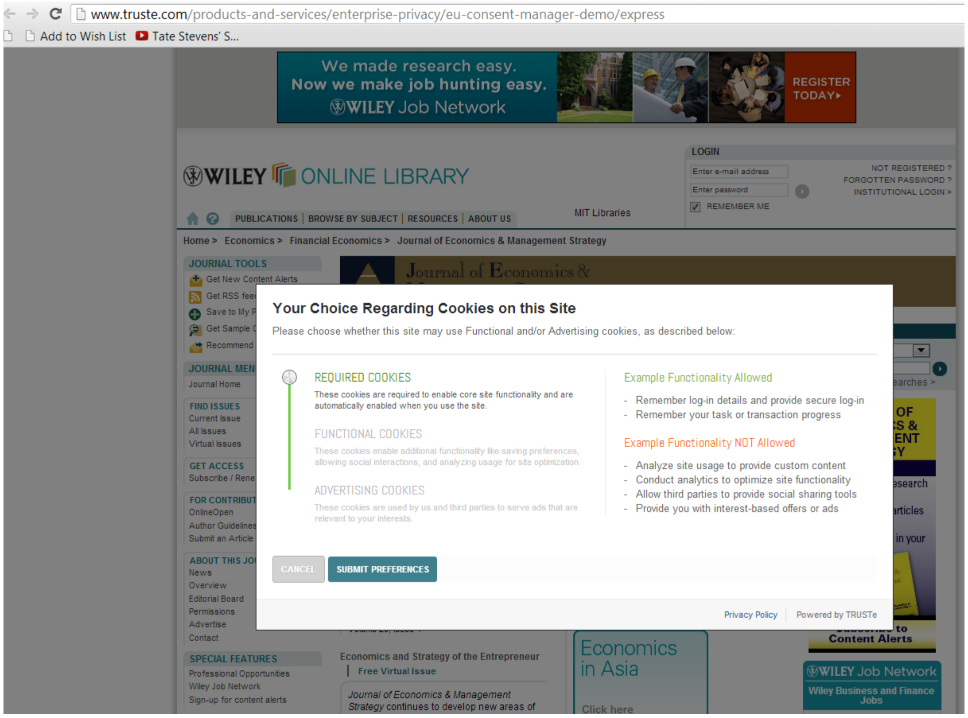

Figure 3: Express Consent

3 The Model

We next formalize how to model this kind of regulation in the context of the advertising-supported internet. We recognize that there are many different sectors, such as health and finance, where privacy regulation is being contemplated. We discuss how our results may apply to these different settings in Section 6.

There is a unit mass of consumers who each have a unit interval of needs. In the case of websites, consumers satisfy these by visiting (advertising-supported) sites with different content. There is a single incumbent firm that supplies content across the whole unit interval; denote this firm G to represent that it is a 'generalist'. A potential entrant produces content across a fraction $1-\alpha$ of the unit interval of content, $\alpha \in [0,1]$; denote this firm S to represent that it is a 'specialist'. The fraction $1-\alpha$ represents the breadth covered by the potential entrant relative to the full spectrum of consumer interests. The quality of content produced by the specialist in its niche is q_S, and the quality of content produced by the generalist across the whole unit interval is $q_G < q_S$. 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.[3]

Each consumer chooses how to address her interests by picking a basket X of firms whose product she will adopt, $X \subseteq \{G, S\}$. A consumer prefers higher quality, and if she consumes both firms' products, the specialist's high quality offering replaces the equivalent part of the generalist's offering over the relevant range $(1-\alpha)$. For example, a consumer may use a range of web applications from a generalist provider, but use a specialist's higher-quality calendar service rather than the generalist's offering. The value v that a consumer derives from choosing both content providers, only the generalist content, and only the specialist content respectively, are:

$$v(G,S) = (1-\alpha)q_S + \alpha q_G \tag{3.1}$$

$$v(G) = q_G \tag{3.2}$$

$$v(S) = (1-\alpha)q_S \tag{3.3}$$

That is, a consumer can obtain one of q_G across the range 1, q_S across the range $1-\alpha$, or q_S across $1-\alpha$ and q_G across the remaining α.[4] These preferences imply that any benefit to a consumer from satisfying all of her interests in one place (by consuming only the generalist's product) is outweighed by the quality premium the specialist offers in its niche. The value to the consumer of the outside option of not consuming either provider is assumed to be zero.

Firms seek revenue. If a firm is not chosen by any consumers, it earns zero revenue. If a firm is chosen, it earns revenue according to (i) the number of consumers who choose it, (ii) the share of consumer interests it satisfies (as described above) and (iii) 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 'targeted' product (T) is enhanced by consumer data and gives revenue of r_T per consumer, and an 'untargeted' product (U) is not enhanced by consumer data and gives revenue of $r_U < r_T$. In the advertising case, r_T represents the revenue per impression of targeted

[3]This is a scope-based interpretation of the two firms' stature, but it is consistent with any setting in which one firm is of higher quality than the other but is initially smaller (perhaps due to some constraint on the speed of growth), even if later it could expand and grow to a larger size.

[4]An alternative environment would be that a consumer adds the specialist's product to the range they consume from the generalist, rather than having it replace the generalist's equivalent. This would mean that a consumer would consider the specialist's standalone value rather than the value premium it offers over the generalist's equivalent, but the spirit of our results would remain under this alternative specification.

advertising T, and r_U represents the revenue per impression of untargeted advertising U.[5]

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 (either α or 1 for the general interest firm and $(1-\alpha)$ for the specialist). Specifically, the revenue earned by a firm that is chosen by a fraction n of consumers and satisfies some arbitrary share β of the representative consumer's interests is given by:

$$R_j = r_A \beta n \tag{3.4}$$

In the case without regulation, consumers suffer harm when they choose to consume any product. Specifically, if some consumer i chooses to consume the product of either or both of the providers, the consumer incurs a utility cost $h_i \geq 0$ that represents this harm. This represents the real or perceived harm a consumer faces when rules on the collection, exploitation and security of their personal data are ill-defined or opaque. In the case without regulation the consumer therefore receives a total utility

$$U_i(X) = v(X) - h_i \tag{3.5}$$

when consuming a menu of products X that is non-empty, and receives zero utility from consuming no products. Denote by $H(h)$ the fraction of consumers who have a value of h_i less than h.

In the case with regulation, the harm h_i to a consumer is removed since rules on data use are well-defined and perceptible. However, for each firm that a consumer chooses to allow to use his data, he incurs a utility cost d. His total payoff is given by

$$U_i(X) = v(X) - nd, \tag{3.6}$$

where n is the number of firms in his chosen basket X that ask for consent in order to use consumer data to streamline their operations. Assume that $d < q_G$, so that a consumer would prefer to give consent and use the generalist's full scope of services rather than consume nothing.

[5]When 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 (2013).

This cost d is a fixed one-time cost that the consumer incurs when giving explicit consent for his 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) and its amendment regarding cookies (2009/135/EC) as well as the language in proposed US privacy regulation (Corbin, 2010; FTC, 2010). One notable feature about new privacy regulation in Europe and potential regulation in the US is the demand that consent be explicit, and require explicit action on the part of consumers, rather than being passive.

Regulation of this type therefore introduces a friction to the consumption experience. Our assumption that such consent requirements represent a fixed cost to consumers reflect the 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, and Lambrecht et al. (2011) show that interruptions caused by a need to ensure data privacy and security in the adoption process can deter consumers.

There are several theories about why requiring explicit opt-in consent such as demonstrated in Figures 1 to 3 deters consumers. One explanation is that opt-in consent forces a consumer to spend time becoming familiar 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. Alternatively, even if consumers do not invest time familiarizing themselves with privacy policies, the interruption and delay in the experience engendered by having to opt-in could impose costs that deter consumers from proceeding. This is particularly likely when opting in involves not simply a single click of agreement but requires that the consumer complete a registration form. This interpretation of the parameter d is consistent with a consumer who is indifferent over whether her data is used to enhance the product, outside of the direct cost of explicit consent, but exerts effort to click the "allow cookies" button. 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).

Another explanation for the deterrence effect found in empirical studies is that opt-in consent requirements make the consumer aware of some particular use of their data to which they object. Under this interpretation, d is consistent with the consumer suffering direct harm due to some facet of information sharing. This could also be a residual of the general harm h the consumer suffered due to the opacity of information policies in the case with no regulation. For example, the privacy theory literature has documented that a subset of people can be harmed by information sharing, perhaps because they are 'bad targets' (Taylor, 2004; Hermalin and Katz, 2006). In other words, they will receive worse products and services if information about them is known. Then if the data policy includes targeting of this kind, consumers suffer some cost when consenting to it.

In analyzing the model we will 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 targeted product T, then a consumer incurs the cost d to adopt the firm's product, while absent regulation there would be no such cost.[6]

We can represent this setting as an extensive form-game. Because each consumer prefers (in the case with regulation) to consume the generalist's product alone rather than nothing when the generalist offers a targeted product, the outcome when the specialist does not enter is fixed for simplicity at the generalist using customer data to optimize its operations and consumers choosing to consume its product in the regulation case and choose optimally in the no-regulation case[7]. 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. Each consumer i chooses a basket X.

[6]In Section 5 we discuss how the structure of consent-gathering may influence the nature of the cost d and its implications.

[7]In 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 4, in which payoffs at terminal nodes are for the specialist, the generalist, and consumer i respectively. The objects of interest will be

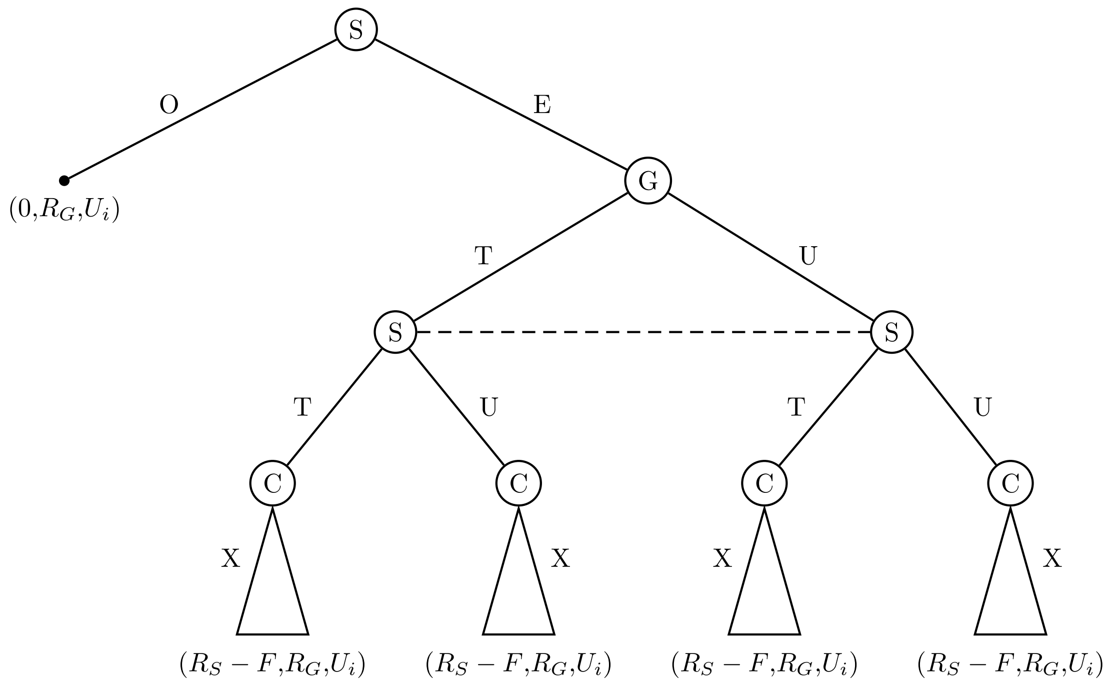

Figure 4: Extensive form

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 $\sigma_S \in \{O, E\} \times \{T, U\}$, an entry decision and a product choice decision. A strategy for the generalist is the product choice decision $\sigma_G = \{T, U\}$. A strategy for the consumer is a menu of products $\sigma_C = X \subseteq \{G, S\}$. Total payoff for consumer i is $U_i(X)$, for the generalist R_G, and for the specialist $R_S - F$ if it enters and zero otherwise.

3.1 No regulation benchmark

First consider the case when firms are not required to obtain the consent of consumers to use prior customer data to enhance their operations. Consumers will therefore not incur consent costs, but instead suffer general harm due to uncertainty over if and how their data will be used. Consider the product offering choice subgame that begins after the specialist firm chooses to enter. If both firms exist, a consumer's optimal choice is independent of the firm's product offering choice, since no consent cost is payable: a consumer who is sufficiently concerned about possible data exploitation (that is, with sufficiently high h) will choose to consume neither product, and a consumer who is less concerned will choose to consume both products.

Specifically, if both firms are active, consumer i will choose to consume both products if

$$h_i < \hat{h} \equiv (1-\alpha)q_S + \alpha q_G, \tag{3.7}$$

or else will choose to consume neither product. Choosing a targeted product is a weakly dominant strategy for both firms in the subgame following entry by the specialist. We can then characterize a qualitatively unique outcome in equilibria of the full game.

Theorem 1. *In equilibrium in the case without regulation:*

1. *The generalist vendor offers a targeted product.*
2. *The specialist vendor enters with a targeted product if* $F < H(\hat{h})(1-\alpha)r_T$*, or else does not enter.*
3. *If the specialist entered, consumer* i *consumes both products if* $h_i < (1-\alpha)q_S + \alpha q_G$*, or else consumes neither firm's product.*

The proof of this result (and others to follow) appears 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 targeted product, and consumers use both products. The revenue of the specialist after entry is increasing in its scope and in the proportion of consumers who are willing to participate in the industry.

3.2 Regulation case

Now consider the case in which privacy regulation, as described above, mandates that firms require the consumer to sign up before they can use data in their operations.

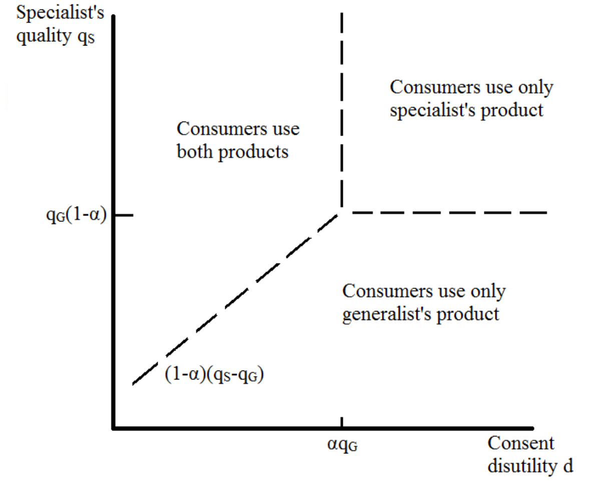

(a) Both use targeted product

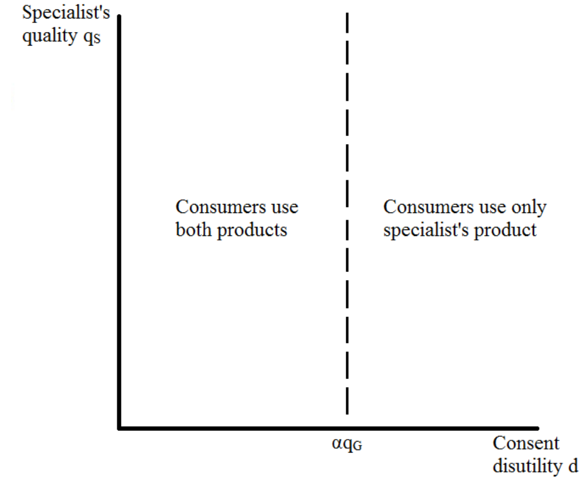

(b) Generalist uses targeted, specialist uses untargeted

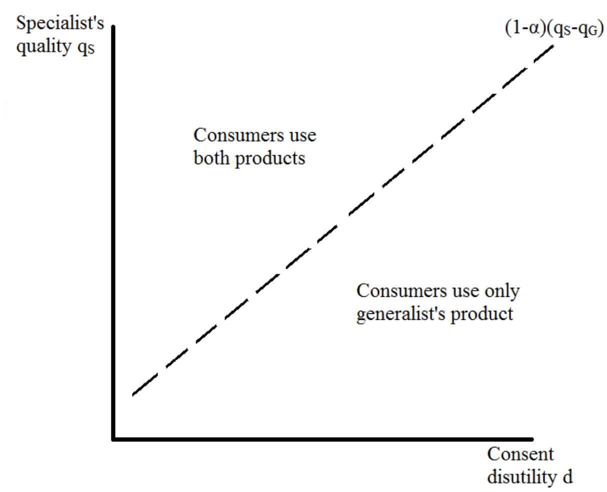

(c) Specialist uses targeted, generalist uses untargeted

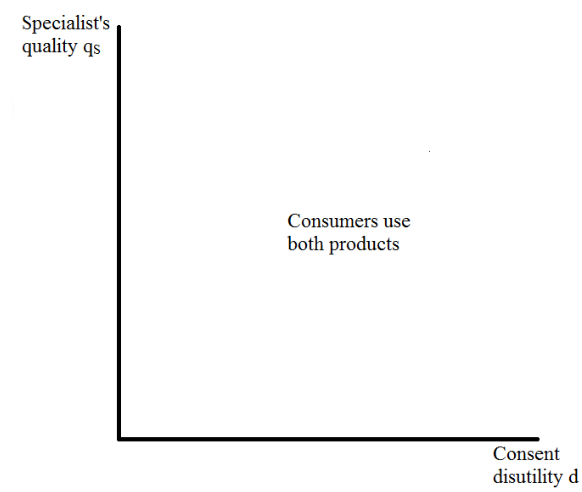

(d) Both use untargeted product

Figure 5: Consumer's optimal basket of sites in the online-advertising subgame, given firm strategies

Figure 5 summarizes a consumer's optimal choice of products 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 5a illustrates the case in which both firms choose to offer a targeted product. In that case, when the cost to a consumer of giving their consent (d) is sufficiently small, and the quality of the specialist (q_S) is sufficiently high, the consumer chooses to consume both products. The specialist then earns revenue $(1-\alpha)r_T$ and the generalist earns αr_T.

These consumer responses thus define payoffs in the product choice subgame to each firm given a pair of product choices.[8] The payoffs to the firms in this stage therefore depend on the parameters q_S, q_G, α and d, because these define the consumer's optimal basket given some pair of choices by the firms.

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

Lemma 1. *In the unique equilibrium in the product choice subgame:*

a. *When $q_G > (1-\alpha)q_S$:*

- **i.** *If $d > \alpha q_G$, both firms choose an untargeted product.*
- **ii.** *If $d \in ((1-\alpha)(q_S - q_G), \alpha q_G)$, the generalist chooses a targeted product and the specialist an untargeted product.*
- **iii.** *If $d < (1-\alpha)(q_S - q_G)$, both firms choose a targeted product.*

b. *When $q_G < (1-\alpha)q_S$:*

- **i.** *If $d > (1-\alpha)(q_S - q_G)$, both firms choose an untargeted product.*
- **ii.** *If $d \in (\alpha q_G, (1-\alpha)(q_S - q_G))$, the specialist chooses a targeted product and the generalist an untargeted product.*
- **iii.** *If $d < \alpha q_G$, both firms choose a targeted product.*

In all equilibria all consumers choose to consume both products.

Equilibrium product offering by each firm depends on the relative strength of the two firms and on the response of the consumers. 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 targeted product, and if distaste is 'large', the unique equilibrium is for both to choose an untargeted product. For intermediate cases, the unique equilibrium is asymmetric, but in which direction depends on the quality premium and scope of the specialist. If

[8] In our model the quality of a firm's product does not depend on whether it chooses a targeted or untargeted product. If this was not the case (perhaps because targeted products are better because they are more relevant, or worse because they are more intrusive), then consumers' optimal choices in the asymmetric cases Figure 5c and 5b would depend on the difference in quality between targeted and untargeted products. However, since firms must offer one or the other type of product, further qualitative results would be unchanged.

the quality premium and the scope of the specialist are relatively small, then the unique equilibrium has the generalist offering a targeted product and the specialist an untargeted product.

Figure 6 illustrates these equilibria in the product choice subgame for various values of the consumers' distaste parameter d and the quality of the specialist q_S.

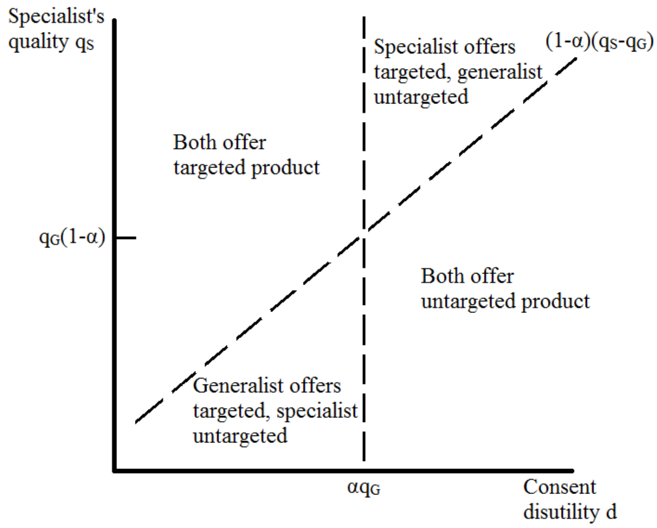

Figure 6: Lemma 1

3.2.1 Implications for extent of privacy intrusion

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

3.3 The equilibrium effect of regulation on 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 targeted product. These equilibria are characterized by the cost of consent to a consumer (d) being 'small enough'. Further, the threshold below which d is 'small enough' is more permissive the larger $(1-\alpha)$ and better $(q_S - q_G)$ 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. *In equilibrium in the case with regulation:*

1. *The generalist vendor offers a targeted product if $d \leq \alpha q_G$ and offers an untargeted product if $d > \alpha q_G$.*

2. *The specialist vendor*

   a. *enters with a targeted product if $F < (1-\alpha) r_T$ and $d < (1-\alpha)(q_S - q_G)$,*

   b. *enters with an untargeted product if $F < (1-\alpha) r_U$, $d > (1-\alpha)(q_S - q_G)$, and*

   c. *does not enter if $F > (1-\alpha) r_T$ or if $F \in [(1-\alpha) r_U, (1-\alpha) r_T]$ and $d > (1-\alpha)(q_S - q_G)$.*

3. *All consumers choose to consume all available products.*

In equilibrium, this strategy profile results in each consumer choosing the menu $\{G, S\}$ whenever both firms operate in the game with regulation. Intuitively, absent regulation, entrants offer a targeted 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, a consumer's costs of giving consent for their data to be used—the parameter d—defines the equilibrium in the product choice subgame played between the two firms. Smaller entrants and entrants that offer a smaller quality premium in their niche are more likely to offer an untargeted product in equilibrium after entry. Since an untargeted 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. *The specialist vendor enters with no regulation and is deterred from entry with regulation if:*

$$(1-\alpha)(q_S - q_G) < d, \tag{3.8}$$

$$(1-\alpha) \in \left[\frac{F}{H(\hat{h})r_T}, \frac{F}{r_U}\right]. \tag{3.9}$$

This follows directly from Theorems 1 and 2. The condition 3.8 is such that the entrant does not offer a targeted product in the equilibrium in the post-entry subgame under regulation, and 3.9 is the condition under which this precludes profitable entry.

The first implication of this result is that if consumers are sufficiently concerned about their data being exploited in the absence of regulation, then the regulation is not an entry deterrent. That is, for regulation to deter entry it is necessary that $H(\hat{h})r_T > r_U$. If this is not the case—if the proportion of consumers who were unwilling to participate in the industry due to concern over the use of their data (high h_i) is high enough—then the entrant never earns less revenue after entry in the game with regulation than in the game without.

If, however, only a few consumers were choosing not to consume anything in the pre-regulation case (that is, if $H(\hat{h})$ is sufficiently large), then the effect of the regulation is to deter entry for some potential entrants. The smaller the quality premium, the smaller the potential entrant, the larger the gap between revenues earned by targeted and untargeted products, and the more the consumer incurs costs giving consent, the larger the set of types 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 targeted product, in order to either freely offer the targeted product without being shunned by the consumer or be profitable despite having to use less lucrative untargeted products. The regulation therefore helps entrench relatively strong generalists.

It can also be the case that a specialist that did not enter before regulation can profitably enter under regulation:

Corollary 2. *The specialist vendor does not enter with no regulation and enters with regulation if*

$$(1-\alpha) < \frac{F}{H(\hat{h})r_T} \tag{3.10}$$

and either:

$$(1-\alpha) > \frac{F}{r_T} \text{ and } (1-\alpha)(q_S - q_G) > d \text{ or} \tag{3.11}$$

$$(1-\alpha) > \frac{F}{r_U} \text{ and } (1-\alpha)(q_S - q_G) < d \tag{3.12}$$

Note first that as h_i tends to zero for all i, the range of parameters for which these conditions are satisfied gets smaller; in the limit, in the absence of any harm in the pre-regulation case, the effect of the regulation is unambiguously anti-competitive.

In the general case, condition 3.10 precludes the specialist from profitable entry before regulation. This is more likely the greater the proportion of consumers who choose not to consume any products due to concern over the security of their data. There are two ways in which a specialist that is precluded from entry in this way will enter under regulation. The first is covered by conditions 3.11. In this case, d is small enough and the specialist broad and good enough that consumers are willing to adopt the specialist's targeted product after regulation. This increases the specialist's revenue after entry. The second way is covered by conditions 3.12. In this case, d is big enough that consumers will not consume a targeted product from the specialist in the game with regulation. However, since so many consumers did not consume the specialist's targeted product under no regulation, the specialist can earn more revenue with an untargeted product under regulation than with a targeted product under no regulation.

Figure 7 summarizes Corollaries 1 and 2 by illustrating equilibrium strategy of the entrant as a function of its quality and the fixed cost it faces. Figure 7a illustrates the case with no regulation and Figure 7b the case with regulation. Two effects are shown: first, the area in which the highest quality entrants can profitably enter increases, due to the regulation having increased consumer participation; second, the area in which entrants of lower quality (though still of greater quality than the generalist incumbent, by assumption) can profitably enter shrinks, due to consumers being unwilling to pay the cost d to consent to the specialist using their data and offering a targeted product. There is a further effect that some specialists enter under both regimes but do worse under regulation since they are forced by consumer preference to adopt untargeted products. The existence of this latter region suggests the potential for competitive distortion even in the case in which profitable entry is not precluded.

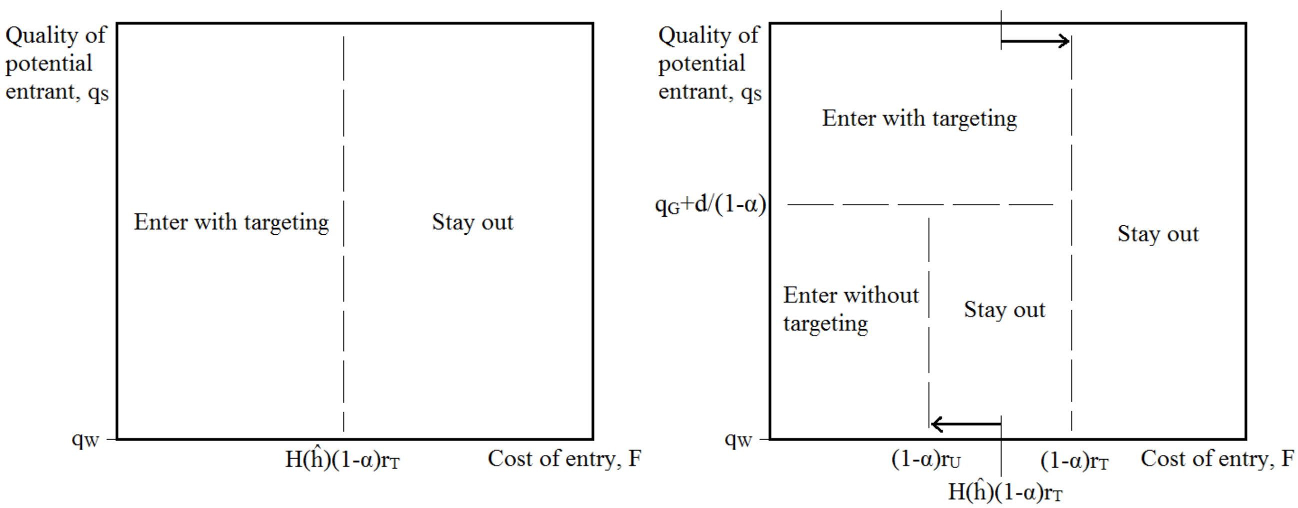

(a) No regulation

(b) With regulation

Figure 7: Equilibrium strategy of the entrant for various q_S, 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 a consumer's distaste for giving consent is sufficiently high then equilibria feature the incumbent choosing untargeted 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 Corollary

1, regulation *benefits* the incumbent, which earns revenue $R_G = r_T$ when the entrant chooses to stay out, and at best $R_G = \alpha r_T$ when the entrant enters. This is more likely for relatively strong incumbents: the stronger the incumbent, the better the marginal entrant must be.

3.4 The effect of regulation on welfare

In this model the introduction of privacy regulation has two effects on welfare. Its direct effect is to remove for consumers the harm h they suffer when participating in an industry that has incentives to exploit consumer data without well-defined rules on data use. The second effect is the competitive effect. As we saw in the previous section, depending on the characteristics of the specialist and on consumer attitudes to potential exploitation of their data, the consent-based regulation we consider can be pro- or anti-competitive.

Focusing first on consumers, whether consumers are better off overall depends on the relationship between the harm they suffer in the absence of regulation and the inconvenience, cost of consent, and change in available products they endure in the presence of regulation. Given the distribution of pre-regulation harm h across the population of consumers, some consumers will do better under regulation since they are reassured about the use of their data and now participate where they would not have before, but some consumers will do worse, since they cared little about the use of their data before and are now inconvenienced. This is true even in the absence of competitive effects, which determine how the menu of available products is affected by the change in regime.

Overall, the balance between the competitive and direct effects can lead to various outcomes. For example, the regulation can be both anti-competitive and welfare-reducing: if consumers were relatively unconcerned with data usage and specialists offer a relatively modest quality premium, regulation does not increase participation by much, but introduces a friction that dissuades the specialist from entry. In the opposite case the regulation can be both pro-competitive and welfare-enhancing: if consumers were very concerned before but experience low consent frictions, they are reassured by regulation and both participate more (to the benefit of the specialist) and are happier when they participate. It is also possible that the regulation can be anti-competitive but welfare-enhancing: if consumers suffered large harm before and consent frictions are relatively high, then before regulation consumers are uneasy but participate anyway, while after regulation they are reassured overall but less willing to consent to adopt the specialist's product.

Which welfare effect prevails therefore depends on relative strength of these various forces. Our model suggests that even if the harm consumers suffer from ill-defined rules on data use—the parameter h in the model—is well understood, a regulation that addresses it can have competitive effects which warrant consideration. In the limit as this harm tends to zero, the effect of the regulation is purely anti-competitive and so strictly welfare-reducing.

4 Extensions

4.1 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 targeted product to compensate consumers for the cost of consent, or raise the price of an untargeted product to make entry profitable.

For simplicity, we focus only on the specialist's problem: fix the generalist's choice to be the targeted product T and assume that the consumer will certainly include this product 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 to use a targeted or untargeted product (T or U).

3. The consumer C chooses between the baskets $\{G\}$ and $\{G, S\}$.

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.

Decompose the payoff r 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, $c_T < c_U$,[9] and p_T and p_U represent a price chosen by the

[9] The assumption $c_T < c_U$ is based on our understanding of industry use of data. Costs are lower for companies that use data because the data enables not just better targeting (through horizontal differentiation) but also more

firm but affects consumer utility. Thus p is the component of net revenue that firms control and that directly affects consumer utility. In contrast, c is the component of net revenue that firms do not control and that does not directly affect consumer utility. This means that the revenue earned by the specialist when it uses product $A \in \{T, U\}$ and is adopted by all consumers is

$$R_S = (p_A - c_A)(1 - \alpha) \tag{4.1}$$

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. *In the model with flexible prices, the specialist vendor enters with no regulation and is deterred from entry with regulation if*

$$c_T + \frac{F}{1-\alpha} < (1-\alpha)(q_S - q_G), \tag{4.2}$$

$$c_U + \frac{F}{1-\alpha} > (1-\alpha)(q_S - q_G), \tag{4.3}$$

$$c_T + \frac{F}{1-\alpha} + d > (1-\alpha)(q_S - q_G). \tag{4.4}$$

The proof appears in the Appendix. Equation 4.2 is the condition for which the specialist can profitably enter absent regulation, and Equations 4.3 and 4.4 are the conditions under which no price that the consumer is 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 r_A into $p_A - c_A$. Then Corollary 1 shows that under fixed prices the specialist

efficient production. For example, in the case of hospitals, data lowers costs because it reduces complications from inappropriate treatment. In the case of the advertising-supported internet, data (in effect) lowers costs because it increases advertising revenue from a given consumer in a way that does not directly affect consumer utility as p does in this subsection.

can no longer profitably enter under privacy regulation if

$$d > (1 - \alpha)(q_S - q_G) \tag{4.5}$$

$$c_U + \frac{F}{1 - \alpha} > p_U. \tag{4.6}$$

Comparing the flexible-prices (i.e. prices that affect consumer utility) 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.

4.2 Many potential entrants

The model above considered a single potential entrant. In this section we extend the model to the case in which there is a pool of potential entrants who each serve different niches. Let consumers and the generalist incumbent be as in Section 3, but now assume a set $M = \{1, ..., m\}$ of potential entrants. Each entrant is a specialist that serves a share γ_j of the consumers' range of interests and has a quality q_j that is strictly greater than the quality of the generalist. Assume no overlap between the ranges served by any two potential entrants, and that the total range covered by the m potential entrants is strictly less than the consumers' full range of interests (so that some of the consumers' interests can be served only by the incumbent generalist).

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

Theorem 4. *Specialist vendor m enters with no regulation and is deterred from entry with regu-*

lation if:

$$\gamma_m(q_m - q_G) < d, \tag{4.7}$$

$$\gamma_m \in \left[\frac{F}{H(h_M)r_T}, \frac{F}{r_U}\right], \tag{4.8}$$

where $h_M = \sum_{m \in M} \gamma_m q_m + (1 - \sum_{m \in M} \gamma_m) q_G$.

As in the single-entrant case, the marginally profitable entrant must be broader and of higher quality in the case with regulation. However, in this case with a range of qualities and scopes across the set of potential entrants, deterred entry can manifest not as unobserved entry (as in the single-entrant case) but as a smaller set of specialist entrants after regulation.

In this setting it is also the case that the ability of one specialist to enter in the pre-regulation case can depend on the quality and scope of potential entrants in other niches. This is because in the absence of regulation the quality of consumer i's whole basket of chosen products must be sufficiently high to overcome h_i, the harm she suffers by consuming any product when data use is ill-defined. Another way to interpret deterred entry in this case is then that an entrant that was 'carried' to profitability by a vibrant overall industry in the pre-regulation case must not stand alone against the cost to the consumer of giving consent.

4.3 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 in the case in which the regulation has anti-competitive effects. To do this, we add an investment stage to the beginning of the game 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 a game as before. 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 $q_G \sim [q_L, \infty)$, or not to invest $(\neg i)$ and maintain quality q_L.

3. If the specialist chose to invest, the firms and consumer play a game as in Section 3 (either with or without regulation) under the parameters q_S and either q_G or q_L depending on whether the incumbent chose to invest.

This defines an extensive game as represented by the game tree in Figure 8.

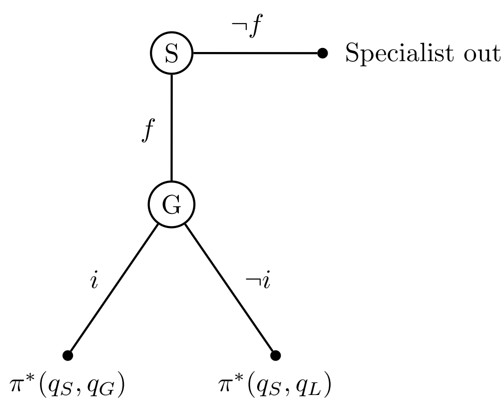

Figure 8: Extensive form of investment choice game

A strategy for the specialist is $\sigma_S \in \{f, \neg f\}$, and for the generalist $\sigma_G \in \{i, \neg i\}$. Payoffs after f are defined by those in equilibrium of the game as in Section 3 with quality parameters as resulting from the investment decisions. In Figure 8 these equilibrium payoffs after f are denoted by π^* as a function of the qualities of the two firms' products.

Fix $(1-\alpha) \in [\frac{F}{H(\hat{h})r_T}, \frac{F}{r_U}]$ so that the entrant can make positive profit only when offering a targeted 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 targeted product and thus to profitably enter.

By assumption $\Pr(q_S > q_L) = 1$. Assume also that $\Pr(q_S > \frac{d}{1-\alpha} + q_L) = 1$. Together these imply that in both the regulation and no-regulation cases the specialist can profitably enter with probability one whenever it chooses to invest and the generalist does not choose to invest. Define the following:

$$\tau \equiv \Pr(q_S > q_G), \tag{4.9}$$

$$\psi \equiv \Pr(q_S > \frac{d}{1-\alpha} + q_G), \tag{4.10}$$

where $\tau, \psi \in [0, 1]$ and $\psi < \tau$ by necessity. With probability τ the specialist obtains a quality high enough to profitably enter in the case without regulation if the generalist invests. With probability ψ the specialist obtains a quality high enough to enter in the case with regulation if the incumbent invests.

The object of interest is investment along the equilibrium path. For the specialist, this is direct; for the generalist to invest along the equilibrium path requires both that the specialist invests and that the generalist chooses also to invest. The following result characterizes the change in investment behavior under the no-regulation and regulation regimes:

Theorem 5. **i.** *If the specialist does not invest without regulation, it does not invest under regulation.*

ii. *The specialist invests without regulation and does not invest under regulation if $I \in [(1-\tau)K, (1-\psi)K]$ and $F > \psi K$.*

iii. *The generalist does not invest without regulation and invests under regulation if $I \in [(1-\tau)K, (1-\psi)K]$ and $F < \psi K$.*

iv. *The generalist invests without regulation and does not invest under regulation if $F \in [\psi K, \tau K]$.*

These results demonstrate that when the specialist can be profitable only with a targeted product and entry is potentially deterred by regulation, the entry-deterring effect is not mitigated by the opportunity to invest in quality. In this situation the specialist never invests more under regulation than under no regulation, and in some cases invests less. The generalist in some cases also invests less, although it can be that in some case the generalist invests more after regulation, since the regulation raises the quality bar for the specialist's profitable entry and so creates a stronger incentive for the generalist to engage in defensive investment. In this framework there is thus no ancillary benefit from the entry-deterring effect of the privacy regulation to investment incentives for prospective specialist entrants; the regulation always increases the difficulty the entrant faces in challenging a relatively strong incumbent.

5 Policy Implications

The above results show 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. While there has not been much progress on this so far, perhaps the closest is the W3C standards setting process for 'Do not track', which aims to set a protocol for how users can prevent firms from tracking their movements online. However, there is little optimism about the ability of this process to reach consensus.[10] To the extent that W3C represents a decrease in d, facilitating and prioritizing this standard-setting process may be a key government action to reduce the anticompetitive implications of regulation.

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

[10] Bott (2012) says for example, 'The debate over the Do Not Track standard has officially moved beyond Alice in Wonderland. These days, I'm not sure whether it's 1984 or Brazil.'

further incremental barrier to entry in the regulatory game relative to the game without regulation. A global opt-in would then have the potential to mitigate the anti-competitive effect. For example, this may be a spur for governments to embrace consent mechanisms that are administered via a single set of browser settings. We note that currently this goes against policymaking in the EU - for example, the Belgian Privacy Commission indicated that consent may not be obtained through current browser settings when preparing their draft bill.[11]

The reason the EU has spurned such 'single opt-in or opt-out" measures is that laws such as the EU Privacy and Electronic Communications Directive, are impossible to implement via a single 'opt-in' because they require a sliding scale of consent that depends on 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. The model in this paper highlights that a consequence of such context-dependent consent requirements is the potential for anti-competitive effects.

6 Conclusion

In this paper we investigate the relationship between offering privacy protection and 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.

As discussed above, our model is motivated by privacy regulation that emphasizes tracking on

[11] Opinion 10/2012 of 21 March 2012.

the advertising-supported internet. For example, it might model competition between a generalist news provider like Google News and a specialist news provider like CelebrityBuzz.com. These websites use data to better target advertising and increase revenues per visitor. Importantly, the ideas and model can extend to other contexts where data use is widespread. The generalist product could be Visa credit cards while the specialist could be a department store credit card. Both use data to perform credit checks in order to reduce defaults. Or, the generalist could be an HMO and the specialist could be a distinctive clinic. Both use data to access medical histories and reduce record-keeping costs.

In a variety of industries, 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.

Armstrong, M., J. Vickers, and J. Zhou (2009). Consumer protection and the incentive to become informed. *Journal of the European Economic Association 7*(2-3), 399–410.

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*.

Bott, E. (2012, October 9). The do not track standard has crossed into crazy territory. *ZdNet*.

Bouckaert, J. and H. Degryse (2006). Opt in versus opt out: A free-entry analysis of privacy policies. Technical report.

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*.

Fudenberg, 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. (2013). Targeted advertising and advertising avoidance. *RAND Journal of Economics*.

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.

Posner, R. A. (1981). The economics of privacy. *American Economic Review 71*(2).

Taylor, C. R. (2004). Consumer privacy and the market for customer information. *The RAND Journal of Economics 35*(4), pp. 631–650.

Tucker, C. (2011). Social networks, personalized advertising, and privacy controls. Mimeo, MIT.

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. For any consumer, $v(G,S) > v(G), v(S)$. For consumer i, $U_i(G,S) > U_i(\emptyset)$ iff $(1-\alpha)q_S + \alpha q_G - h_i > 0$. Thus if both firms are active i will consume both firms' products if $h_i < \hat{h} \equiv (1-\alpha)q_S + \alpha q_G$ or else will consume neither product. For consumer i, $U_i(G) > U_i(\emptyset)$ iff $q_G - h_i > 0$. Thus if only the generalist is active i will consume the generalist's product if $h_i < q_G$ or else will consume no product.

By backward induction, in the product choice subgame the specialist earns $H(\hat{h})(1-\alpha)r_T$ by playing T and $H(\hat{h})(1-\alpha)r_U$ by playing U. Similarly the generalist earns $H(\hat{h})\alpha r_T$ by playing T and $H(\hat{h})\alpha r_U$ by playing U. Since $r_T > r_U$, T is strictly dominant in the product choice subgame for both firms.

Again by backward induction, after entry the specialist earns $H(\hat{h})(1-\alpha)r_T$. Since entry costs F and the payoff to staying out is zero, the specialist enters iff $H(\hat{h})(1-\alpha)r_T > F$ and stays out otherwise. □

A.2 Lemma 1

Proof. Note that $q_G > (1-\alpha)q_S$ iff $\alpha q_G > (1-\alpha)(q_S - q_G)$.

a.iii. and b.iii. Since $d < \min\{\alpha q_G, (1-\alpha)(q_S - q_G)\}$, $X^* = \{W, S\}$ for any σ_G, σ_S and so $\sigma_G = T$, $\sigma_S = T$ are dominant strategies for W and S in the product choice subgame with regulation. At $\sigma_G = T$, $\sigma_S = T$, $X^* = \{G, S\}$.

a.i and b.i. Since $d > \max\{\alpha q_G, (1-\alpha)(q_S - q_G)\}$, $j \notin X^*$ if $\sigma_j = T$ and so $\sigma_G = U$, $\sigma_S = U$ are dominant strategies for G and S in the product choice subgame with regulation. At $\sigma_G = U$, $\sigma_S = U$, $X^* = \{G, S\}$.

a.ii. Since $d < \alpha q_G$, $X^* = \{G, S\}$ when $\sigma_G = T$, $\sigma_S = U$. Since $d > (1-\alpha)(q_S - q_G)$ and $q_G > (1-\alpha)q_S$, $X^* = \{G\}$ when $\sigma_G = T$, $\sigma_S = T$. Thus $\sigma_G = T$ is dominant for G in the product choice subgame with regulation, and the specialist's best response is $BR_S(\sigma_G = T) = U$. At $\sigma_G = T$, $\sigma_S = U$, $X^* = \{G, S\}$.

b.ii. Since $d < (1-\alpha)(q_S - q_G)$, $X^* = \{G, S\}$ when $\sigma_G = U$, $\sigma_S = T$. Since $d > \alpha q_G$ and $q_G < (1-\alpha)q_S$, $X^* = \{S\}$ when $\sigma_G = T$, $\sigma_S = T$. Thus $\sigma_S = T$ is dominant for S in the product choice subgame with regulation, and $BR_W(\sigma_S = T) = U$. At $\sigma_G = U$, $\sigma_S = T$, $X^* = \{G, S\}$.

□

A.3 Theorem 2

Proof. If $d > (1-\alpha)(q_S - q_G)$, then when $\sigma_S = T$ in the product choice subgame, $X^* = \{G\}$ and so $R_S = 0$, but when $\sigma_S = U$, $X^* = \{G, S\}$ and so $R_S = (1-\alpha)r_U > 0$. Thus if $d > (1-\alpha)(q_S - q_G)$ then $\sigma_S^* = U$ in the product choice subgame in any SPE.

If $d < (1-\alpha)(q_S - q_G)$, then when $\sigma_S = T$ in the product choice subgame, $X^* = \{G, S\}$ and so $R_S = (1-\alpha)r_T$, but when $\sigma_S = U$, $X^* = \{G, S\}$ and so $R_S = (1-\alpha)r_U < (1-\alpha)r_T$. Thus if $d < (1-\alpha)(q_S - q_G)$ then $\sigma_S^* = U$ in the product choice subgame in any SPE.

If $\sigma_S^* = T$ in the product choice subgame and $F > (1-\alpha)r_T$, then if $\sigma_S = E, T$ then the specialist's payoff is $R_S - F < 0$, but if $\sigma_S = O, T$ then the specialist's payoff is 0. Thus if $d < (1-\alpha)(q_S - q_G)$ and $F > (1-\alpha)r_T$ then $\sigma_S^* = O, T$. If $\sigma_S^* = T$ in the product choice subgame and $F < (1-\alpha)r_T$, then if $\sigma_S = E, T$ then the specialist's payoff is $R_S - F > 0$, but if $\sigma_S = O, T$ then the specialist's payoff is 0. Thus if $d < (1-\alpha)(q_S - q_G)$ and $F < (1-\alpha)r_T$ then $\sigma_S^* = E, T$.

If $\sigma_S^* = U$ in the product choice subgame and $F > (1-\alpha)r_U$, then if $\sigma_S = E, U$ then the specialist's payoff is $R_S - F < 0$, but if $\sigma_S = O, U$ then the specialist's payoff is 0. Thus if $d > (1-\alpha)(q_S - q_G)$ and $F > (1-\alpha)r_T$ then $\sigma_S^* = O, U$. If $\sigma_S^* = U$ in the product choice subgame and $F < (1-\alpha)r_U$, then if $\sigma_S = E, U$ then the specialist's payoff is $R_S - F > 0$, but if $\sigma_S = O, U$ then the specialist's payoff is 0. Thus if $d > (1-\alpha)(q_S - q_G)$ and $F < (1-\alpha)r_T$ then $\sigma_S^* = E, U$. □

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

$$U_i(G,S) - U_i(G) = [(1-\alpha)q_S + \alpha q_G] - [q_G] \tag{A.1}$$
$$= (1-\alpha)(q_S - q_G). \tag{A.2}$$

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

$$F < R_j = (p_T - c_T)\beta \tag{A.3}$$
$$= ((1-\alpha)(q_S - q_G) - c_T)(1-\alpha) \tag{A.4}$$
$$\Rightarrow c_T + \frac{F}{1-\alpha} < (1-\alpha)(q_S - q_G) \tag{A.5}$$

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 G is

$$U_i(G,S) - U_i(G) = [(1-\alpha)q_S + \alpha q_G - 2d] - [q_G - d] \tag{A.6}$$
$$= (1-\alpha)(q_S - q_G) - d. \tag{A.7}$$

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

$$
\begin{aligned}
F < R_j &= (p_T - c_T)\beta && \text{(A.8)}\\
&= ((1-\alpha)(q_S - q_G) - d - c_T)(1-\alpha), && \text{(A.9)}\\
\Rightarrow c_T + \frac{F}{1-\alpha} + d &< (1-\alpha)(q_S - q_G). && \text{(A.10)}
\end{aligned}
$$

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

$$
\begin{aligned}
U_i(G,S) - U_i(G) &= [(1-\alpha)q_S + \alpha q_G - d] - [q_G - d] && \text{(A.11)}\\
&= (1-\alpha)(q_S - q_G). && \text{(A.12)}
\end{aligned}
$$

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

$$
\begin{aligned}
F < R_j &= (p_U - c_U)\beta && \text{(A.13)}\\
&= ((1-\alpha)(q_S - q_G) - c_U)(1-\alpha), && \text{(A.14)}\\
\Rightarrow c_U + \frac{F}{1-\alpha} &< (1-\alpha)(q_S - q_G). && \text{(A.15)}
\end{aligned}
$$

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

A.5 Theorem 5

Proof. Denote the game without regulation J and the game with regulation J'. The following result defines equilibria in the games J and J'.

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

$$
\begin{aligned}
&(f, i) \text{ if} && I < (1-\tau)K, F < \tau K, && \text{(A.16)}\\
&(f, \neg i) \text{ if} && I > (1-\tau)K, && \text{(A.17)}\\
&(\neg f, i) \text{ if} && I < (1-\tau)K, F > \tau K. && \text{(A.18)}
\end{aligned}
$$

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

$$
\begin{aligned}
&(f, i) \text{ if} && I < (1-\psi)K, F < \psi K, && \text{(A.19)}\\
&(f, \neg i) \text{ if} && I > (1-\psi)K, && \text{(A.20)}\\
&(\neg f, i) \text{ if} && I < (1-\psi)K, F > \psi K. && \text{(A.21)}
\end{aligned}
$$

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

For the incumbent in J', similarly, but now if G plays $\neg i$, $q_G = q_L$ and $q_L + \frac{d}{1-\alpha} < q_S$ with probability 1; $R_G = \alpha r_T$ with probability 1. If G plays $\neg i$, $q_G + \frac{d}{1-\alpha} < q_S$ with probability ψ; $R_G = \alpha r_T$ with probability ψ and $R_G = r_T$ with probability $(1-\psi)$. The expected payoff to playing i, $\psi(\alpha r_T) + (1-\psi)r_T$, exceeds the payoff to $\neg i$, (αr_T), if $I < (1-\psi)K$. Thus if $I < (1-\psi)K$, $\sigma_G^*(J') = i$, else $\sigma_G^*(J') = \neg i$.

Next, consider the entrant in J. By backward induction, since $F < K$ by assumption, $\sigma_S^*(J) = f$ when $I > (1-\tau)K$, since $\sigma_G^*(J) = \neg i$, $q_S > q_G$ for sure and $R_S = K$. If $I < (1-\tau)K$ so that $\sigma_G^*(J) = i$, $q_S > q_G$ 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 $\sigma_G^*(J) = i$, $\sigma_S^*(J) = f$ if $I < (1-\tau)K$, $F < \tau K$, and $\sigma_G^*(J) = i$, $\sigma_S^*(J) = \neg f$ if $I < (1-\tau)K$, $F > \tau K$.

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

If $I > (1-\psi)K$, $\sigma_G^*(J') = \neg i$. If the entrant plays f, $q_G = q_L$ and $q_G + \frac{d}{1-\alpha} < q_S$ with probability 1, and so in expectation $R_S = K$. If the entrant plays $\neg f$, its payoff is zero. Thus if $I > (1-\psi)K$ (and since $F < K$ by assumption), then $\sigma_G^*(J') = \neg i$, $\sigma_S^*(J') = f$. $\square$

Theorem 5 follows directly. $\square$