Towards a Resolution of the Privacy Paradox

Kristof Madarasz (LSE) and Marek Pycia (U Zurich)*

Abstract

We consider a seller who sells an item to a buyer over time. The seller can invest in collecting information about the buyer's preferences and the buyer can counter-invest to decrease the precision of the seller's information thus protecting his privacy. We show that the buyer is never willing to pay for such privacy protection. Furthermore, if the buyer can choose between different privacy platforms we show more generally that the buyer always chooses the cheapest privacy platform irrespective of its informational characteristics or that of the available alternatives. Finally, we show that when a privacy platform who contracts with the seller offers some content to users and can offer different privacy platforms from an arbitrary set and price them endogenously, while the buyer can also chose to just shop offline from the seller, in equilibrium the platform prices all privacy platforms equally, the buyer participates on a platform with probability one, and the set of platforms offered minimize the buyer's information rent. The same holds when different platform providers compete.

1 Introduction

George Stigler (1980) opens his piece on the economics of privacy by claiming that "the enormous increase in the interest in privacy in our society is evident in the public press and in the statute books." He continues that "In some

*This draft February 2020. Incomplete. For their comments, we would like to thank Dilip Abreu, Elisabetta Iossa, Stephen Morris, Rakesh Vohra, and the audiences at Spring 2016 Job Talk Lunch Seminar at LSE and at doctoral courses at UCLA and U Zurich. We first drafted this paper while visiting the William S. Dietrich II Economic Theory Center at Princeton University, and we would like to thank the Center for its hospitality; in addition, Pycia would like to than the Center for financial support. Keywords: Privacy Protection, Tracking, Defaults, Dynamic Pricing.

respects this interest in privacy is paradoxical, for the average citizen has more privacy — more areas of his life in which his behavior is not known by his fellows — than ever before." While the former claim still holds, the latter is unlikely to be true. In fact, currently there is an extensive discussion of a different paradox. The so-called privacy paradox refers to the fact that research on online behavior has revealed *discrepancies* between user attitude and their actual behavior: While users claim to be very concerned about their privacy, they nevertheless undertake very little to protect their personal data (see e.g., Barth and deJong 2017 or Kokolakis 2017 for reviews).

Indeed, the presence of unverifiable private information remains central to the interaction between firms and consumers. With the rise of the internet and traceable online search, however, the ease of recording and processing data about people has increased substantially. Many business models in the economy are based on collecting, storing, and processing such data about observable individual behavior. While buyers preferences are the buyers' soft information, such activities, in principle, allow firms to uncover information about otherwise unobservable and unverifiable personal preferences from patterns of behavior. It is then no surprise that firms value such information and engage in costly effort to obtain it. Information about preferences decreases the information rent that a potential seller needs to offer to a buyer when selling, which, in turn, it increases the seller's expected profit and in many contexts such rise in profits is likely to also lead to a decrease consumer surplus.[1]

People are generally aware of such activities. Furthermore, they appear concerned about firms' investments in technologies that serve the purpose to obtain, store, and interpret information about them. For example, Turow et al. (2009) report, based on a representative survey, that "contrary to what many marketers claim, most adult Americans (66%) do not want marketers to tailor advertisements to their interests. Moreover, when Americans are informed of three common ways that marketers gather data about people in order to tailor ads, even higher percentages - between 73% and 86% - say they would not want such advertising." They also report that 63% believe advertisers should be required by law to immediately delete information about their internet activity.

Despite being concerned about the loss of their privacy, however, con-

[1] For example, Mikians et al (2012) suggest that considerable price differences, 10-30%, exist for identical products based on a variety of online characteristics, such as location or browser configuration.

sumers appear to do little to protect the privacy of their behavior , e.g., conceal their behavioral patterns, pay extra for apps or products that do not track or verify them, request access to various aspects of their data, e.g., Rose (2005), Acquisti and Grossklags (2005). For example, Rose (2005) finds that although most survey respondents reported that they were concerned about their privacy, only 47 percent of them expressed a willingness to pay any amount to ensure the privacy of their information. Similarly, although the evidence is at times harder to interpret, others, e.g., Tsai et al. (2011) and Beresford et al. (2012) also provide findings consistent with the general consensus that people appear to be willing to pay very little extra to use platforms that provide a greater level of privacy.

To explain this apparent paradox, various authors have proposed that lack of information, or some form of bounded rationality or psychological factors such as a preference for immediate gratification, loss aversion or miscalibration of probabilities to be at play, e.g., Acquisti and Grossklags (2005), Acquisiti, John, and Loewenstein. (2013). While such behavioral factors are likely to be important, attempts to theoretically explain and practically solve this puzzle are still scarce and the subject may deserves more research attention. Indeed, the above cited surveys also suggest that further research in this area is likely to be key.

In this paper we emphasize another channel that may imply that a buyer's willingness to pay extra for platforms that limit the collection of their data, thus allowing the seller to learn about their private and unverifiable preferences, may be small. In particular, we describe a simple, but robust reason that might help explain this phenomenon given the underlying strategic interaction.

We first consider the classic problem where an uninformed seller faces a private-informed buyer over time. The seller can invest in increasing the precision of his/her information about the buyer over time. The buyer can also invest in decreasing the precision of the seller's information. Despite the fact that buyer has substantial information rents to collect and suffers a great loss of such rents when her privacy is lost, we show that the buyer is never willing to invest in a greater protection of his privacy and, as long as such protection is costly, does not limit the seller's learning process.

More generally, we then describe an informational irrelevance result. Suppose that a consumer, who is privately informed about her preferences, can choose between different experiments, privacy platforms. Each privacy platform is associated with some signal structure that is transmitted to the seller. The information that is revealed by the buyer's type may depend on the buyer's type in an arbitrary way. We allow for an arbitrary finite set of experiments. Suppose then that each experiment has some cost associated with it that the buyer has to incur. For example, protecting privacy may be more costly, then not protecting it at all. Such costs may also depend on the buyer's type, but there is a common default or least costly experiment to all types. We show that in this case all buyer types pool and choose this least costly experiment irrespective of its information content. In other words, it does not matter whether this platform offers full privacy protection or leads to full revelation of the buyer's preferences, or provides information about some types but not about a set of other types, all buyer types choose the same platform. The buyer always sticks to the default privacy platform irrespective of its informational characteristics or that of the available alternatives.

We then endogenize the choice and the pricing of privacy platforms by a platform provider. Suppose first that the platform provider is a monopolist who can choose from an arbitrary set of technologically feasible platforms. The monopolist can decide which platforms to offer and at what price each. The platform provider also contracts with the seller and agrees on some profit-sharing, that is, on a contract where in exchange for the information provided to the seller, the seller's transfer back to the platform provider is increasing in the boost in the seller's profit given the information provided. Finally, a buyer then decides which platform to choose, may decide to shop 'offline directly from the seller, or exit the market. In this context we show that in equilibrium the platform provider offers platforms that maximizes the value of the information passed to the seller.

Finally, maintaining the same setting we also allow for competition between platform providers. In equilibrium, again, the platform providers offer platforms that maximize the value of the information passed to the seller. The results continue to hold when allowing for competition between sellers.

Related Literature [To be Completed] Our setup relates to various strands in the literature. It relates to classic mechanism design in that a privately informed party with non-verifiable information chooses between different alternatives and that choice may help the uninformed party learn her type. In our setting these alternatives are different information revelation technologies offered at potentially different – and potentially type-dependent — benefits / costs. Here, we show a general pooling result. It also relates to the literature on information design, e.g., Kamenica and Gentzkow (2011), Bergemann and Morris (2019), where the sender commits to an experiment to persuade a receiver. We relate to this literature in that the platform designer offers a set of platforms experiments. We also relate to the voluntary disclosure literature of an informed party, e.g., Grossman and Hart (1980) and Milgrom (1981) who show an unraveling result that under certain assumptions (including: common knowledge that the agent is privately informed, disclosing is costless, and information is verifiable) all types disclose their private information. Among others, Jovanovich (1982) considers costly disclosure and shows that if disclosure is costly, then in some settings high quality sellers disclose and low quality sellers do not. Instead, in our setting, the informed party has no verifiable information and instead can invest in others not learning her private information. We thus show that an unravelling logic resembling Grossman and Hart (1980) is much more general then in standard treatments.

2 The Value of Protecting Privacy

Consider the following classic monopoly problem. The seller owns an object that he produces at a normalized cost of zero. The buyer's privately knows her valuation which is denoted θ and assumed to be positive; its cdf $F(\theta)$ is strictly increasing on some positive support $[\underline{\theta}, \overline{\theta}]$. In each period $t = 1....$ the seller makes a price offer that the buyer then accepts or rejects. If the buyer accepts, she obtains her valuation net the price and the seller receives the price and the game ends. If the buyer rejects, then the game continues. The seller and the buyer discount the future at same fixed rate. We employ the standard notion of perfect BNE (equilibrium or PBE henceforth) as our solution concept.

In this classic setup Myerson (1983) shows that, for the seller who can commit to a strategy, the optimal selling mechanism is for the seller to maintain the same posted price in all periods. Such an optimal mechanism then also characterizes the information that the seller can optimally elicit from the privately-informed buyer in an incentive-compatible manner. The seller's profit is constrained by the information rent that the seller needs to offer to the buyer in terms of consumer surplus to be able to sell.

We study the privacy of buyer's information in this classic problem. We assume that the type θ, e.g., describing the buyer's preferences, cannot be credibly revealed to the seller (it is not a hard information), but there are some signal processes that over time provide information about θ; we model these processes below. We are interested in the question whether the buyer would choose to influence these signal processes in a way that would protect the buyer's privacy. We hence assume that at the interim stage, that is, before the seller makes any price offer, but once the buyer has private information, the buyer can invest into information (privacy) protection. As our focus is on the buyer's privacy protecting decisions, we simplify the seller's information collecting decisions and assume they will be costless to the seller or to the intermediaries (platform providers) we study in subsequent sections. Let us note however, that unlike the buyer, the seller (and the platform providers) would be willing to pay for information collection. Finally, we make no assumption as to whether the seller can commit to a price sequence or if his pricing decisions need to satisfy sequential rationality as our results hold in both cases.

2.1 Example

As an example, suppose that at the beginning of each period t there is an independent leakage probability α such that the buyer's valuation leaks to the seller, that is, the seller privately learns θ with this probability. For example, the seller, or the operator of an online platform on her behalf, may be able to figure out the buyer's preferences for the seller's product from observing the buyer's activity online. We can think of this leakage probability as one publicly set by the seller at some cost initially. Since the result does not depend on the exact values of α we do not model this stage explicitly.

Suppose that, having learned her type, but before the seller can observe anything or make any offers, the buyer can invest $m \in 0 \cup [a, b]$ to protect his privacy, where $b > a > 0$. Protecting his privacy means that this investment decreases the arrival probability; i.e., $\alpha(m)$ is decreasing in m. Such protection is costly and the cost function $c(m)$ is such that $c' > 0$ on $[a, b]$ and there exists $\varepsilon > 0$ such that $c(a) > \varepsilon$. We interpret $m = 0$ as the no investment decision, leaving the leakage probability unaffected, and normalize $c(0) = 0$. In other words, we assume that there is an (arbitrarily small but positive) cost to pay for protecting privacy. We also assume that the event $m > 0$ (and hence $m = 0$) is observable by the seller.

Proposition 1 *The buyer never invests in protecting her privacy.*

This and all subsequent proofs are provided in the Appendix. The logic of the proof is as follows. Note that a pooling equilibrium with no investment, that is when no buyer type protects his privacy, exists. This is true because given such expectations if any buyer type were to invest, the seller can always attribute the investment coming from the type with the highest willingness to pay. In turn, the seller would not drop the price below that level but only once leakage occurs which, however, still leaves the deviating type with no surplus. In contrast, a separating equilibrium cannot exist because if it did, then one can always consider the lowest type, or the infimum of types in case the set of types who invest is not closed, who decides to invest some positive amount. Such a type, however, can never gain from investment since it is also the infimum of the seller's posterior conditional on the buyer's equilibrium investment choice. This type would then benefit by deviating and not investing and thereby not protecting his privacy such benefits include a lower cost of investment and a (weakly) lower sales price.

2.2 General information revelation processes

The above proposition does not hinge on the details of the example and remain valid for general partial dynamic information revelation. Specifically, suppose that in each period the seller obtains a signal $s_t : [\underline{\theta}, \overline{\theta}, \{\widehat{s}_k\}_{k=1}^{t-1}] \to \Delta Z$ with probability α_t where Z is a finite realization space. We also allow the signal realization in period t to be a function of the past signal realizations, $\{\widehat{s}_k\}_{k=1}^{t-1}$, in case signals in the past have arrived. We assume now that the signal is no longer fully revealing and assume that each realization of s_t occurs with positive probability given each element of the domain, that is, observing a signal s_t per se, leaves the seller with a posterior that has full support. Of course in equilibrium, the seller's posterior in period t may rule out a positive measure of types, but such a feature of the seller's beliefs is a consequence of equilibrium behavior directly. Investment is costly just as before, but we can also let $\alpha_t(m)$ be an arbitrary function which is strictly positive on $[a, b]$.

Proposition 2 *The buyer never invests in protecting her privacy.*

The logic of the above result is the same as the one described before.

2.3 Privacy Platforms

Building on the logic of the above, suppose that the buyer can choose between N different privacy platforms (approaches, experiments) $\{s_1, ...s_N\}$. Each privacy platform, $s_j : \Theta \to \Delta(Z_1 \times Z_2 \times ....)$ is a function which assigns to each type a probability distribution over a set of, possibly correlated, signal realizations over time. We again assume that the each probability distribution has full support. Let $c(s, \theta)$ be the cost of *choosing* platform s for type θ. The privacy platform chosen is observed by the seller. We assume that there is a platform that is the cheapest for all types, that is, there exist $s^* \in S$ such that $c(s^*, \theta) < c(s', \theta)$ for all θ and s'. We can interpret this platform as the default as without loss of generality we can assume that if the buyer does nothing then he or she uses this platform (notice that this is just a normalization calling the choice of the cheapest platform as doing nothing). The motivation for this terminology is that defaults are typically less costly to choose than other options.

Proposition 3 *The buyer always chooses the default privacy platform independent of its information content of the other element of the choice set.*

Remark 1 *true—with no change in the proof—if we allow a richer set of privacy platforms in which having selected privacy platform s_j the agent can further choose a message (or an action) from a set M_j which is costless and not observed by the seller. Thus the privacy platform can be defined as a mapping $s_j : \Theta \times M \to \Delta(Z_1 \times Z_2 \times ....)$ which assigns to each type and message pair a probability distribution over a set of, possibly correlated, signal realizations over time. We can interpret these privacy platforms as contracts which tell the buyer that signals z_1, z_2 etc. will be communicated to the seller at respective times 1, 2, etc.*

3 Design of privacy platforms

So far we took the existence of privacy platforms to be exogenous. What happens when they are designed by a privacy platform provider, such as a search website or other online intermediary? Our main insight—on the buyers inaction to protect their information in equilibrium—carries over to such a richer environment. In particular, we show that the observability of buyer's choice of privacy platform emerges endogenously in equilibrium.

We start by considering a monopolistic designer of the privacy platforms and we address competing platform providers in the next section. The platform provider can offer privacy platforms from an arbitrary set P of technologically feasible privacy platforms. Each platform consists of both the signal generating process and a contract with the buyer that specifies what type of information might be passed by the provider to the seller of the good the buyer may want to buy, cf. Section 2.3 above. We assume that the set of feasible platforms is finite; but this assumption can be relaxed as long as the topology of the problem ensures that the provider has well-defined optimal choice. We also assume that all platforms are equally costly to the provider; this assumption can be fully relaxed.[2]

After selecting any subset, the designer decides at what price to offer each to the buyer. In other words, we now allow the cost of each platform to be determined endogenously in equilibrium by a profit-maximizing platform provider. Finally, the platform providers communicate to the seller the buyer's choice and all other information about the buyer that is allowed by the contractual agreement with the consumer. After offering a subset of the technologically available platforms and corresponding prices, the buyer picks whichever she prefers.

We allow the buyer to choose to shop 'offline' directly from the profit-maximizing seller, i.e., not use the platform. This option then allows for a, in principle, type-dependent outside option. We may also assume that the buyer can also decide to choose neither a platform from the provider, nor purchase offline, i.e., to exit the market. We furthermore assume that being on the platform offers some, potentially arbitrary small, pure benefit $b > 0$ which we can interpret as e.g., search or entertainment value of the platform.

The model so far did not specify the bargaining between the platform provider and the seller. To close the model let us assume that the platform provider makes a take-it-or-leave-it offer to the seller and then the seller responds, both events happening before the provider makes the offer to the buyer. The timing details turn out not to matter as long as the bargaining between the platform provider and the seller occurs while both parties are still uninformed. In particular, we could also assume that the provider makes a take-it-or-leave offer to the seller simultaneously with the offer to the buyer. The above bargaining procedure between the platform provider and the seller is also not crucial: the no-privacy protection insight goes through for all bargaining procedures in which the expected transfer from the seller to the platform is strictly monotonic in the seller's expected revenue, e.g. if the platform captures a fixed positive share of the seller's revenue. Below we assume without loss of generality that if a platform is never chosen by the buyer in equilibrium it is also not offered by the provider.

Proposition 4 *In equilibrium all platforms offered have a price of b, the set*

[2] After we relax this assumption, we need to adjust the proposition below so that the platform provider chooses the platform that maximizes the provider's share of the seller profit net of provider's cost of operating the platform.

of platforms provided maximizes the seller's revenue among all feasible platform sets, and with probability 1 *the buyer participates in one of the platforms provided.*

The above implies that when the full information extraction (no privacy) platform is feasible for the platform provider, then the following is true:

Corollary 1 *Suppose that the platform provider has a full revelation platform available. Then that is the platform that will be provided at the price of* b*. In turn, the seller and the platform provider will extract all the surplus from the market and the equilibrium is essentially unique.*

4 Competition between platform providers

While we derived our insight—that the cheapest privacy platform is chosen—assuming a monopolistic seller and an exogenous set of privacy platform or a monopolistic designer of the platforms, the insight remains valid when two or more sellers compete for the buyer's custom as well as when two or more privacy-platform designers (tracking institutions) offer their services to buyers.

4.1 Competition amongst Platform Providers

Consider now the setting of Section 3 except that there are multiple platform providers competing for the buyers. As with the monopoly platform provider, the details of the bargaining between the platforms and the seller do not matter as long as the expected transfer from the seller to the platform is strictly monotonic in the seller's expected revenue, e.g. if the platforms capture a fixed positive share of the seller revenue from the buyers matched with the seller by the platform. We assume that each platform provider has access to the same set of feasible platforms but otherwise allow asymmetric platform providers; thus different platforms might have different ability to negotiate with the seller.[3] The providers simultaneously offer menus of platforms to the buyers.[4] As before we allow transfers between the buyer and the platform providers; we refer to negative prices as subsidies.[5] The buyer sees the

[3] Allowing different platform providers access to different sets of feasible platforms does not change things substantially but its full analysis would require taking a less agnostic stance on the bargaining between the seller and platform providers as we then need to analyze the seller's preference to contract with the provider who can offer the most information-extracting platform.

[4] This assumption is to focus attention only. In particular, our results hold if the providers move sequentially.

[5] Below we separately consider the case when subsidies are not allowed.

menus of each provider before choosing one of the platforms or choosing to shop offline or exit the market. Finally, the platform providers communicate to the seller the buyer's choice and all other information about the buyer that is allowed by the contractual agreement with the consumer.

Proposition 5 *In equilibrium, the platforms that maximize the seller's revenue are offered and, with probability 1, the buyer chooses to join one of these platforms.*

For instance, if the competing platform providers have access to the same set of feasible platforms and platform i captures the share λ_i of the seller's revenue, then each provider with highest λ_i chooses the platforms that maximize the seller's revenue and provides the subsidy equal to the product of the second highest λ_i times the expected revenue of the seller. Furthermore, with probability 1, the buyer chooses to join one of the platforms offered by providers with highest λ_i.

We thus see that the competition shifts some of the surplus towards the buyers while not affecting the insight that the privacy is not protected.

4.1.1 The role of subsidies

So far we allowed the platforms to subsidize the buyers. Suppose now that such subsidies (beyond the benefit of using platforms b) are not possible. We may want to impose this assumption e.g. if we are concerned that in practice subsidies may lead the buyers to collect the transfers—possibly multiple times, by posing as multiple buyers (shills)—without buying the product. No subsidy is then a simplifying assumption made in lieu of specifying a richer model.

In the presence of the no-subsidy assumption, multiple essentially different equilibria of the privacy game become possible. In each equilibrium, all platform providers offer their platforms for free, and, with probability 1, the buyer chooses to join one of the platforms. There is an equilibrium in which the privacy is not protected and each platform provider chooses the platforms that maximize the seller's revenue but there are now also other equilibria, including an equilibrium in which the most privacy-protecting platforms are offered by providers and chosen by the buyers. The no-privacy equilibrium maximizes the joint surplus of the platform providers and as such a focal one. We formulate this insight for the case of symmetric platform providers.

Proposition 6 *Suppose that subsidies are not allowed. In the equilibrium that maximizes the joint surplus of the platform providers, each provider* *chooses the platforms that maximize the seller's revenue and offers it at price* 0. *Furthermore, with probability 1, the buyer chooses to join one of the platforms offered.*

4.2 Competition among sellers

The insights of Sections 2, 3, and 4.1 remain valid when there is more than one seller as long as the sellers have some profits from their interaction with the buyer (gross of the payment to platform provider). If the competition among the sellers is such that all surplus accrues to the buyer then the no-privacy outcome remains a Perfect Bayesian Equilibrium, but there are other equilibria as well.

5 Regulation

[[**Incomplete**]] The results imply a general divergence between privacy attitudes and privacy choice of consumers and the idea that buyers will pool on the cheapest privacy platform irrespective of its informational content. This insight is relevant in light of the great policy interest in regulating the ways firms need to ask for the consent of those whose information they collect, store and sell. The results imply that if firms simply are required to give the option to consumers to protect their privacy may be completely ineffective. Instead direct regulation of privacy may be a more effective way to safeguarding. This might help inform such regulations as GDPR. By shifting the (cheapest for the buyer) default from opt out to opt in, the regulation can shift the equilibrium from no privacy to the regulations prescribed privacy provided the platforms offer no subsidies. With subsidies, and assuming that the sellers are willing to subsidy opt-in sufficiently to cover the costs of an active choice, the old equilibrium (opt in) prevails with the buyers merely compensated for the active opt-in choice. At the same time our logic suggests that regulation that takes away the property rights of platform of the data collected, e.g., mandatory sharing of the data, may work since it breaks the link between the platforms' and the seller's payoffs.

6 Conclusion

Privacy is a major public concern. In the context of classic monopoly selling we show that despite that enjoying privacy may be greatly valuable to a consumer nevertheless the consumer is never willing to pay to protect her privacy. The result also suggest that regulating default privacy settings can have a very powerful effect provided the subsidies that can be offered by the privacy platform providers are also controlled.

References

[1] Acquisiti, A. and J. Grossklag (2005). “Privacy and Rationality.” In: *Individual Decision Making IEEE. Sec Priv* 3: 24-30.

[2] Acquisiti, A. L. K. John, and G. Loewenstein (2013). What is Privacy Worth. *Journal of Legal Studies.*

[3] Acquisti, A C. Taylor and L. Wagman (2016). The Economics of Privacy, Journal of Economics Literature.

[4] Ali, N., G. Lewis, and S. Vasserman (2019). “Voluntary Disclosure and Personalized Pricing.” Working Paper.

[5] Barth and DeJong (2017). “The privacy paradox – Investigating discrepancies between expressed privacy concerns and actual online behavior – A systematic literature review.” *Telematics and Informatics*, 1038-1058.

[6] Beresford, A., D. Kubler, and S. Preibusch (2012). “Unwillingness to Pay for Privacy: A Field Experiment.” *Economics Letters* 117:25-27.

[7] Bergemann, D,, and S. Morris (2019). “Information Design: A Unified Perspective.” *Journal of Economic Literature* 57(1): 44–95

[8] Board, S and J. Lu, 2018 “Competitive Information Disclosure in Search Markets.” *Journal of Political Economy* 126 (5), 1965-2010.

[9] Grossman, S., and O. Hart. 1980. ”Disclosure Laws and Takeover Bids.” Journal of Finance 35 (2): 323-34.

[10] Gutman, I. Kremer, and A. Skrzypacz (2014) “Not Only What but Also When: A Theory of Dynamic Voluntary Disclosure,” *American Economic Review*

[11] Hermalin, B., Katz, M., (2006). Privacy, Property Rights and Efficiency: The Economics of Privacy as Secrecy. Quantitative Marketing and Economics 4(3), pp. 209–239.

[12] Kamenica, E., and M. Gentzkow (2011). “Bayesian Persuasion.” *American Economic Review.* 101(6), 259–2615.

[13] Kokolakis, S. (2017). “Privacy attitudes and privacy behaviour: A review of current research on the privacy paradox phenomenon.” *Computers & Security.* Pp. 122-134.

[14] Milgrom, Paul (1981) “Good News and Bad News: Representation Theorems and Applications.” *Bell Journal of Economics* 12, 380–391.

[15] Myerson, Roger B. (1981). ”Optimal Auction Design”. *Mathematics of*

Operations Research 6 (1): 58–73.

[16] Stigler, G. J. (1980). "An Introduction to Privacy in Economics and Politics." *Journal of Legal Studies* 9, pp. 623–644

[17] Taylor C. (1999). "Time-on-the-Market as a Sign of Quality." *Review of Economic Studies* 66, 555-78.

[18] Turow, J., King, J., Hoofnatle, C.J., et al. (2009). "Americans Reject Tailored Advertising and Three Activities That Enable It." Departmental Papers (ASC), 137.

[19] Rose, Ellen A. (2005). "Data Users versus Data Subjects: Are Consumers Willing to Pay for Property Rights to Personal Information?" *Proceedings of the 38th Annual Hawaii International Conference on System Sciences*. Washington, D.C.: IEEE Computer Society Press. doi:10.1109/HICSS.2005.184.

[20] Tsai, Janice Y., S. Egelman, L. Cranor, and Alessandro Acquisti. 2011. "The Effect of Online Privacy Information on Purchasing Behavior: An Experimental Study." *Information Systems Research* 22:254–68.

7 Appendix

Proof. [Proof of Proposition 2] First note that an equilibrium in which all types pool on $m = 0$ exists. This is true because if any type would deviate to some choice m' the seller can attribute this to coming from $\overline{\theta}$ given the full-support assumption on each s_t. Since given such beliefs it is an equilibrium off-path strategy to then keep charging $p = \overline{\theta}$, no profitable deviation exists. It remains to prove that no different equilibrium exists. Because type $\underline{\theta}$ never obtains a surplus from trading in equilibrium, this type will choose investment $m = 0$. This argument rules out any pooling equilibrium on some investment choice $m' \neq 0$. Consider now the infimum θ of types that chooses some investment $m' \neq 0$. By also choosing m' type θ would not obtain a surplus from trading in equilibrium and would prefer to deviate to investment $m = 0$ because $c(m') > 0$. As θ is the infimum of types choosing m' for any $\epsilon > 0$ there must be type $\theta_\epsilon \in (\theta, \theta + \epsilon)$ that chooses m'. Let $\epsilon \in (0, c\,(a))$. Then type θ_ϵ profit from trading is smaller than ϵ as the price for buyers choosing m' is at best θ and this type would benefit by deviating to investment $m = 0$. The presence of this deviation rules out any semi-separating equilibrium. ■

Proof. [Proof of Proposition 3] As before, $\underline{\theta}$ must choose $x^* = \arg\min c(x, \underline{\theta})$. Suppose that some types choose an experiment $x' \neq x^*$. Consider again the infimum of the types choosing x^*. This type, or a type arbitrarily close to it, receives essential no surplus from trading. Since $c(x', \widehat{\theta}) > c(x^*, \widehat{\theta})$, the result follows from the previous argument. ■

Proof. [Proof of Proposition 4] Note first that irrespective of the provider's choice in equilibrium no type will buy offline. To show this, suppose that a set of types chose to shop offline. Take the infimum of these types. In equilibrium, such a type can have no surplus from buying offline given a profit maximizing seller. At the same time, this type loses benefit $b > 0$ net of the price of the platform. As long as the price is strictly below b, this type has a strict incentive to deviate, a contradiction. Similarly were the price strictly higher than b, no types would go through the platform, and thus the platform provider charges at most b. Finally, if the platform provider charged b and a buyer would go offline with positive probability then the platform could deviate by charging slightly less than b.

Suppose now that the platform provider offered two different platforms at different prices. By our previous argument, all types choose the cheaper platform. Hence, in equilibrium all platforms will be offered at price b. Furthermore, since the price is the same for all equilibrium platforms and the expected transfer the platform provider obtains from the seller is strictly increasing in the seller's expected revenue, in equilibrium the platform provider will maximize the seller's revenue. ■

Proof. [Proof of Proposition 5] By the same argument as in earlier environments studied, all buyer types buy from one of the cheapest platforms irrespective of their informational content. Thus, the platform providers are engaged in the Bertrand competition on subsidies and each provider whose platform is chosen with positive probability (also referred to as a winning provider) offers to the buyer the maximum of the subsidies that other providers are able to offer. By our assumption on the bargaining between platforms and the seller, the platforms that maximize the seller's revenue also maximize the platform provider's profits and their ability to subsidize the buyer; thus these platforms are chosen by the winning providers. ■