

The Effect of Privacy on Market Structure and Prices

Daniel Bird*

Eitan Berglas School of Economics, Tel Aviv University

Zvika Neeman

Eitan Berglas School of Economics, Tel Aviv University

Abstract

Protection of consumers' privacy is often motivated by the fear that, without it, consumers may be exploited via personalized pricing. We explain how privacy may affect prices in search markets through a different channel, namely, the effect privacy has on market structure. If privacy is not protected, then in addition to consumer search, firms may engage in targeted advertising. We show that privacy protection reduces consumer surplus if firms price discriminate between the search and advertising markets. Absent such discrimination, privacy protection increases consumer surplus only if the advertising market is large and noncompetitive. We relate our results to the "privacy paradox."

JEL CODES: D40, D83, L10.

KEYWORDS: privacy, price discrimination, search, targeted advertising, interest-based advertising, privacy paradox.

*Eitan Berglas School of Economics, Tel Aviv University (e-mail: dbird@tauex.tau.ac.il).

We are very grateful to three anonymous referees and the Editor for constructive comments and suggestions. We have also benefited from comments by Yair Antler, Heski Bar-Isaac, Winand Emons, Itay Fainmesser, Igor Letina, Raphael Levy, Roe Levy, Ariel Rubinstein and seminar participants at Bar-Ilan University, Bern University, HEC, the Technion and Tel Aviv University, as well as conference participants at NASMES 2022, EARIE 2022, and CESC 2023. Bird and Neeman acknowledge financial support from the Sapir Center for Economic Development.

1 Introduction

The ongoing debate about the need to protect consumers' privacy in retail markets is, to a large extent, driven by the fear that sellers might use consumer data to exploit consumers via personalized pricing.¹ In this paper, we suggest an alternative channel through which privacy may affect consumer welfare. Specifically, we explain how privacy protection, or the lack thereof, may affect the distribution of market prices through its effect on market structure.

In many retail markets consumers or firms (or both) must search for one another before they can transact. Protection of consumer privacy effectively shuts down the possibility for firms to search for consumers via targeted advertising. Hence, if consumer privacy is protected, then only a *search market*, in which (only) consumers search for the firms that sell the goods they are interested in, may exist. If, on the other hand, consumer privacy is not protected, then firms may acquire relevant information about consumers, which facilitates targeted advertising. This implies that alongside search markets, there may also exist *advertising markets*, in which firms make direct offers to consumers who are interested in the goods they sell. Equilibrium prices in the search market obviously depend on whether firms may also reach consumers through targeted advertising, and so privacy, or the lack thereof, affects both market structure and pricing.

We consider a simple but prevalent form of targeted advertising, namely, "interest-based advertising," that is, advertising that is used by a firm that knows nothing about consumers beyond the fact that they are possibly interested in the general category of goods that it produces.^{2,3} Interest-based advertising does not allow firms to screen between different types of consumers. Thus, in essence, in this paper we analyze how limitations of firms' ability to reach consumers can impact market prices. Accordingly, our

¹See, e.g., the survey by Acquisti, Taylor and Wagman (2016).

²The Interactive Advertising Bureau estimates interest-based advertising revenues in the United States alone to be equal to \$107.5 billion in 2018. These revenues are expected to continue to grow rapidly.

³ The website youradchoices.com (see <https://youradchoices.com/choices-faq#jr02>) that describes interest-based advertising explains that "a typical set of information associated with a user's web browser might include:

Gender: Male

Age range: 25–34

Geography: Washington DC metro area

Interested in baseball

Interested in travel to Europe

Car shopper."

model can also be used to analyze the impact of other policies that affect the effectiveness of advertising, such as “no solicitation” laws, regulation of the use of online influencers, or regional advertising in television networks.

We study the effect of consumer privacy on consumer surplus, firms’ profits, and social welfare when trade can potentially occur via either a competitive search market or interest-based advertising. To study these questions, we consider a model with a large number of consumers and firms that produce a certain good. A small fraction of the consumers in the population are interested in this good and their willingness to pay for the good is fixed and commonly known by firms. Consumers have access to a search market in which they engage in a “noisy search” for the firms that produce the good that they are interested in. This means that each consumer’s search effort yields a random number of relevant price quotes, which is possibly larger than one. If consumer privacy is not protected, then firms can also reach interested consumers through interest-based advertising. As with search markets, we assume that the number of relevant interest-based ads that reach an interested consumer is random. Importantly, we allow the distribution of the number of relevant price quotes that a consumer obtains through search to be different from the distribution of the number of price quotes that are observed through ads.

We show that the effect of privacy protection crucially depends on whether firms price discriminate between the search and advertising markets (henceforth, “price discriminate”). That is, the effect of privacy protection depends on whether firms may quote one price to consumers that they reach through interest-based advertising and another price to unsolicited consumers who reach them through an active search. Such price discrimination is certainly very easy to implement with available technology, for example, by including coupon codes in firms’ ads or by using different landing pages for search engines and ads. Nevertheless, in practice, firms often refrain from such price discrimination (see the discussion at the end of Section 4).

We obtain the following main results. First, if firms price discriminate, then protection of consumer privacy hurts consumers. Intuitively, this is because if firms price discriminate, then equilibrium prices in the search market are unaffected by the opening of the advertising market, and so privacy protection denies consumers the option of buying goods through the advertising market, which hurts them. The same intuition implies that privacy protection hurts firms if and only if the advertising market is less competitive than the search market, where the competitiveness of a market is determined by the distribution of the number of price quotes in that market (see formal definition of compet-

itiveness in Section 2.1). Since consumers probably pay less attention to the ads that they are exposed to than to what they discover through their own search, the number of price quotes consumers obtain by search is likely to be larger than the number of relevant price quotes obtained from interest-based ads.⁴ Hence, in practice, the advertising market is likely to be less competitive than the search market. It follows that, in practice, relaxation of consumer privacy protection is likely to lead to a Pareto improvement.

Second, if firms do not price discriminate, then privacy protection helps consumers if and only if the advertising market is “large enough,” where the size of the advertising market is measured by the mass of consumers that are exposed to at least one relevant ad. Moreover, we find that the critical size of the advertising market is increasing in its competitiveness. Again, privacy protection hurts firms if and only if the advertising market is less competitive than the search market. As before, we find that relaxation of consumer privacy protection may yield a Pareto improvement. However, even when privacy protection increases consumer surplus, each individual consumer benefits from sharing his own private information with firms. This insight provides a possible explanation for the privacy paradox: consumers understand that they benefit collectively from privacy protection, but each consumer personally prefers to share his information with firms.⁵

Finally, we show that price discrimination increases consumer surplus if and only if it reduces firms’ aggregate profits. The intuition for this result is that price discrimination determines the distribution of surplus between consumers and firms, but the volume of trade and the number of searches that are performed until a consumer finds an acceptable price offer are identical. This result suggests that if price discrimination benefits firms, it should be constrained by regulators who are concerned about consumer surplus. On the other hand, if price discrimination hurts firms, then they would like to commit to not engage in price discrimination.

⁴ For example, Acquisti, Taylor and Wagman (2016) write that “despite the large sums of money spent on targeted advertising [...] its effectiveness is unclear” (p. 464) and Lewis and Rao (2015) write that “given the total volume of advertising [that] a typical consumer sees across all media, even an intense campaign only captures about 2% of a user’s advertising ‘attention’ ”(p. 1948).

⁵See Acquisti, Taylor and Wagman (2016) and the references therein. In recent work, Madarasz and Pycia (2021) provide an alternative explanation for the privacy paradox by considering the equilibrium of a model in which firms invest in collecting information about consumers’ preferences and consumers counterinvest to prevent firms from doing so.

Related Literature

Privacy. For an excellent recent survey of the economics of privacy, see Acquisti, Taylor and Wagman (2016). They observe that privacy protection may both benefit consumers by limiting firms' ability to extract their consumer surplus through price discrimination, and hurt consumers by increasing their search costs and denying firms information that would allow them to better cater to consumers' tastes (Bergemann and Bonatti, 2011). Thus, privacy protection may either benefit or hurt consumers, depending on which of these two effects is stronger. In some environments, sellers may benefit from committing to full privacy (Calzolari and Pavan, 2006), and in others they may benefit from committing not to engage in price discrimination (Ichihashi, 2020).⁶

In yet other cases, it may be possible to design consumer privacy to maximize consumer surplus. Romanyuk and Smolin (2019) characterize the information design policy that maximizes surplus in a matching market without price discrimination. They show that the privacy provided by simple censorship rules is often constrained efficient. Bird and Neeman (2022) show how it is possible to design consumer privacy to maximize consumer surplus in monopolistic markets.⁷ Elliott et al. (2022) study how to design firms' information about consumers in a way that weakens competition between firms and enables firms to extract the full surplus from trade without collusion or any long-term incentives. Mauring (2022) analyzes the effect of privacy when firms can obtain information about consumers' search costs rather than about their willingness to pay for a good. She finds that, in such settings, privacy protection generally benefits consumers. Finally, Fainmesser, Galeotti and Momot (2023) study the incentives of businesses to collect and protect users' data. In particular, they study how a business' revenue model affects its decisions to collect and protect the data.

Relative to the literature on consumer privacy, our contribution consists of showing how privacy regulations can affect market structure, and therefore pricing (as defined by the number of price quotes observed by consumers and their costs of obtaining them).

Search and Price Dispersion. Our model of the search market is based on Burdett and Judd's (1983) model of "noisy search." Baye et al. (2006) survey the extensive literature on consumer search and price dispersion. They divide the theoretical part of the literature

⁶Relatedly, Ali, Lewis and Vasserman (2022) and Pram (2021) show that in order to improve their terms of trade, informed buyers may wish to give up his privacy by disclosing hard information about themselves.

⁷Bergemann, Brooks and Morris (2015) and Haghpanah and Siegel (2023) characterize the distribution of surplus that can be attained via different market segmentations in a canonical buyer-seller setting.

into two strands. One strand shows that price dispersion can arise in sequential search models in which consumers obtain a deterministic number of price quotes, at a cost, in every period. The other strand employs information clearinghouse models (e.g., Baye and Morgan, 2001) in which some consumers have access to the entire set of firms that sell the good they want, for a small cost. Notice that Burdett and Judd's (1983) model of noisy search belongs to neither of these strands of the literature. Baye et al. (2006) also report on the large empirical literature that shows that price dispersion is extremely common in practice.

Targeted Advertising and Search Markets. Butters (1977) is probably the first to have considered a model with price dispersion in which a search market and an advertising market for a single good operate side by side. Robert and Stahl (1993) consider a model with homogeneous products and study the implications of price advertising when shopping trips are costly to consumers. They do not consider privacy as such, but it is of course possible to interpret lower advertising costs in their model as weaker privacy protection. They show that as advertising costs vanish, prices become competitive. By contrast, in our setting, less privacy can hurt consumers in some cases, which suggests that lowering advertising costs can hurt consumers.

More recently, Rhodes and Zhou (2022) show that in the short run, when the number of firms in the market is fixed, the impact of personalized pricing hinges on the degree of firms' market coverage. If coverage is high, then personalized pricing intensifies competition and so hurts firms but benefits consumers, whereas the opposite is true if coverage is low. However, in the long run, when the number of firms is determined endogenously, personalized pricing always benefits consumers because it induces the socially optimal level of firm entry.

Braghieri (2019) analyzes firms' pricing decisions and consumers' adoption of anonymizing technologies in markets where advertising slots are sold by a two-sided intermediary, but consumers' willingness to pay (for each good) varies. In equilibrium, firms engage in price discrimination, there is no price dispersion, and the introduction of tracking technologies makes all consumers better off. However, unlike us, Braghieri assumes that the quality of the match between consumers and firms is lower in the search market than in the advertising market, which implies that prices in the former are lower than those in the latter. Also, unlike in our model, the search market in his model collapses if firms are not allowed to price discriminate. Another paper in which targeted advertising intensifies price competition when combined with consumer search is that of De Cornière

(2016). He considers a setting with heterogeneity in consumers' willingness to pay and a standard sequential search technology. In his model, advertisers target consumers based on their search queries, after they complete their searches unsuccessfully. He shows that such targeting reduces consumers' search costs, improves matches, and intensifies price competition. Competition among search engines in his model can either increase or decrease welfare, depending on the extent of multi-homing by advertisers.

Zhu, Moorthy and Shi (2022) develop a general equilibrium model of informative advertising in order to examine the implications of privacy regulations on consumer welfare. They characterize the conditions under which the privacy preferences of different types of consumers are aligned or opposed. They show that if ad prices are *exogenous* to the privacy regime (in our model, ad prices do depend on the privacy regime) then all consumers benefit from no privacy.

Relative to the literature on targeted advertising and search, our contribution consists of the analysis of the case where firms are unable to price discriminate between consumers, which leads to potentially different welfare implications. The reasons that firms may be unable to price discriminate are discussed at the end of Section 4 below.

Price Discrimination. Intuitively, if the set of consumers who are served by firms remains fixed, then price discrimination benefits firms at the expense of consumers. In his survey of the relevant literature, Armstrong (2006) notes that an increased ability to engage in price discrimination generally boosts the profit of a monopolistic firm, unless the firm cannot commit to its pricing policy. In competitive markets, the effects of price discrimination on consumer surplus, firms' profits, and social welfare depend on the kinds of information and/or tariff instruments available to firms. If firms agree about whether specific information about consumers implies that they should set higher or lower prices, as is the case in our setting, then price discrimination typically increases firms' profits. But this is not so if firms disagree. Interestingly, Preuss (2023) shows that price discrimination hurts both firms' profits and social welfare in search markets in which firms can track whether consumers are aware of their willingness to pay, and hence are likely to have inelastic demand when consumers' search costs are small. The opposite happens when consumers' search costs are large, because in that case price discrimination discourages consumers from searching to learn their willingness to pay.

The rest of the paper proceeds as follows. In Section 2 we present the model. In Section 3 we analyze equilibria both in the case where consumer privacy is protected, i.e., where

only the search market exists, and in the case where consumer privacy is not protected, i.e., where both the search and advertising markets exist. In the latter case, we distinguish between the subcases in which firms engage and do not engage in price discrimination. In Section 4 we compare equilibria in terms of their induced consumer surplus, firms' profits, and social welfare. In Section 5 we offer concluding remarks.

2 Model

We consider a market for a good that is produced by a large number of firms at a marginal cost of zero.⁸ There exists a large population of consumers of which only a small fraction are interested in the good. These interested consumers each have a willingness to pay of $v > 0$ for (one unit of) the good. We normalize the measure of interested consumers per firm to one.⁹

Each firm quotes a search market price for the good. Consumers can search for firms by using a noisy sequential search technology (as in Burdett and Judd, 1983). In each round of search, each consumer pays a cost of $c > 0$ to obtain a random number of new price quotes. After observing the realized price quotes, each consumer decides either to buy the good at the best price obtained up to then, to quit the market, or to pay the search cost c again and obtain still more new price quotes. For simplicity, we assume that the number of price quotes obtained in each round of search is at most two. We denote the probability of obtaining $n \in \{0, 1, 2\}$ price quotes in each round of search by q_n^S . We further assume that $0 < q_1^S, q_2^S < 1$.¹⁰

If consumer privacy is not protected, then firms can contact the consumers who are interested in the good that they produce through costless interest-based advertising. Nevertheless, due to the limited attention consumers pay to firms' ads (and their limited screen

⁸Assuming instead that the (constant) marginal costs are positive would simply shift all market prices up by the size of the marginal costs.

⁹Specifically, we are thinking of the following process. Let N be an integer and $\mu \in (0, 1)$. Suppose that there are N firms and $\frac{N}{\mu}$ consumers, and that each consumer is interested in the good with probability μ . Moreover, suppose that when consumers search they are allocated randomly to firms, and that the allocation of different consumers are independent of one another. Then, as N tends to infinity, the measure of consumers who are interested in the good per firm becomes deterministic and equal to 1.

¹⁰We discuss the assumption that consumers obtain no more than two price quotes in Section 2.2 and explain why this assumption does not impact our qualitative results in Section 5. The assumption that the probabilities are interior allows us to abstract away from corner cases that complicate the exposition without adding content.

size), the fact that firms can contact consumers does not necessarily imply that consumers will notice that they have done so. Hence, as in the case of consumer search, we assume that each consumer is exposed to a random number of relevant ads, and, for symmetry with the search market, we also assume that the number of ads that a consumer observes is at most two. We denote the probability that each consumer receives $n \in \{0, 1, 2\}$ relevant ads by q_n^A , and, as in the case of consumer search, assume that $0 < q_1^A, q_2^A < 1$. Finally, we assume that firms quote an ad price that may be different from their search market price, and that consumers receive interest-based ads once and that they receive them *before* deciding whether or not to search.

We analyze three distinct market regimes. Under the first regime, consumer privacy is protected. That is, consumers may search for firms that produce the good, but there is no advertising. We denote this regime by \mathcal{P} for “privacy.” Under the second and third regimes, consumer privacy is not protected and firms have access to consumer data, which enables interest-based advertising. Under the second regime, firms may price discriminate between the search and advertising markets by offering a different price in each of these two markets. We denote this regime by \mathcal{D} for “discrimination.” The third regime is identical to the second, except that firms are required to charge identical prices in the search and advertising markets. We denote this regime by \mathcal{ND} for “no discrimination.” In Appendix B, we also consider a variation of the model, in which the level of privacy protection is a continuous variable that is measured by the probability that consumers receive relevant ads. The qualitative results obtained from that variant are identical to those obtained from the baseline model.

Fix a regime. The timing of the model is as follows. First, firms set prices in the relevant markets. Next, consumers receive interest-based ads (unless the regime is \mathcal{P}). Finally, consumers decide either to buy the good through an ad they received (if they received one), to search for the good themselves, or to quit the market. Importantly, we assume that even though consumers receive ads only once, they can search as many times as they want. Moreover, we assume that the search is with perfect recall. That is, after each round of the search consumers can purchase from every price quote they obtained in the past (through ads or searches).

2.1 Strategies and Equilibrium

Consumers

By standard arguments, under all three regimes, consumers' strategies can be described by a single reservation price \tilde{p} such that a consumer who has obtained price quotes (through ads or searches) that are less than or equal to \tilde{p} buys the good immediately at the lowest such price.¹¹ The consumers' reservation price \tilde{p} is determined such that consumers are indifferent between purchasing the good at price \tilde{p} and the better option of quitting the market and searching one more time for additional price quotes, and then purchasing the good at the lowest price observed up to then.¹² This implies that the reservation price \tilde{p} satisfies the following condition:

$$v - \tilde{p} = \max \{0, v - c - \mathbb{E}(\min\{\tilde{p}, p_{min}\})\}, \quad (1)$$

where p_{min} is the lowest price quote that is obtained in one additional round of search (if no quotes are obtained in the next round of search, then we set $p_{min} = \infty$).

Firms

Firms' profits in each market (under any regime) depend on the distribution of the number of price quotes observed by consumers before purchasing (Burdett and Judd, 1983). We denote this endogenously generated equilibrium distribution by β . We distinguish between three different distributions of price quotes: a distribution of the number of price quotes in the search market (under regimes \mathcal{P} and \mathcal{D}), a distribution of the number of price quotes in the advertising market (under regime \mathcal{D}), and a distribution of the number of price quotes in the "merged market" (under regime \mathcal{ND}). We refer to the latter as the merged market since in the case where firms advertise and consumers search but firms quote the same price in both the search and advertising markets, it is possible to interpret the equilibrium as involving just one, merged, market. The formal description of these distributions is given in the relevant sections below.

The expected profit of a firm that quotes a price p , when the consumers' reservation price is \tilde{p} , in market $z \in \{S, A, M\}$ (where S, A, M represents the search, advertising, and

¹¹Because a consumer will compare the best available offer with the distribution of offers from search, the same reservation price is maintained in both the advertising and search markets.

¹²At the reservation price \tilde{p} a consumer is exactly indifferent between searching and not searching. Thus, it may be assumed that if the next round of the search does not generate a price that is lower than \tilde{p} , the consumer will not search again.

merged markets, respectively) is given by

$$\Pi^z(p) = \begin{cases} p\Omega^z \sum_{n=1}^{\infty} \beta_n^z n (1 - F^z(p))^{n-1} & \text{if } p \leq \tilde{p} \\ 0 & \text{if } p > \tilde{p} \end{cases} \quad (2)$$

where β_n^z is the equilibrium probability that a consumer observes n price quotes before purchasing, Ω^z is the mass of consumers that buy in market z , and $F^z(\cdot)$ is the distribution of prices in market z .¹³ Note that, in equilibrium, firms will offer prices that are no higher than \tilde{p} . Thus, a consumer will not search if he has received an ad, or will search until the first time his search yields at least one price quote. Therefore, the number of price quotes obtained by a consumer before purchasing is no more than two. That is, $\beta_n^z = 0$ for every $n \geq 3$ and $z \in \{S, A, M\}$.

Equilibrium

Our model admits the existence of an equilibrium with a degenerate search market in which consumers do not search because firms quote high prices, and firms quote high prices because consumers do not search.¹⁴ To abstract away from such a degenerate equilibrium, we focus only on equilibria in which consumers do actively search for the good in the search market. That is, we focus our attention on equilibria in which a consumer who has observed no price quotes would rather search for price quotes than quit the search market. Our model also admits a plethora of equilibria where search is limited due to the freedom to specify consumers' beliefs about firms' pricing after observing an out-of-equilibrium price offer. To overcome this issue we adopt the assumption of *passive beliefs* (see, e.g., Hart et al., 1990) commonly made in such settings. That is, if consumers observe an out-of-equilibrium price offer, they maintain their original beliefs.

In an equilibrium with active consumer search, it holds that

$$v - \tilde{p} = v - c - \mathbb{E}(\min\{\tilde{p}, p_{min}\}),$$

and so the consumers' reservation price solves the following simplified version of Equation (1):

$$\tilde{p} = c + \mathbb{E}(\min\{\tilde{p}, p_{min}\}). \quad (3)$$

¹³Observe that consumers who obtain multiple price quotes are over represented in the population of consumer-price quote pairs.

¹⁴These equilibria of the search market are similar to those described by Diamond (1971) in which firms quote the monopolistic price and consumers refrain from incurring the cost of search.

In equilibrium, firms take the consumers' reservation price \tilde{p} as given, and set prices optimally in each of the markets that exist under the prevailing regime. Under regime \mathcal{P} , only the search market exists; under regime \mathcal{D} , both an advertising market and a search market exist; and under regime \mathcal{ND} , only a merged advertising and search market exists. We denote the set of markets that exist under regime R by $Z(R)$.

Definition (Equilibrium with Active Consumer Search). *A tuple $\langle \tilde{p}, \{F^z(\cdot)\}_{z \in Z(R)} \rangle$ is an equilibrium with consumer search under regime R if:*

1. *The consumers' reservation price \tilde{p} solves Equation (3) given the distribution of firms' price quotes in the search market.¹⁵*
2. *Every price p that is quoted by firms under regime R maximizes firms' profits in the market in which it is quoted.*
3. *A search yields a nonnegative expected payoff to consumers who have received no previous price quotes.*

Note that an equilibrium with an active search market need not exist for all parameters. In the Section 3 we derive conditions that determine when an equilibrium with active consumer search exists under each of the three regimes that we consider. Henceforth, we refer to equilibria with active consumer search under regime R as equilibria.

Intensity of Competition

The price quoted by a firm in a given market depends on the number of additional price quotes that a consumer who has been exposed to the firm's own price quote is expected to observe. In the degenerate case in which a firm believes that it will be the only firm that a consumer is exposed to, the firm will act as a monopolist and quote the highest price that consumers are willing to pay, i.e., \tilde{p} . In the other extreme case, in which a firm believes that a consumer who is exposed to its price quotes will also observe (at least) one other price quote, the firm will act as it would in a competitive market and quote a price that is equal to its marginal cost, i.e., zero. In our model, firms are uncertain about whether a consumer was exposed to an additional price quote. As shown by Burdett and Judd (1983), in this case firms adopt a mixed strategy in equilibrium. In this equilibrium, both the price distribution and the consumers' reservation price are jointly determined by the

¹⁵Under regime \mathcal{ND} , these are the price quotes in the merged market.

distribution of the number of price quotes obtained by the consumer before purchasing the good.

As the probability that consumers obtain multiple price quotes increases, the market becomes more competitive and market prices decrease. To gain more insight into the roles played by the probabilities $\{q_0^S, q_1^S, q_2^S\}$ and $\{q_0^A, q_1^A, q_2^A\}$, note that the probability that consumers obtain no price quotes per search, q_0^S , is essentially a multiplier of the consumers' search costs because if consumers do search, then they will do so until they obtain at least one price quote. Thus, the consumers' "effective search costs" are given by $\frac{c}{1-q_0^S}$ and the "effective probabilities" of observing one or two price quotes per consumer search are given by $\frac{q_1^S}{q_1^S+q_2^S}$ and $\frac{q_2^S}{q_1^S+q_2^S}$, respectively. Similarly, the probability that consumers receive no ads, q_0^A , is also a parameter that is not directly related to the competitiveness of the advertising market. However, unlike q_0^S , it is possible to interpret q_0^A as an inverse measure of the "size of the advertising market" because the measure of consumers who receive at least one ad and, in equilibrium, go on to purchase the good without searching for it is $1 - q_0^A$. As in the search market, the "effective probabilities" of observing one or two price quotes through interest-based ads are $\frac{q_1^A}{q_1^A+q_2^A}$ and $\frac{q_2^A}{q_1^A+q_2^A}$, respectively.

The above discussion suggests that once the consumers' effective search costs and the size of the advertising market are fixed, we may use $\frac{q_2^S}{q_1^S+q_2^S}$ and $\frac{q_2^A}{q_1^A+q_2^A}$ as measures of the competitiveness of the search and advertising markets, respectively. Accordingly, for the rest of this paper we use these ratios as measures of the competitiveness of the search and advertising markets, respectively.

2.2 Discussion of Assumptions

At Most Two Price Quotes. Our assumption that a consumer receives at most two price quotes in each round is a simplifying assumption that enables us to obtain closed-form solutions to equilibrium outcomes.¹⁶ In Section 5 we explain why our main qualitative results do not depend on this assumption. That being said, we note that previous work suggests that the probability with which a consumer obtains exactly one price quote is the most important statistic of the distribution of the number of price quotes. Burdett and Judd (1983) show that this probability is what determines whether a firm will make a

¹⁶In particular, this assumption enables us to derive a closed-form solution of the reservation price in regimes \mathcal{P} and \mathcal{D} . Moreover, it allows us to derive Lemma 1 in which we calculate the ratio between the expected price paid by a consumer and the reservation price in regime \mathcal{ND} .

profit, and, if the reservation price is held constant, what this profit will be. More recently, Bergemann, Brooks and Morris (2021) have shown that the probability that the consumer observes one price quote forms an upper bound on firms' profits that is robust to the firms' information structure.

Sequence of Ads and Searches. Our assumption that consumers receive ads before they search aims to capture the following sequence of events. First, an event that creates an interest in some specific good occurs. For example, a consumer realizes that he will need to purchase a good by some future date. At the time when the interest arises, the consumer plans to search (and purchase) the good at some later point in time. In the interim, until the time when the consumer plans to actively search for the good, the consumer may engage with (relevant) ads that he receives. Our model abstracts away from the choice of when to purchase the good, but implicitly assumes that there is some fixed time by which the good needs to be purchased. The existence of such a deadline prevents the consumer from obtaining an infinite number of ads before purchasing, and is consistent with our assumption that the consumer can choose to search multiple times even though he receives ads only once. A noteworthy consequence of this assumption is that, in our model, firms cannot learn about a consumer's preferences through his search activities (such learning is the main focus of De Cornière, 2016).

Impact of Privacy on the Effectiveness of Search. In practice, the efficiency of consumer search when firms have information about the consumer may be greater than when they do not have this information. In such a case, privacy protection would hurt consumers and benefit firms. However, so long as the effect of privacy on search effectiveness is small, this would not have a qualitative impact on our results.

Privacy Protection is Not a Consumer's Choice. In our model the decision of whether or not to protect consumer privacy is a social choice. It is not a decision that each consumer makes individually for himself. In fact, in our model each consumer would benefit (individually) from sharing his data because this would not affect equilibrium prices, and would provide the consumer with more opportunities to purchase the good. Thus, our notion of privacy should be interpreted as allowing a regulator to control the manner in which firms use consumer data, rather than as allowing each individual consumer to control how his data is used by firms. Hence, our notion of privacy is distinct from the way privacy is treated in many of the recent laws that protect consumer privacy. For example, the EU's General Data Protection Regulation—in particular, Chapter 3 thereof—

establishes the right of the “data subject” to control how his data is used. We discuss this further in our concluding remarks (Section 5).

3 Equilibrium Analysis

In this section we characterize the equilibria of our model under all three regimes.

3.1 Equilibrium Analysis under Regime \mathcal{P}

When consumer privacy is protected, only the search market exists, and so the analysis is almost identical to that performed by Burdett and Judd (1983). In fact, a minor modification of their Theorem 4 establishes the existence of a unique equilibrium with active consumer search. Our assumption that the number of price quotes obtained in each round of consumer search is at most two enables us to derive a simple condition that characterizes when it is profitable for consumers to engage in a search and, moreover, to derive closed-form expressions of firms’ profits and consumer surplus in this equilibrium.

Proposition 1. *Consider regime \mathcal{P} . There exists an equilibrium with active consumer search if and only if*

$$v \cdot q_2^S \geq c. \quad (4)$$

Moreover, if Condition (4) holds, then there exists a unique equilibrium with active consumer search in which the consumer surplus is

$$CS^{\mathcal{P}} = v - \frac{c}{q_2^S},$$

and the firms’ profits are

$$\pi^{\mathcal{P}} = \frac{q_1^S}{1 - q_0^S q_2^S} \frac{c}{q_2^S}.$$

Condition (4) has a simple interpretation: consumers engage in an active search if and only if the probability of obtaining two price quotes is high enough. In other words, the consumers engage in an active search if and only if the search market is competitive enough to avoid unraveling as in the Diamond paradox (Diamond, 1971), where firms quote the monopolistic price and consumers refrain from incurring the cost of search.

Proposition 1 shows that if consumers engage in an active search, then the total surplus in the search market is given by

$$CS^{\mathcal{P}} + \pi^{\mathcal{P}} = v - \frac{c}{1 - q_0^S}.$$

Because in equilibrium all the consumers purchase the good and receive a payoff of v , social welfare is equal to v minus the consumers' expected search costs until obtaining at least one price quote, i.e., $\frac{c}{1-q_0^S}$. It follows that if the value of q_0^S is held constant, then increasing the competitiveness of the market by increasing the value of q_2^S at the expense of q_1^S transfers surplus from firms to consumers.

3.2 Equilibrium Analysis under Regime \mathcal{D}

To characterize the equilibrium under this regime (in which both the search and advertising markets coexist), note that the consumers' reservation price is determined by the equilibrium distribution of prices in the search market. Moreover, the search market under this regime is identical to the search market under regime \mathcal{P} , except that the mass of consumers in the search market is q_0^A rather than one.

Because the size of the search market (the fraction of consumers who engage in an active search) does not affect the firms' pricing decisions, the consumers' reservation price and the firms' pricing decisions in the search market are identical under regimes \mathcal{P} and \mathcal{D} . Moreover, the same calculations that were used to establish Proposition 1 can be used to determine when an equilibrium with active consumer search exists, to characterize the equilibrium in the advertising market, and to derive the consumer surplus and firms' profits.

Proposition 2. *Consider regime \mathcal{D} . There exists an equilibrium with active consumer search if and only if Condition (4) holds. If Condition (4) holds, then there exists a unique equilibrium with active consumer search in which the consumer surplus is*

$$CS^{\mathcal{D}} = v - \frac{c(1 - q_2^A)}{q_2^S},$$

and the firms' profits are

$$\pi^{\mathcal{D}} = \left(q_1^A + q_0^A \frac{q_1^S}{1 - q_0^S} \right) \frac{c}{q_2^S}.$$

Note, first, that social welfare is given by

$$CS^{\mathcal{D}} + \pi^{\mathcal{D}} = v - \frac{q_0^A}{1 - q_0^S} c.$$

Thus, social welfare is increasing in the size of the advertising market $1 - q_0^A$. This result is due to the fact that in equilibrium consumers purchase the good when they receive

their first price quote. Hence, those consumers who receive an ad avoid the search cost, while those who do not receive an ad incur an expected search cost of $\frac{c}{1-q_0^S}$.

Second, as was the case under regime \mathcal{P} , if the consumers' effective search costs and the size of the advertising market are held constant, increasing the competitiveness of either market by increasing the value of q_2^S at the expense of q_1^S or by increasing the value of q_2^A at the expense of q_1^A transfers surplus from firms to consumers.

3.3 Equilibrium Analysis under Regime \mathcal{ND}

The characterization of the equilibrium in the merged market is more complicated than under the previous two regimes. The difficulty is due to the fact that the distribution of the number of price quotes that a consumer observes before buying the good (which is the object that determines the firms' pricing decisions) is different from the distribution of the number of price quotes generated by an additional round of search (which is the object that determines the consumers' reservation price).

To derive the equilibrium, we first characterize the expected price that a consumer pays for the good if he decides to search for it. That is, we derive the expectation of the lowest price quote that the consumer receives, conditional on receiving at least one price quote. Let q^S denote the vector $\langle q_0^S, q_1^S, q_2^S \rangle$ and recall that β_1 denotes the fraction of consumers that observe one price quote before buying the good.

Lemma 1. *In equilibrium, the ratio between the expected price that is paid by a consumer who searches for price quotes and his reservation price is*

$$\kappa(q^S, \beta_1) = \frac{\beta_1 \left(2q_2^S \left(\beta_1 \left(1 + \tanh^{-1}(1 - \beta_1) \right) - 1 \right) + (1 - \beta_1)q_1^S \log \left(\frac{\beta_1}{2 - \beta_1} \right) \right)}{2(1 - \beta_1)^2(q_0^S - 1)}$$

and $\tanh^{-1}(\cdot)$ denotes the inverse of the hyperbolic tangent function.

Lemma 1 implies that $\mathbb{E}[p_{min} | \text{at least one price quote}] = \tilde{p} \cdot \kappa(q^S, \beta_1)$. It allows us to characterize when it is profitable for consumers to engage in an active search, and to derive the consumer surplus and the firms' profits. If a consumer engages in an active search then he obtains exactly one price quote before he purchases the good in two cases: when he receives one ad, and when he receives no ads and the first round of the search that generates a price quote yields exactly one price quote. It follows that, when the search market exists, $\beta_1^{\mathcal{ND}} = q_1^A + q_0^A \frac{q_1^S}{1 - q_0^S}$.

Proposition 3. *Consider regime \mathcal{ND} . There exists an equilibrium with active consumer search if and only if*

$$v \geq \frac{c}{(1 - q_0^S) (1 - \kappa(q^S, \beta_1^{\mathcal{ND}}))}. \quad (5)$$

If Condition (5) holds, then there exists a unique equilibrium with active consumer search in which the consumer surplus is

$$CS^{\mathcal{ND}} = v - \frac{c}{1 - q_0^S} \left(\frac{\kappa(q^S, \beta_1^{\mathcal{ND}})}{1 - \kappa(q^S, \beta_1^{\mathcal{ND}})} + q_0^A \right),$$

and the firms' profits are

$$\pi^{\mathcal{ND}} = \frac{c}{1 - q_0^S} \frac{\kappa(q^S, \beta_1^{\mathcal{ND}})}{1 - \kappa(q^S, \beta_1^{\mathcal{ND}})}.$$

Note that when consumers engage in an active search, social welfare is the same as under regime \mathcal{D} , and it is given by $v - \frac{c}{1 - q_0^S} q_0^A$. This is because under both regimes all the consumers purchase the good and only those consumers who do not receive an ad must incur the search cost. Moreover, Lemma A.1 in Appendix A shows that $\kappa(q^S, \beta)$ is increasing in β . It follows that, as in the previous two regimes, if the consumers' effective search costs and the size of the advertising market are held constant, increasing the competitiveness of either the advertising market or the search market transfers surplus from firms to consumers.

4 Welfare Analysis

In this section we analyze how protection of consumer privacy affects consumer surplus and firms' profits. In particular, we compare the equilibrium payoffs under regime \mathcal{P} with those of regimes \mathcal{D} and \mathcal{ND} . In the remainder of the paper we assume that Conditions (4) and (5) hold, i.e., that there exists a unique equilibrium with active consumer search under all three regimes.

The Value of Privacy to Consumers

First, we evaluate the benefit of privacy to consumers when firms price discriminate between the search and advertising markets. As explained in Section 3, the equilibrium in the search market is identical under regimes \mathcal{P} and \mathcal{D} and, in particular, the consumers'

(optimal) reservation price is the same under both regimes. It follows that privacy protection hurts consumers because it denies them opportunities to buy the good, without affecting their value from search.

Proposition 4. *Suppose that Condition (4) holds. If firms price discriminate between the search and advertising markets, then privacy protection hurts consumers.*

When firms price discriminate privacy protection has an unambiguous negative effect on consumer surplus, whereas when firms do not price discriminate (regime \mathcal{ND}) the effect of privacy protection on consumer surplus depends on the size and competitiveness of the advertising market. On the one hand, privacy protection prevents firms from initiating contacts with consumers, which increases consumers' search expenditures. On the other hand, privacy protection changes the competitiveness of the market (as measured by the fraction of consumers who observe more than one price quote before purchase). If the advertising market is more competitive than the search market, then privacy protection decreases market competitiveness and unambiguously hurts consumers. However, if the advertising market is less competitive than the search market, (which, as discussed in the Introduction, seems to be the more likely case in practice), then privacy protection increases market competitiveness and benefits consumers unless the implied increase in their search expenditure outweighs the increase in market competitiveness.¹⁷

Proposition 5. *Fix the values of v, c , and q^S . Suppose that Conditions (4) and (5) hold, and assume that firms do not price discriminate.*

1. *If the advertising market is less competitive than the search market, then privacy protection increases consumer welfare if and only if the size of the advertising market $(1 - q_0^A)$ is greater than some critical threshold, which is increasing in the competitiveness of the advertising market.*
2. *If the advertising market is (weakly) more competitive than the search market, then privacy protection decreases consumer welfare.*

The first part of Proposition 5 is illustrated in Figure 1 below. In this figure, we depict how the size and competitiveness of the advertising market impact the value of privacy to consumers: in the red region, privacy protection hurts consumers; in the blue region, privacy protection benefits consumers; and in the white region consumers do not search.

¹⁷Observe that market competitiveness is given by $\frac{q_2^S}{q_1^S + q_2^S}$ under regime \mathcal{P} , and by $(1 - q_0^A) \frac{q_2^A}{q_1^A + q_2^A} + q_0^A \frac{q_2^S}{q_1^S + q_2^S}$ under regime \mathcal{ND} .

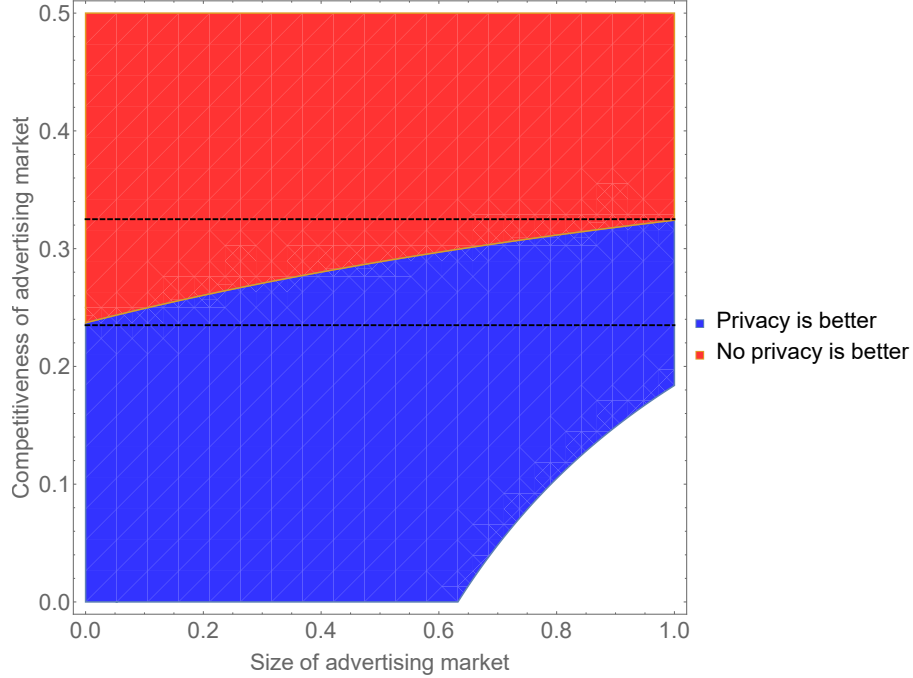


Figure 1: Value of privacy to consumers for $v = 5, c = 1, q_1^S = q_2^S = 1/2$.

Figure 1 highlights the fact that the effect of privacy protection on consumer welfare depends not only on the competitiveness of the advertising market but also on its size. Intuitively, without privacy protection, firms' pricing choices are determined by the *competitiveness of the merged market*. When the putative advertising market is the less competitive market, opening it up (by removing privacy protection) will reduce market competitiveness. Moreover, as the size of this putative market increases – while keeping its competitiveness fixed – the reduction in competitiveness will be more pronounced. However, a priori, a larger putative advertising market need not hurt consumers, as the consumers' search expenditure is decreasing in the size of the advertising market.

If the competitiveness of the advertising market is sufficiently similar to that of the search market (above the upper dashed line in Figure 1), then, regardless of the size of the advertising market, the reduction in search expenditure is the dominant force, whereas if the advertising market is extremely noncompetitive (below the lower dashed line in Figure 1), then the reduction in market competitiveness is the dominant force. Importantly, the second observation implies that the mere existence of a small and noncompetitive advertising market reduces consumer welfare.

If the level of competitiveness in the advertising market is intermediate (between the two dashed lines in Figure 1), then whether or not privacy protection is beneficial for

consumers depends on the size of the advertising market. Our result shows that if the advertising market is small, the reduction in search expenditure is enough to compensate consumers for the lower level of competition in the merged market. However, if the advertising market is large, the reduction in search costs is not enough to compensate consumers for the decrease in competition. That is, for intermediate levels of competitiveness, privacy protection increases consumer welfare only if the advertising market is large enough.

Figure 1 also suggests that there exists a consumer-optimal size of the advertising market. Intuitively, consider an intermediate level of competition in the advertising market. If the size of the advertising market is exactly equal to the critical threshold mentioned in Proposition 5, then consumer welfare is the same as when the advertising market does not exist. However, for any size between these two extremes, consumers strictly benefit from the existence of the advertising market. This suggests that, for a given level of competition in the advertising market, there is a consumer-optimal size of the advertising market. This idea is visualized in Figure 2 below, which depicts consumer surplus as a function of the size of the advertising market for regimes \mathcal{P} and \mathcal{ND} .




Figure 2: Affect of advertising market size on consumer surplus for $v = 5, c = 1, q_1^S = q_2^S = 1/2$, and $\frac{q_2^A}{q_1^A + q_2^A} = 0.3$.

If we were to assume that privacy protection is continuous – i.e., that regulation could determine the mass of consumers who receive ads (see formal definition in Appendix B) – then this idea would indicate that when the advertising market is less competitive than the search market, there exists an optimal (interior) level of privacy protection.

The Value of Privacy for Firms

Firms' profits in the search and advertising markets depend on the competitiveness of these markets: As the market becomes less competitive, i.e., as $\frac{q_1^z}{q_1^z + q_2^z}$, $z \in \{S, A\}$, in-

creases, firms' profits in that market increase as well. The following proposition establishes a stronger result. It shows that firms benefit from advertising if and only if the search market is more competitive than the advertising market. That is, if the search market is more competitive than the putative advertising market, then privacy protection decreases firms' profits regardless of whether or not firms price discriminate between the search and advertising markets, the magnitude of the consumers' search cost, and the size of the advertising market.

Proposition 6. *Suppose that Conditions (4) and (5) hold. Privacy protection decreases firms' profits if and only if $\frac{q_1^A}{q_1^A + q_2^A} > \frac{q_1^S}{q_1^S + q_2^S}$.*

To see why Proposition 6 holds true, note that if the search and advertising markets are distinct (as in regime \mathcal{D}), then the consumers' reservation price does not change if consumer privacy is not protected. Thus, opening up the advertising market (by removing privacy protection) benefits firms if and only if the advertising market is less competitive than the search market. This is the case because the consumers' reservation price and the prices in the search market are unaffected by the opening up of the advertising market, and some of the transactions performed in the search market move to the advertising market, where prices are higher on average.

If firms do not price discriminate and the search and advertising markets are merged (regime \mathcal{ND}), then the explanation becomes more involved because if firms do not price discriminate, then the removal of privacy protection can affect the consumers' reservation price and hence the distribution of prices in the search market. However, if the advertising market is less competitive than the search market, then removal of privacy protection leads to an increase in market prices, which, in turn, reduces the value of each round of search, increases the consumers' reservation price, and leads to an additional increase in prices. Similarly, if the advertising market is more competitive than the search market, then removal of privacy protection leads to a decrease in the consumers' reservation price, which further reduces firms' profits.

The Social Value of Privacy

Even without consumer privacy protection, the information that firms hold about consumers is extremely general and vague (see the discussion in the Introduction and in particular footnote 3), and so it is plausible that it should be easier for consumers to find firms that produce the goods that they are interested in than for firms to find the con-

sumers who are interested in the specific goods they produce through interest-based advertising. Therefore, it stands to reason that, in many settings, the search market would be more competitive than the advertising market. That is, in many settings, the if and only if condition in Proposition 6 is likely to be satisfied. In such cases, combining the results of Propositions 4 and 6 implies that if firms engage in price discrimination between the search and advertising markets, then the protection of consumer privacy is Pareto inferior to non-protection.

Corollary 1. *Assume that Condition (4) holds and that $\frac{q_1^S}{q_1^S+q_2^S} < \frac{q_1^A}{q_1^A+q_2^A}$. Equilibrium outcomes under regime \mathcal{D} Pareto dominate equilibrium outcomes under regime \mathcal{P} .*

The intuition for this result is based on three simple observations. First, privacy protection decreases welfare because it increases consumers' search expenditures. Second, if firms price discriminate, then privacy protection cannot benefit consumers. Third, if the search market is more competitive than the advertising market, $\frac{q_1^S}{q_1^S+q_2^S} < \frac{q_1^A}{q_1^A+q_2^A}$, then, for any fixed reservation price, firms' profits are greater in the advertising market than in the search market.

It is worth emphasizing that even if the advertising market is almost monopolistic, privacy protection is still Pareto dominated. For example, if $q_1^A = 9/10, q_0^A = q_2^A = 1/20, v = 5, c = 1$, and $q_1^S = q_2^S = 1/2$, then privacy protection lowers consumer surplus from 3.1 to 3, and decreases firms' profits from 1.85 to 1.

Even if firms do not engage in price discrimination, privacy protection can be Pareto dominated. We illustrate this claim graphically in Figure 3 below. In the figure, the green region indicates the advertising markets for which equilibrium outcomes in regime \mathcal{ND} Pareto dominate equilibrium outcomes in regime \mathcal{P} .

In the area above the lower boundary of the green region, consumers prefer regime \mathcal{ND} to regime \mathcal{P} , whereas in the area below the upper boundary of the green region, firms prefer regime \mathcal{ND} to regime \mathcal{P} . To see why these two regions must intersect – and create a nonempty green region – observe that if the advertising market is exactly as competitive as the search market ($\frac{q_1^A}{q_1^A+q_2^A} = \frac{q_1^S}{q_1^S+q_2^S}$), then market competitiveness is unaffected by the protection of consumer privacy. It follows that the distribution of prices is the same when privacy is protected and when it is not. This, in turn, implies that firms' profits are identical under regimes \mathcal{P} and \mathcal{ND} . However, at this critical level of competitiveness consumers strictly prefer regime \mathcal{ND} to regime \mathcal{P} because their search expenditures are lower under the former regime than under the latter regime. It follows that if the competitiveness of the advertising market market is slightly below this critical

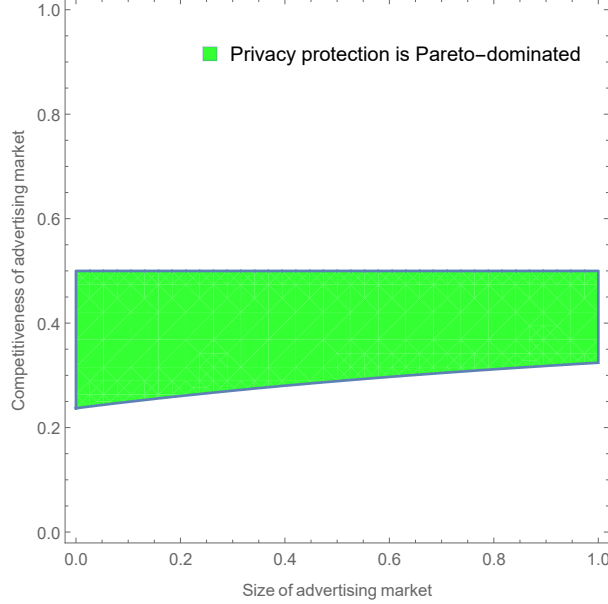


Figure 3: Social value of no privacy protection when $v = 5, c = 1, q_1^S = q_2^S = 1/2$.

level, then privacy protection is Pareto dominated by no privacy protection regardless of whether or not firms price discriminate.

So far we have focused on the question of whether and under what market conditions the protection of privacy benefits consumers. In some cases, it may be impossible to protect consumer privacy but possible to prevent price discrimination. Whether doing so is beneficial to consumers depends on the comparison between regimes \mathcal{D} and $\mathcal{N}\mathcal{D}$.

Corollary 2. *Suppose that Conditions (4) and (5) hold. Consumer surplus is higher under regime \mathcal{D} than under regime $\mathcal{N}\mathcal{D}$ if and only if firms' profits are lower under regime \mathcal{D} than under regime $\mathcal{N}\mathcal{D}$.*

Corollary 2 shows that if price discrimination benefits firms, then it hurts consumers. Thus, in this case, regulators who are concerned about consumer surplus should limit firms' ability to price discriminate. On the other hand, if price discrimination hurts firms, then they would like to commit to not engage in price discrimination.

Importantly, the last observation relies on the assumption that firms collude in their refusal to engage in price discrimination: if the regime is $\mathcal{N}\mathcal{D}$, then in the generic case where the search and advertising markets differ in their level of competitiveness, it is profitable for each individual firm to price discriminate. However, because it seems that in many markets price discrimination is frowned upon by consumers, it may not be possible for an individual firm to engage in price discrimination in markets in which this is

not an accepted practice. For example, Amazon – the largest online retailer in the world – publicly committed to not price discriminate among consumers.¹⁸ Once a market leader such as Amazon commits to refraining from price discrimination, it may be difficult for other, smaller, firms to price discriminate. Arguably, such an effect would be stronger in markets that are more subject to public scrutiny because of their size, prevalence, or general importance.

Furthermore, empirical evidence suggests that price discrimination is rather limited in practice. For example, a recent Ipsos, London Economics and Deloitte (2018) field survey conducted on behalf of the European Union found that personalized pricing “was observed in only 6% of matched identical product pairs. Even when price differences were observed, the differences were small, with the median difference being less than 1.6%” (p. 260).¹⁹

5 Concluding Remarks

More Than Two Price Quotes. We have assumed that consumers cannot receive more than two ads, or observe more than two price quotes per search. In no case does the intuition for our main qualitative results rely on this assumption. This suggests that this assumption, which allowed us to obtain an explicit solution to the consumers’ reservation price, is not essential for our results to hold.

Indeed, note that the result that the combination of price discrimination and no privacy protection is beneficial to consumers (Proposition 4) relies on the fact that interest-based advertising expands the consumers’ choice set without impacting equilibrium outcomes in the search market. Hence, this result clearly holds in more general environments as well. Moreover, the result that if the size of the advertising market is not too large, then consumers are hurt by privacy protection even if firms do price discriminate should also generalize. To see this, note that regardless of the number of possible price quotes received by consumers, a firm’s equilibrium profit is equal to the product of the probability that a consumer observes one price quote and his reservation price. Thus, if $\frac{q_1^A}{1-q_0^A} = \frac{q_1^S}{1-q_0^S}$,

¹⁸In September 2000 Amazon famously outraged consumers by charging different prices to different consumers for the same DVDs. In response to the scandal that erupted, Amazon’s CEO Jeff Bezos committed to charge a single price for each good on Amazon (see, e.g., “Amazon apologizes for price-testing program that angered customers,” by Todd R. Weiss for Computerworld (9/28/2000) and “Web sites change prices based on customers’ habits,” by Anita Ramasastry for CNN.com (6/24/2005)).

¹⁹This finding is consistent with previous findings. See the references in the aforementioned report.

consumers are hurt by protection of privacy regardless of the size of the market. Moreover, since consumer surplus is continuous in the size of the advertising market and its competitiveness, it follows that for lower levels of competitiveness in the merged market, there exists a critical size of the advertising market under which consumers will benefit from advertising.

Second, the result that firms are hurt by privacy protection if and only if the probability that consumers observe one price quote is higher in the advertising market than in the search market (Proposition 6) relies on the fact that a firm's equilibrium profit is equal to the product of the probability that a consumer observes one price quote and his reservation price, which, as mentioned above, holds generally. This, in turn, suggests that the results that show that consumer privacy protection is Pareto-dominated (Corollary 1 and the following claim) should also hold more generally.

Costly Advertising. In this paper we abstracted away from the interaction between firms and the advertisers who create the ads that firms use to contact consumers. In particular, we did not consider how the cost of advertising impacts our welfare results.²⁰ The value of a single ad to a firm is the expected number of consumers who will observe the ad multiplied by the minimum price charged in the advertising/merged market. Hence, if the marginal cost of creating each ad is below this value, our characterization of the equilibrium in the interaction between firms and consumers remains unchanged. Costly advertising clearly decreases the firms' profits obtained from removing privacy protection, and so it may impact our welfare analysis. However, if the marginal cost of advertising is small, our qualitative results remain unchanged. Similarly, we assumed that it is costless for a consumer to examine a product for which he received an advertisement. If it were costly for a consumer to do so, then the reduction in search expenditure due to the removal of privacy protection would be lower than in our baseline model. Nevertheless, if the cost of examining ads is smaller than the effective cost of search, our qualitative results would remain valid.

The Social Aspect of Privacy. An important difference between the aspect of consumer privacy that we studied in this paper and the aspect of consumer privacy that is related to standard price discrimination through personalized pricing is that in our setting the individual consumer always benefits from sharing his personal data. Intuitively, in a general

²⁰The cost of advertising may also include the cost of processing consumer information in order to target the ads.

equilibrium environment an individual consumer's choice does not affect the pricing decisions of firms, and so sharing data and receiving interest-based ads has no downside for an individual consumer. However, as we have shown, if all consumers share their data then this affects market structure and prices in a way that may hurt consumers. This insight suggests that privacy laws that give consumers control over their personal information (e.g., the EU's "General Data Protection Regulation") may be insufficient to protect consumers and that more comprehensive regulation that considers the equilibrium consequences of consumer privacy regulation may be needed.²¹

Product Differentiation and Consumer Heterogeneity. In our model there is one good in which the relevant consumers are interested. Privacy protection determines firms' access to these consumers. A model with multiple goods, in which each consumer is interested in a single good, gives rise to another natural notion of privacy. In such a model, the information that could be kept private is the specific good in which each consumer is interested.

It is straightforward that if the economy were made up of multiple replicas of our model – one replica for each good – then the analysis of each replica would be identical to the analysis in this paper. This implies that all of our results would also hold in a setting with the aforementioned type of differentiated goods. In fact, even if consumers' were to obtain a low value from consuming their non-preferred goods, our results would hold so long as, in equilibrium, consumers prefer searching for their preferred good to buying one of their non-preferred goods.

An additional simplifying assumption in our model is that all the consumers have the same valuation for the good. However, the intuition that suggests that relaxing privacy protection is beneficial to consumers would remain valid in settings in which there is heterogeneity in consumers' valuations. Moreover, these forces may be augmented by an additional mechanism. In particular, if consumers have a heterogeneous value for the good, consumers with lower valuations will search more often. It follows that opening up the advertising market (by removing privacy protection) will result in the search market having a lower distribution of valuations than in the general population. If lower consumers' valuations lead to a decrease in firms' prices, then opening up the advertis-

²¹The "public good" aspect of privacy has also been established by Garratt and van Oordt (2021). They show that keeping payment histories private can increase consumer surplus by preventing price discrimination through personalized pricing, even though, if the firm's pricing strategies are held constant, each consumer individually benefits from disclosing his payment history.

ing market should lead to a reduction in prices in the search market, provided that firms engage in price discrimination. That is, heterogeneity in consumers' willingness to pay should increase the value of removing privacy protection when firms price discriminate.

A Appendix: Proofs

Proof of Proposition 1

If consumers are willing to pay the search cost to obtain a random number of price quotes, they search until they obtain at least one price quote. Hence, the fraction of consumers that observe n price quotes before purchasing is

$$\beta_n^{\mathcal{P}} = \frac{q_n^S}{1 - q_0^S}.$$

Burdett and Judd (1983) establish that since $\beta_1^{\mathcal{P}} \in (0, 1)$, in equilibrium, the price distribution is continuous, and that its support has an upper bound of $\tilde{p}^{\mathcal{P}}$, where $\tilde{p}^{\mathcal{P}}$ is the reservation price under regime \mathcal{P} . Moreover, by setting this price, the firm will sell only if it is the only quote a consumer obtains. Hence, since the firms must be indifferent between all prices in the support of $F^S(\cdot)$, the expected profit for each firm is $\beta_1^{\mathcal{P}} \tilde{p}^{\mathcal{P}}$. By the definition of a firm's profit given in Equation (2), it follows that the equilibrium price distribution is given by

$$F^S(p) = \frac{(2 - \beta_1^{\mathcal{P}})p - \beta_1^{\mathcal{P}} \tilde{p}^{\mathcal{P}}}{2(1 - \beta_1^{\mathcal{P}})p},$$

where the lower bound of the support of this distribution is $\underline{p}^{\mathcal{P}} = \frac{\tilde{p}^{\mathcal{P}} \beta_1^{\mathcal{P}}}{2 - \beta_1^{\mathcal{P}}}$.

Next, we derive the consumers' reservation price under regime \mathcal{P} . Denote by $f^S(p)$ the density of $F^S(p)$ and by $j^S(p) = \frac{d(2F^S(p) - (F^S(p))^2)}{dp}$ the density of the minimum of two (independent) price quotes from the distribution $F^S(p)$. The consumer's expected payment for the good after one round of search that generates at least one quote is given by

$$\beta_1^{\mathcal{P}} \int_{\underline{p}^{\mathcal{P}}}^{\tilde{p}^{\mathcal{P}}} p f^S(p) dp + \beta_2^{\mathcal{P}} \int_{\underline{p}^{\mathcal{P}}}^{\tilde{p}^{\mathcal{P}}} p j^S(p) dp = \tilde{p}^{\mathcal{P}} \beta_1^{\mathcal{P}}, \quad (\text{A.1})$$

where the equality is obtained by plugging the previously derived expression into the integral and solving it.

Since firms' prices are never higher than $\tilde{p}^{\mathcal{P}}$, the reservation price can be calculated by assuming that if the next search generates no quotes, then the consumer will pay $\tilde{p}^{\mathcal{P}}$

for the good and otherwise he will pay as per one of the new quotes. Since his expected payment will be $\tilde{p}\beta_1^{\mathcal{P}}$, Equation (3) simplifies to

$$v - \tilde{p}^{\mathcal{P}} = q_0^{\mathcal{S}}(v - \tilde{p}^{\mathcal{P}}) + (1 - q_0^{\mathcal{S}})(v - \tilde{p}^{\mathcal{P}}\beta_1^{\mathcal{P}}) - c,$$

and the reservation price is

$$\tilde{p}^{\mathcal{P}} = \frac{c}{q_2^{\mathcal{S}}}.$$

Since consumers might have to search multiple times to obtain (at least) one quote, their expected search cost is $\frac{c}{1-q_0^{\mathcal{S}}}$, which, in turn, implies that consumer surplus is given by

$$CS^{\mathcal{P}} = v - \frac{c}{1-q_0^{\mathcal{S}}} - \tilde{p}^{\mathcal{P}}\beta_1^{\mathcal{P}} = v - \frac{c}{q_2^{\mathcal{S}}} \geq 0.$$

Finally, to obtain the firms' equilibrium profits recall that these profits are given by $\beta_1^{\mathcal{P}}\tilde{p}^{\mathcal{P}} = \frac{q_1^{\mathcal{S}}}{1-q_0^{\mathcal{S}}}\frac{c}{q_2^{\mathcal{S}}}$. \square

Proof of Proposition 2

First, note that firms will set a price no higher than $\tilde{p}^{\mathcal{D}}$ for consumers that arrive via searches or ads, where $\tilde{p}^{\mathcal{D}}$ is the reservation price under regime \mathcal{D} . Note that the search market is the same under this regime as under regime \mathcal{P} , except that it has a mass $q_0^{\mathcal{A}}$ of consumers rather than a mass of one. However, this scaling does not affect equilibrium in the search market, which, in turn, implies that $\tilde{p}^{\mathcal{D}} = \tilde{p}^{\mathcal{P}}$. Hence, the consumer surplus that originates from the fraction $q_0^{\mathcal{A}}$ of consumers that do not receive ads is $q_0^{\mathcal{A}}CS^{\mathcal{P}}$.

A fraction $1 - q_0^{\mathcal{A}}$ of consumers receive ads and purchase a good through the cheapest ad they receive. From the same calculation used to solve Equation (A.1), it follows that the expected payment such consumers make for the good is $\frac{q_1^{\mathcal{A}}}{1-q_0^{\mathcal{A}}}\tilde{p}^{\mathcal{D}}$. Hence, the consumer surplus that originates from the fraction $1 - q_0^{\mathcal{A}}$ of consumers that receive ads is $(1 - q_0^{\mathcal{A}})(v - \frac{q_1^{\mathcal{A}}}{1-q_0^{\mathcal{A}}}\tilde{p}^{\mathcal{D}})$. Thus, consumer surplus is given by

$$CS^{\mathcal{D}} = q_0^{\mathcal{A}}CS^{\mathcal{P}} + (1 - q_0^{\mathcal{A}})(v - \frac{q_1^{\mathcal{A}}}{1-q_0^{\mathcal{A}}}\tilde{p}^{\mathcal{D}}) = v - \frac{c(1 - q_2^{\mathcal{A}})}{q_2^{\mathcal{S}}}.$$

Next, we calculate the firms' profits. The derivation in Proposition 1 implies that with a measure one of firms and a measure S of consumers that engage in a noisy search using a reservation price \tilde{p} , and in which β of the consumers observe one price quote before purchasing, firms' profits are given by $S\beta\tilde{p}$.

The size of the advertising market is $(1 - q_0^{\mathcal{A}})$ and in this market $\beta_1^{\mathcal{A}} = \frac{q_1^{\mathcal{A}}}{1-q_0^{\mathcal{A}}}$. Hence, the firms' profits in this market are $(1 - q_0^{\mathcal{A}})\beta_1^{\mathcal{A}}\tilde{p}^{\mathcal{D}} = q_1^{\mathcal{A}}\tilde{p}^{\mathcal{D}}$. Similarly, the size of the search

market is q_0^A and in this market $\beta_1^S = \frac{q_1^S}{1-q_0^S}$, and so the firms' profits in the search market are $q_0^A \beta_1^S \tilde{p}^D = q_0^A \frac{q_1^S}{1-q_0^S} \tilde{p}^D$. Thus, firms' profits are given by

$$\pi^D = q_1^A \tilde{p}^D + q_0^A \frac{q_1^S}{1-q_0^S} \tilde{p}^D = (q_1^A + q_0^A \frac{q_1^S}{1-q_0^S}) \frac{c}{q_2^S}.$$

□

Proof of Lemma 1

To calculate the best expected price obtained in one round of search, conditional on obtaining at least one price quote, we must first calculate the firms' pricing strategy. For any given β_1 and \tilde{p} , an analogous argument to the one used in the proof of Proposition 1 shows that the equilibrium price distribution is

$$F(p) = \frac{(2 - \beta_1)p - \beta_1 \tilde{p}}{2(1 - \beta_1)p}.$$

Denote by \underline{p} the lower bound of the support of this distribution (recall that the upper bound is the reservation price \tilde{p}), by $f(p)$ the density of $F(p)$, and by $j(p) = \frac{d(2F(p) - (F(p))^2)}{dp}$ the density of the minimum of two (independent) price quotes from the distribution $F(p)$. The consumer's expected payment from buying after one round of search that generated at least one quote is thus given by

$$\begin{aligned} & \frac{q_1^S}{1-q_0^S} \int_{\underline{p}}^{\tilde{p}} p f(p) dp + \frac{q_2^S}{1-q_0^S} \int_{\underline{p}}^{\tilde{p}} p j(p) dp = \\ & \tilde{p} \times \left(\frac{\beta_1 \left(2q_2^S \left(\beta_1 \left(1 + \tanh^{-1}(1 - \beta_1) \right) - 1 \right) + (1 - \beta_1) q_1^S \log \left(\frac{\beta_1}{2 - \beta_1} \right) \right)}{2(1 - \beta_1)^2 (q_0^S - 1)} \right). \end{aligned}$$

□

Proof of Proposition 3

If the consumers' value from searching one more time is nonnegative, then the reservation price can be calculated by assuming that a consumer that receives a price quote of \tilde{p} (through ads or previous searches) will search one more time and then buy the good. Hence, Equation (3) simplifies to

$$v - \tilde{p}^{ND} = q_0^S (v - \tilde{p}^{ND}) + (1 - q_0^S) (v - \tilde{p}^{ND} \kappa(q^S, \beta_1^{ND})) - c,$$

and so the reservation price is

$$\tilde{p}^{ND} = \frac{c}{(1 - q_0^S) (1 - \kappa(q^S, \beta_1^{ND}))}.$$

The search market exists if the value of the good minus the sum of the expected search cost and the expected price paid is nonnegative. That is,

$$v - \frac{c}{1 - q_0^S} - \tilde{p}^{\mathcal{N}\mathcal{D}} \kappa(q^S, \beta_1^{\mathcal{N}\mathcal{D}}) \geq 0.$$

Hence, the search market exists if

$$v \geq \frac{c}{(1 - q_0^S)(1 - \kappa(q^S, \beta_1^{\mathcal{N}\mathcal{D}}))}.$$

If the search market exists, then the consumers' expected search costs are $q_0^A \frac{c}{1 - q_0^S}$ and by Lemma 1 their expected payment for the good is $\tilde{p}^{\mathcal{N}\mathcal{D}} \kappa(q^S, q_1^A + q_0^A \frac{q_1^S}{1 - q_0^S})$. Thus, consumer surplus is given by

$$CS^{\mathcal{N}\mathcal{D}} = v - \frac{c}{1 - q_0^S} \left(\frac{\kappa(q^S, \beta_1^{\mathcal{N}\mathcal{D}})}{1 - \kappa(q^S, \beta_1^{\mathcal{N}\mathcal{D}})} + q_0^A \right).$$

Since, in this case, all the consumers buy the good, firms' profits are equal to the expected price paid by consumers, i.e.,

$$\pi^{\mathcal{N}\mathcal{D}} = \tilde{p}^{\mathcal{N}\mathcal{D}} \kappa(q^S, \beta_1^{\mathcal{N}\mathcal{D}}) = \frac{c}{1 - q_0^S} \frac{\kappa(q^S, \beta_1^{\mathcal{N}\mathcal{D}})}{1 - \kappa(q^S, \beta_1^{\mathcal{N}\mathcal{D}})}.$$

□

Proof of Proposition 4

To establish this result we must show that $CS^{\mathcal{D}} > CS^{\mathcal{P}}$. Plugging in the expressions derived above shows that this is equivalent to

$$v - \frac{c(1 - q_2^A)}{q_2^S} > v - \frac{c}{q_2^S} \Leftrightarrow q_2^A > 0.$$

□

Proof of Proposition 5

To determine whether privacy protection is beneficial to consumers, we must sign the expression $CS^{\mathcal{N}\mathcal{D}} - CS^{\mathcal{P}}$. Plugging in the values derived in the previous section shows that this expression is equal to

$$c \left(\frac{q_0^A}{1 - q_0^S} + \frac{1}{1 - q_0^S} \frac{\kappa(q^S, \beta_1^{\mathcal{N}\mathcal{D}})}{1 - \kappa(q^S, \beta_1^{\mathcal{N}\mathcal{D}})} - \frac{1}{q_2^S} \right). \quad (\text{A.2})$$

If the ad market is exactly as competitive as the search market, $\beta_1^{N^D} = \frac{q_1^S}{1-q_0^S}$. Since $\kappa(q^S, \frac{q_1^S}{1-q_0^S}) = \frac{q_1^S}{1-q_0^S}$, in this case, Equation (A.2) simplifies to

$$-\frac{c(1-q_0^A)}{1-q_0^S},$$

and so privacy protection hurts consumers. If the ad market is more competitive than the search market, then $\beta_1^{N^D} < \frac{q_1^S}{1-q_0^S}$. Since by Lemma A.1 $\kappa(q^S, \beta_1)$ is increasing in β_1 , it follows that in such cases privacy protection hurts consumers. This establishes part 2 of the proposition.

Consider an iso-curve of $\beta_1^{N^D}$ in the probability simplex q^A , that is, the set of pairs $\langle q_0^A, q_1^A \rangle$ for which $q_1^A + q_0^A \frac{q_1^S}{1-q_0^S}$ is constant. Along such a curve, expression (A.2) is linear and increasing in q_0^A . Therefore, for every $\beta_1^{N^D}$, there exists $q^*(\beta_1^{N^D})$ such that removing privacy protection increases consumer surplus if and only if²² $q_0^A > q^*(\beta_1^{N^D})$. Moreover, since by Lemma A.1 $\kappa(q^S, \beta_1)$ is increasing in β_1 , it follows immediately that $q^*(\beta_1^{N^D})$ is decreasing in $\beta_1^{N^D}$.

To establish the first part of the proposition, note that the size of the ad market is $1 - q_0^A$, and that, if the advertising market is less competitive than the search market, increasing the size of the ad market decreases $\beta_1^{N^D}$. \square

Proof of Proposition 6

First, consider the case in which firms price discriminate. Consumer privacy protection decreases firms' profits if $\pi^D > \pi^P$. Plugging in these expressions yields

$$(q_1^A + q_0^A \frac{q_1^S}{1-q_0^S}) \frac{c}{q_2^S} > \frac{q_1^S}{1-q_0^S} \frac{c}{q_2^S} \Leftrightarrow \frac{q_1^A}{1-q_0^A} > \frac{q_1^S}{1-q_0^S}.$$

Second, consider the case in which firms do not price discriminate. Consumer privacy protection decreases firms' profits if $\pi^{N^D} > \pi^P$. Plugging in these expressions and simplifying shows that this is equivalent to

$$\frac{\kappa(q^S, \beta_1^{N^D})}{1 - \kappa(q^S, \beta_1^{N^D})} > \frac{q_1^S}{q_2^S}.$$

²²It may be the case that $q^*(\beta_1^{N^D}) < 0$ for some level of $\beta_1^{N^D}$, in which case removing privacy protection is profitable for all advertising technologies that are consistent with $\beta_1^{N^D}$ given q^S . Similarly, if $q^*(\beta_1^{N^D}) > 1$ for some level of $\beta_1^{N^D}$, then removing privacy protection is not profitable for all advertising technologies that are consistent with $\beta_1^{N^D}$ given q^S .

Evaluating $\kappa(q^S, \beta_1^{\mathcal{N}\mathcal{D}})$ at $\beta_1^{\mathcal{N}\mathcal{D}} = \frac{q_1^S}{1-q_0^S}$ yields $\kappa(q^S, \frac{q_1^S}{1-q_0^S}) = \frac{q_1^S}{q_2^S}$, and so for this value of β_1 the firms' profits under regimes \mathcal{P} and $\mathcal{N}\mathcal{D}$ are the same. In Lemma A.1 we showed that κ is increasing in $\beta_1^{\mathcal{N}\mathcal{D}}$, and so consumer privacy protection decreases firms' profits if and only if

$$\frac{q_1^S}{1-q_0^S} < \beta_1^{\mathcal{N}\mathcal{D}} = q_1^A + q_0^A \frac{q_1^S}{1-q_0^S} \Leftrightarrow \frac{q_1^A}{q_1^A + q_2^A} > \frac{q_1^S}{q_1^S + q_2^S}.$$

□

Proof of Corollary 2

If Conditions (4) and (5) hold, then Propositions 2 and 3 jointly imply that total welfare is the same under regimes \mathcal{D} and $\mathcal{N}\mathcal{D}$. The proposition follows from this observation. □

Lemma A.1. $\kappa(q^S, \beta_1)$ is strictly increasing in β_1 .

Proof. Differentiating κ with respect to β_1 gives

$$\frac{\partial \kappa(q^S, \beta_1)}{\partial \beta_1} = \frac{(\beta_1 - 1) \left(-4q_0^S + 2(\beta_1 - 3)q_1^S + (\beta_1 - 2)q_1^S \log \left(\frac{\beta_1}{2-\beta_1} \right) + 4 \right) + 4(\beta_1 - 2)\beta \tanh^{-1}(1 - \beta_1)(q_0^S + q_1^S - 1)}{2(\beta_1 - 2)(\beta_1 - 1)^3(q_0^S - 1)}.$$

Note that this expression is linear in q_1^S , and so to establish that this derivative is positive it suffices to show that it is positive when evaluated at $q_1^S = 0$ and $q_1^S = 1 - q_0^S$.

The derivative of κ evaluated at $q_1^S = 0$ is

$$\frac{-2\beta_1 + 2(\beta_1 - 2)\beta_1 \tanh^{-1}(1 - \beta_1) + 2}{(\beta_1 - 2)(\beta_1 - 1)^3},$$

an expression that is positive for all $\beta_1 \in (0, 1)$.

The derivative of κ evaluated at $q_1^S = 1 - q_0^S$ is

$$\frac{2(\beta - 1) + (\beta - 2) \log \left(\frac{\beta}{2-\beta} \right)}{2(2 - \beta)(\beta - 1)^2}.$$

The denominator of this expression is clearly positive. The numerator evaluated at $\beta = 1$ is zero, and the first derivative thereof is also zero at that point. Moreover, the numerator is a convex function of β_1 . Thus, for any $\beta_1 \in (0, 1)$, Taylor's theorem with the remainder written in the Lagrange form establishes that the numerator is positive. □

B Appendix: Continuous Privacy

In our welfare analysis (and baseline model) we assumed that the level of privacy protection is binary: either privacy is protected or it is not. However, in practice, privacy regulation is often more nuanced and specifies the level of privacy protection in a more continuous manner. In this section, we generalize our analysis to such settings in which the level of privacy protection can be chosen continuously.

We define an ϵ -increase in the level of privacy protection as a change that decreases the mass of consumers who are exposed to interest-based ads by ϵ , without altering the level of competition in the advertising market. Formally, for a given distribution of the number of ads received by each consumer, $\{q_0^A, q_1^A, q_2^A\}$, an ϵ -increase in the level of privacy protection changes that distribution to $\tilde{q}_0^A = q_0^A + \epsilon$, $\tilde{q}_1^A = q_1^A - \frac{q_1^A}{q_1^A + q_2^A} \epsilon$, and $\tilde{q}_2^A = q_2^A - \frac{q_2^A}{q_1^A + q_2^A} \epsilon$. Note that since, in equilibrium, consumers who receive an ad do not search, an ϵ -increase in the level of privacy protection, in essence, increases (decreases) the size of the search (advertising) market by ϵ .

The Marginal Value of Privacy to Consumers

If firms engage in price discrimination, then privacy protection denies consumers access to additional options to purchase the good, without affecting their options in the search market. Intuitively, this suggests that an ϵ -increase in the level of privacy protection is harmful to consumers as it makes it less likely that they will receive this option.

Proposition 4'. *Suppose that Condition (4) holds. If firms price discriminate between the search and advertising markets, then an ϵ -increase in the level of privacy protection hurts consumers.*

Proof. In Proposition 2 we established that when firms engage in price discrimination consumer welfare is $v - \frac{c(1-q_2^A)}{q_2^S}$. Hence, marginally increasing the level of privacy protection reduces consumer welfare by $\frac{cq_2^A}{(1-q_0^A)q_2^S}$. \square

On the other hand, if firms do not price discriminate, privacy protection impacts consumer welfare in two ways: through changes in their expected search costs, and through the change in the competitiveness of the market. As in the baseline model, reducing the level of privacy protection decreases expected search costs, and increases the expected price if and only if the advertising market is less competitive than the search market. Thus, increasing the level of privacy protection is beneficial for consumers if and only if the search market is sufficiently more competitive than the advertising market.

Proposition 5'. Fix the values of v, c , and q^S . Suppose that Conditions (4) and (5) hold, and assume that firms do not price discriminate.

1. If the advertising market is less competitive than the search market, then an ϵ -increase in the level of privacy protection increases consumer welfare if and only if the advertising market is sufficiently large.
2. If the advertising market is (weakly) more competitive than the search market, then an ϵ -increase in the level of privacy protection decreases consumer welfare.

Proof. In Proposition 3 we established that when firms do not price discriminate consumer welfare is $v - \frac{c}{1-q_0^S} \left(\frac{\kappa(q^S, \beta_1^{ND})}{1-\kappa(q^S, \beta_1^{ND})} + q_0^A \right)$. Hence, the marginal value of increasing the level of privacy protection is

$$-\frac{c}{1-q_0^S} \left(1 + \frac{1}{(1-\kappa(q^S, \beta_1^{ND}))^2} \frac{\partial \kappa(q^S, \beta_1^{ND})}{\partial \beta_1^{ND}} \left(\frac{q_2^A}{q_1^A + q_2^A} - \frac{q_2^S}{q_1^S + q_2^S} \right) \right).$$

By Lemma A.1 the derivative of $\kappa(q^S, \beta_1^{ND})$ with respect to β_1^{ND} is positive. Hence, if the advertising market is more competitive than the search market, increasing privacy protection hurts consumers. It can be shown that $\frac{1}{(1-\kappa(q^S, \beta_1^{ND}))^2} \frac{\partial \kappa(q^S, \beta_1^{ND})}{\partial \beta_1^{ND}}$ is increasing in β_1^{ND} , and so, if the advertising market is more competitive than the search market, then the marginal value of protecting privacy is positive if and only if β_1^{ND} is large enough. Since $\beta_1^{ND} = q_1^A + q_0^A \frac{q_1^S}{1-q_0^S}$, this will occur if the advertising market is sufficiently large. \square

The Marginal Value of Privacy to Firms

Firms' profits are determined by the level of competitiveness in the market. In particular, since firms make higher profits in less competitive markets, firms benefit from full privacy protection if and only if the advertising market is more competitive than the search market. Clearly, if the advertising market is less competitive than the search market, an ϵ -increase in the level of privacy protection, which decreases the size of the advertising market while increasing the size of the search market, would also be beneficial to firms.

Proposition 6'. Suppose that Conditions (4) and (5) hold. An ϵ -increase in the level of privacy increases firms' profits if and only if $\frac{q_2^A}{q_1^A + q_2^A} > \frac{q_2^S}{q_1^S + q_2^S}$.

Proof. If firms engage in price discrimination, then by Proposition 2 the firms' aggregate profits are $\left(q_1^A + q_0^A \frac{q_1^S}{1-q_0^S}\right) \frac{c}{q_2^S}$. Hence, marginally increasing the level of privacy protection increases firms' profits by $\frac{c}{q_2^S} \left(\frac{q_2^A}{q_1^A + q_2^A} - \frac{q_2^S}{q_1^S + q_2^S}\right)$.

If firms do not engage in price discrimination, then by Proposition 3 their profits are $\frac{c}{1-q_0^S} \frac{\kappa(q^S, \beta_1^{ND})}{1-\kappa(q^S, \beta_1^{ND})}$. Hence, marginally increasing the level of privacy protection increases firms' profits by $\frac{c}{1-q_0^S} \left(\frac{1}{(1-\kappa(q^S, \beta_1^{ND}))^2} \frac{\partial \kappa_2(q^S, \beta_1^{ND})}{\partial \beta_1^{ND}} \left(\frac{q_2^A}{q_1^A + q_2^A} - \frac{q_2^S}{q_1^S + q_2^S}\right)\right)$. The result follows since $\kappa(q^S, \beta_1^{ND})$ is increasing in β_1^{ND} (Lemma A.1). \square

References

- Acquisti, Alessandro, Curtis Taylor, and Liad Wagman.** 2016. "The Economics of Privacy." *Journal of Economic Literature*, 54(2): 442–492.
- Ali, S Nageeb, Greg Lewis, and Shoshana Vasserman.** 2022. "Voluntary Disclosure and Personalized Pricing." *The Review of Economic Studies*, 90(2): 538–571.
- Armstrong, Mark.** 2006. "Recent Developments in the Economics of Price Discrimination." In *Advances in Economics and Econometrics 3: Theory and Applications*, ed. R. Blundell, W. Newey and T. Persson, 41–97. Cambridge University Press.
- Baye, Michael R., and John Morgan.** 2001. "Information Gatekeepers on the Internet and the Competitiveness of Homogeneous Product Markets." *American Economic Review*, 91(3): 454–474.
- Baye, Michael R., John Morgan, Patrick Scholten, et al.** 2006. "Information, Search, and Price Dispersion." In *Handbook on Economics and Information Systems*. Vol. 1, ed. T. Hendershott, 323–375. Elsevier.
- Bergemann, Dirk, and Alessandro Bonatti.** 2011. "Targeting in Advertising Markets: Implications for Offline versus Online Media." *RAND Journal of Economics*, 42(3): 417–443.
- Bergemann, Dirk, Benjamin Brooks, and Stephen Morris.** 2015. "The Limits of Price Discrimination." *American Economic Review*, 105(3): 921–957.
- Bergemann, Dirk, Benjamin Brooks, and Stephen Morris.** 2021. "Search, Information, and Prices." *Journal of Political Economy*, 129(8): 2275–2319.

- Bird, Daniel, and Zvika Neeman.** 2022. "What Should a Firm Know? Protecting Consumers' Privacy Rents." *American Economic Journal: Microeconomics*, 14(4): 257–295.
- Braghieri, Luca.** 2019. "Targeted Advertising and Price Discrimination in Intermediated Online Markets." *Working Paper*.
- Burdett, Kenneth, and Kenneth L. Judd.** 1983. "Equilibrium Price Dispersion." *Econometrica*, 51(4): 955–969.
- Butters, Gerard R.** 1977. "Equilibrium Distributions of Sales and Advertising Prices." *Review of Economic Studies*, 44: 465–491.
- Calzolari, Giacomo, and Alessandro Pavan.** 2006. "On the Optimality of Privacy in Sequential Contracting." *Journal of Economic Theory*, 130(1): 168–204.
- De Cornière, Alexandre.** 2016. "Search Advertising." *American Economic Journal: Microeconomics*, 8(3): 156–188.
- Diamond, Peter A.** 1971. "A Model of Price Adjustment." *Journal of Economic Theory*, 3(2): 156–168.
- Elliott, Matthew, Andrea Galeotti, Andrew Koh, and Wenhao Li.** 2022. "Market Segmentation through Information." *Working Paper*.
- Fainmesser, Itay P., Andrea Galeotti, and Ruslan Momot.** 2023. "Digital Privacy." *Management Science*, 69(6): 3157–3173.
- Garratt, Rodney J., and Maarten R. C. van Oordt.** 2021. "Privacy as a Public Good: A Case for Electronic Cash." *Journal of Political Economy*, 129(7): 2157–2180.
- Haghpanah, Nima, and Ron Siegel.** 2023. "Pareto-Improving Segmentation of Multi-product Markets." *Journal of Political Economy*, 131(6): 1546–1575.
- Hart, Oliver, Jean Tirole, Dennis W. Carlton, and Oliver E. Williamson.** 1990. "Vertical integration and Market Foreclosure." *Brookings Papers on Economic Activity. Microeconomics*, 21(1990): 205–286.
- Ichihashi, Shota.** 2020. "Online Privacy and Information Disclosure by Consumers." *American Economic Review*, 110(2): 569–595.

- Ipsos, London Economics, and Deloitte.** 2018. "Consumer Market Study on Online Market Segmentation through Personalised Pricing/Offers in the European Union."
- Lewis, Randall A., and Justin M. Rao.** 2015. "The Unfavorable Economics of Measuring the Returns to Advertising." *Quarterly Journal of Economics*, 130(4): 1941–1973.
- Madarasz, Kristof, and Marek Pycia.** 2021. "Towards a Resolution of the Privacy Paradox." *Working Paper*.
- Mauring, Eeva.** 2022. "Search and Price Discrimination Online." *Working Paper*.
- Pram, Kym.** 2021. "Disclosure, Welfare and Adverse Selection." *Journal of Economic Theory*, 197: 105327.
- Preuss, Marcel.** 2023. "Search, Learning, and Tracking." *RAND Journal of Economics*, 54(1): 54–82.
- Rhodes, Andrew, and Jidong Zhou.** 2022. "Personalized Pricing and Competition." *Working Paper*.
- Robert, Jacques, and Dale O. Stahl.** 1993. "Informative Price Advertising in a Sequential Search Model." *Econometrica*, 61(3): 657–686.
- Romanyuk, Gleb, and Alex Smolin.** 2019. "Cream Skimming and Information Design in Matching Markets." *American Economic Journal: Microeconomics*, 11(2): 250–276.
- Zhu, Ruizhi, Sridhar Moorthy, and Xianwen Shi.** 2022. "Advertising Platforms and Privacy." *Working Paper*.