Privacy Trading in the Surveillance Capitalism Age
Viewpoints on 'Privacy-Preserving' Societal Value Creation

Ranjan Pal, Jon Crowcroft
University of Cambridge
rp631@cam.ac.uk,ranjanpal9@gmail.com,jac22@cam.ac.uk

This article is an editorial note submitted to CCR. It has NOT been peer reviewed.
The authors take full responsibility for this article's technical content. Comments can be posted through CCR Online.

ABSTRACT

In the modern era of the mobile apps (part of the era of *surveillance capitalism*, a famously coined term by Shoshana Zuboff), huge quantities of data about individuals and their activities offer a wave of opportunities for economic and societal value creation. However, the current personal data ecosystem is mostly de-regulated, fragmented, and inefficient. On one hand, end-users are often not able to control access (either technologically, by policy, or psychologically) to their personal data which results in issues related to privacy, personal data ownership, transparency, and value distribution. On the other hand, this puts the burden of managing and protecting user data on profit-driven apps and ad-driven entities (e.g., an ad-network) at a cost of trust and regulatory accountability. Data holders (e.g., apps) may hence take commercial advantage of the individuals' inability to fully anticipate the potential uses of their private information, with detrimental effects for social welfare. As steps to improve social welfare, we comment on the the existence and design of efficient consumer-data releasing ecosystems aimed at achieving a maximum social welfare state amongst competing data holders. In view of (a) the behavioral assumption that humans are 'compromising' beings, (b) privacy not being a well-boundaried good, and (c) the practical inevitability of inappropriate data leakage by data holders upstream in the supply-chain, we showcase the idea of a regulated and radical privacy trading mechanism that preserves the *heterogeneous* privacy preservation constraints (at an aggregate consumer, i.e., app, level) upto certain compromise levels, and at the same time satisfying commercial requirements of agencies (e.g., advertising organizations) that collect and trade client data for the purpose of behavioral advertising. More specifically, our idea merges *supply function economics*, introduced by Klemperer and Meyer, with differential privacy, that, together with their powerful theoretical properties, leads to a stable and efficient, i.e., a maximum social welfare, state, and that too in an algorithmically scalable manner. As part of future research, we also discuss interesting additional techno-economic challenges related to realizing effective privacy trading ecosystems.

CCS CONCEPTS

• **Security and Privacy**; • **Economics**;

KEYWORDS

Privacy, Trading, Mobile Apps, Market, Ecosystem, Efficiency

1 INTRODUCTION

Mobile technology is a major driving force behind the modern digital society including business small and large, personal lifestyles, as well as the state-of-the-art IoT/CPS systems. All this is made possible through mobile apps (applications) that enable the functioning of operations in this ecosystem. ln-app advertising is an essential part of this digital ecosystem of free mobile applications, where the ecosystem entities comprise the consumers, consumer apps, ad-networks, advertisers, and retailers (see Figure 1 for a simplified representation for the ad-network and advertisers case). In reality, advertisers and retailers could be directly linked to the consumer apps in sell-buy relationships. As a popular example, *Evite.com* may sell lists of their consumers attending a party in a given location directly to advertisers. As another example, the gene testing company *23andMe* might sell their clientele information directly to pharmaceutical companies in order for the latter to develop medical drugs. As a social objective, a *win-win* situation among the entities of this ecosystem is desired, where (a) the privacy preferences of consumers can be preserved, (b) the free apps can display appropriate targeted consumer ads [21], (c) the ad-networks can publish the right ads on their billboards, and (d) the advertisers can target as large as possible and more importantly the right set of consumers. The basic requirement for this 'win-win' ecosystem to exist in the first place, is the flow of personalized information from the consumer to the advertisers via the ad-networks for effective/profitable ad placements, that subsequently motivate the latter to collect personal data about consumers via apps. The vision and benefits for such an ecosystem were laid down by a certain school of information economists [1][16], in favor of having increased aggregate societal welfare. More specifically, according to the authors in [1], in return for personal data, advertisers and marketers will benefit the individuals through monetary compensation (e.g., discounts, Facebook Libre Coins) and intangible benefits (e.g., personalization and customization of information content), and price discrimination. Furthermore, the authors state that the lack of use of personal data will lead to opportunity costs and market inefficiencies. To this end, three important questions that draw our attention are: *can such an ecosystem exist in the digital society?; if yes, then how should it be designed?; and what is the implication of such a design to policy, economic science, and technology?*

Organization of the Paper - In this note, we first provide arguments for the virtually inevitable need of such ecosystems, despite positive-minded barriers put in place to prevent commercial use of personal data in practice (see Section 2). We then provide an overview of supply function economics as an appropriate tool, that when combined with differential privacy, can be used to design an ecosystem resulting in stable and efficient economic outcomes from privacy trading (see Section 3). Finally, we also discuss interesting additional techno-economic challenges related to realizing effective privacy trading ecosystems in practice.

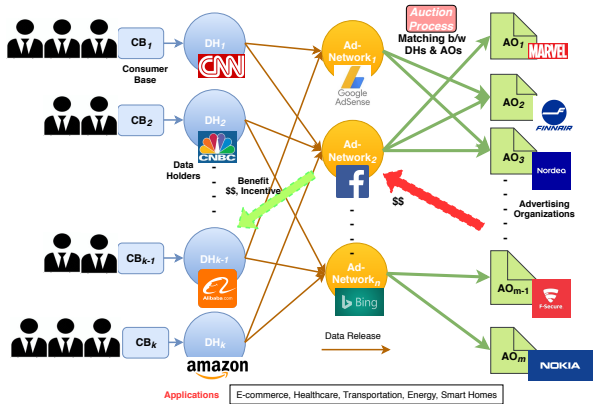

Figure 1: Illustration of Mobile In-App Ad Ecosystem

2 LIKELY INEVITABILITY OF DATA RELEASE

In this section, we state the likely inevitable need of data release ecosystems citing multiple reasons, despite positive-minded barriers put in place to prevent commercial use of personal data in practice. This subsequently calls for putting privacy trading methodologies in place as one possible direction to ensure improved economic welfare. As section organization, we first discuss the inevitability of data release frameworks in the presence of positive-minded barriers to commercial use of personal data, thereby paving the way for trading of data (privacy). We then brief the reader about the two-decade old but a newly-coined term, 'surveillance capitalism', and how such a capitalism era we are in bolsters the inevitability of data release. Further, we take a quick look into privacy issues that might arise in the presence of privacy trading itself. Finally, we comment on a mechanism that mitigates privacy issues from trading personal data.

Barriers to Commercialization of Personal Data - As a regulatory corrective step to prevent commercialization of personal data, data protection laws, a high-profile example being the General Data Protection Regulation (GDPR) [24], impose constraints, rights and obligations regarding personal data and its use. However, it is questionable as to whether the psychological approach of many apps—in offering a binary opt in/out, often after presenting pages of legalese—results in user empowerment with respect to making the proper choice between gaining utility from an app versus not using it. Indeed, we see that individuals are increasingly using ad-blocking technology[1] as a means to 'push-back', alongside deciding to gain utility from apps. However, ad blocking firms like *Eyeo*, maker of the popular *AdBlock Plus* product, has achieved such a position of leverage that it gets Google et.al., to pay it to have their ads whitelisted by default - under its self-styled 'acceptable ads' program [17] - clearly going against the core functionality principle of ad-blockers. In addition, any attempt of territorial governments to enforce privacy regulations could increase the likelihood of data-driven companies (whose profits depend significantly on data) to employ legal arbitrage, and in extreme cases drive out firms benefiting from the data economic chain, reducing tax revenues of countries. Finally, there is further evidence (courtesy Federal Trade Commission Reports) that despite the high transaction costs and risks in holding personal data (e.g., via penalty mechanisms enforced by GDPR), re-identification by data brokers of data through the connection of disparate datasets, finds an efficient market. This, despite a plethora of popular anonymization and aggregation techniques aimed to prevent personal data breach attempts. *Thereby, it is fair to assume that, with a high likelihood, there would be an inevitable breach of personal consumer information in general to satisfy the economics behind the working of the current ad ecosystem*, thereby leading to failing of the privacy-preservation property[2]. Some recent studies [17] have stacked up cases against behavioral advertising post GDPR, stating that advertising firms can make more revenues from traditional advertising channels such as TV and newspapers, compared to online/mobile advertising. However, the firms under question in the studies are big popular ones like NYT, and the same 'increased revenue' reasoning cannot be said to hold for small to medium sized firms who rely heavily on behavioral advertising for generating revenues. Thus, in view of the above mentioned arguments, an ecosystem possessing all the desired properties (a) - (d) mentioned in Section 1, cannot exist in practice. Mathematically, even for a relaxed version of this ecosystem, this has been proved by the authors in [19] - their main result being that the standard utilitarian social objective cannot be optimized for the case when aggregate consumer privacy preferences, represented via the popular differential privacy metric, are ideally homogeneous across the apps.

The Age and Rise Of Surveillance Capitalism - In her recent book [25], Shoshana Zuboff states with numerous real-life surveillance examples. since the early 2000's, of how our daily life activities are all recorded, rendered as behavioral data, processed, analysed, bought, bundled, and resold like sub-prime mortgages in a *behavioral futures* market. The litany of appropriated experiences is repeated so often and so extensively that we become numb, forgetting that this is not some dystopian imagining of the future, but the present. Originally intent on organising all human knowledge, Google ended up controlling all access to it (the process starting post the 9/11 attacks when the US government became liberal on surveillance of human data for security purposes and also coinciding with Google needing to boost their ROI for their glamorous investors) ; we do the searching, and are searched in turn. Setting out merely to connect us, Facebook found itself in possession of our deepest secrets. And in seeking to survive commercially beyond their initial goals, these companies realised they were sitting on a new kind of asset: our 'behavioural surplus', the totality of information about our every thought, word and deed, which could be traded for profit in new markets based on predicting our every need - or producing it. In a move of such audacity that it bears comparison to the enclosure of the commons or colonial conquests, the tech giants unilaterally declared that these previously untapped resources were theirs for the taking, and brushed aside every objection. While insisting that their technology is too complex to be legislated, there are companies that have poured billions into lobbying against oversight, and while building empires on publicly funded data and the details of our private lives they have repeatedly rejected established norms of societal responsibility and

[1] https://pagefair.com/blog/2017/adblockreport/

[2] According to Shoshana Zuboff, author of *The Age of Surveillance Capitalism: The Fight for a Human Future at the New*, we are at a critical juncture where we still have the power to decide what kind of world we want to live in, and what we decide now will shape the rest of the century. Our choices allow technology to enrich the few and impoverish the many, or harness it and distribute its benefits. Most critically, it shows how we can protect ourselves and our communities ensuring we are the masters of the digital rather than its slaves.

accountability. What is crucially different about this new form of exploitation and exceptionalism is that beyond merely strip-mining our intimate inner lives, it seeks to shape, direct and control them. Their operations transpose the total control over production pioneered by industrial capitalism to every aspect of everyday life. The extraction is so grotesque, so creepy, that it is almost impossible to see how anyone who really thinks about it lives with it âĂŞ and yet we do. There is something about its opacity, its insidiousness, that makes it hard to think about. Likewise the benefits of faster search results and turn-by-turn directions mask the deeper, destructive predations of what Shoshana Zuboff terms 'surveillance capitalism', a force that is as profoundly undemocratic as it is exploitative, yet remains poorly understood. Ignorance of its operation is one of the central strategies of this regime, and yet the tide is turning: more and more people express their unease about the surveillance economy and, disturbed by the fractious, alienated and trustless social sphere it generates, are seeking alternatives. It will be a long, slow and difficult process to extricate ourselves from the toxic products of both industrial and surveillance capitalism. *Till then, a workable solution might be to trade data in a fashion that benefits all in the data release ecosystem, and not just the data greedy firms.*

Structure of Data to be Traded - To cite an example of the structure of data that could be traded by sellers (e.g., mobile apps) having access to aggregate consumer data from their client base, parts of it that is assumed to be private, we choose a database to be a likely structure. As popular practical examples, the firm *BookYourData* (BYD) offers upstream buyers ready-made lists of contacts of business individuals across different industries, job titles, job functions, and job levels. A record in a list consists of contact information such as name, email address, job function, job department, country etc. The organization *SalesLead* (SL) maintains a variety of datasets of American businesses in the form of profession-based lists and state/province-based lists - the *Accountant Sales Leads* dataset contains records of US-based accountants, whereas the *Alabama Sales Leads* dataset contains records of different businesses (accountants, real-estate agents, etc.) based in Alabama. Each record in a dataset consists of contact information such as mailing address, geo-location, email address, phone number, etc. As another major example, the telemarketing company *TelephoneLists* specializes in offering its buyers phone lists as datasets that consists of information on consumers (contact details, demographics, etc.) as well as businesses (number of employees, sales, etc.) in North America.

Privacy Issues Arising from Trading - Given the just mentioned inevitability of data release from mobile ad-tech systems, trading data and subsequently privacy is a way to make most value out of consumer data. To emphasize this point, a recent survey conducted by the authors in [3] advocate consumers willing to trade data for incentives. However, one of the main privacy preservation barriers to realizing such trading, is the functioning methodology of the free mobile apps themselves. As such, free apps (the incentive) are only free in monetary terms; they come with the price of potential unwanted privacy leakage of the consumer, due to their reliance on personal consumer information to generate additional revenues, thereby leading to negative externalities being imposed on society. Now-a-days, mobile devices are a lot more intimate to users[3]; they are carried around at all times and are being used more and more for sensitive operations like personal communications, dating, banking, etc., each of which are conducted through multiple similar apps available on the app stores. Therefore, potential privacy leakage concerns arising from information collected by these apps for ad-personalization are more serious. One could argue here that paying for apps[4] would mitigate this issue, however, statistics prove that consumers around the world are more keen on using free apps compared to paid apps[5], and are also quite neutral to the collection of cookies by third parties, during browsing activities[6]. In addition, people in general exhibit the well known *privacy paradox* [2], wherein privacy conscious people arguably give up personal information with or without the presence of benefits, in a somewhat voluntary fashion. *Thus, in view of these reasons, preserving consumer privacy is a big challenge to realizing data trading ecosystems.* Recently, in an unpublished but thoroughly investigated research argument [17], Acquisiti, in contrast to his article in [1], states that behaviorally targeted advertising might increase the data holder's revenue but only marginally - thereby discouraging the idea of privacy trading purely from a profit standpoint. At the same time ad-marketers might having to pay orders of magnitude more to buy these targeted ads, despite the minuscule additional revenue they might generate for a data holder. However, he does add that the lack of privacy regulations in most parts of the world, the existence of"opaque blackbox" ad-exchanges, and lack of rigorous interrogatory research into the benefits of data release, make it very difficult to convince stakeholders of the un-necessity of data trading ecosystems - *thereby again making a case for the current inevitability of data release,* at least till the time when transparency of information flows is made a law and enforced well.

Towards a Positive Direction for Society - A deeper look into the results in [19] reveals that the inability to achieve a social optimal state in data trading ecosystems lie in (i) the hardness to satisfy strict consumer privacy preferences, and (ii) the inability to internalize the negative externalities due to privacy leakage, e.g., recent Facebook-Cambridge Analytica data scandal [23]. Thus, from a micro-economic perspective, one possible direction towards optimizing social welfare, i.e., efficiency, is to relax the strictness of privacy preserving preferences [20], thereby allowing consumers to compromise their ideal privacy requirements in return for benefits (e.g., monetary and non-monetary incentives). These benefits contribute to resolving the issue in (ii). The weight behind this idea lies in the fact that from a psychological perspective, most human beings are acceptable to making varied levels of compromises in real-life, especially for goods like privacy that have non-clear boundaries [3]. Note from Figure 1 that privacy compromises by consumers would result in apps selling more relevant personalized information to ad-networks (and thereby generating more revenue), the latter able to sell more ad-space to advertisers at an increased revenue, and the advertisers being able to target a broader personalized set of consumers. Thus, we have a win-win situation among all ecosystem entities. The big question then is: *what is the optimal way to compromise aggregate consumer privacy?* To answer this question, we propose a radical idea of combining the use of supply function

[3]The average American adult spends 2 hours and 51 minutes of time behind apps per day; *Source: Hackernoon*

[4]There are quite a few services that already offer some level of choice/configuration between full subscription (no ads, thus no third party privacy exposure) and fully advertisement/analytics paid for (i.e. "free"). Consequently there's the possibility of doing an empirical study to populate a model of peoples'(not yet evident that they are privacy-rational) "willingness to pay" in terms of utility function/curves for privacy/money.

[5]https://www.appsflyer.com/resources/state-app-spending-global-benchmarks-data-study/

[6]Statistic.com

framework from micro-economic theory [14] with mathematically rigorous information-theoretic privacy-preserving measures such as differential privacy (DP) to execute effective compromise in aggregate consumer privacy.

Privacy in Ad-Ecosystems without Trading - As an orthogonal concept to ours, the line of work proposed in [12] develops an inexpensive cryptography-based scalable system that allows targeted online advertisements to reach appropriate users thereby increasing click rates for advertisers; is privacy compliant with needs of standards organizations such as EFF, ACLU, and FTC; and satisfies the business needs of the ecosystem comprising data brokers, ad-networks, users, data brokers, and advertisers. However their system does not comment on the suitability to non-app online settings satisfying the "same-origin" policy. In addition, it does not deal with provable privacy-preserving mechanisms to handle inferential privacy attacks. In contrast, our proposed supply-function framework supports the use of composition-induced information-theoretic privacy preserving measures that are immune to inferential attacks.

3 THE SUPPLY FUNCTION FRAMEWORK

In this section, we introduce the idea of the supply function framework as an appropriate *regulated*[7] and rigorous[8] economic method to trade *group privacy* - the privacy of a group of app clients, rather than individual clients themselves,[9] in a manner so as to satisfy objectives (a) - (d) mentioned in Section 1, for a 'win-win' privacy trading ecosystem. A regulated entity here could be a government, induced via policies such as the GDPR. To this end, we first provide the conceptual working of the framework. We then provide a strong rationale behind the framework being an appropriate one for our problem at hand. Finally, we state practical use-cases where the supply function framework will be acceptable to consumers willing to trade on data.

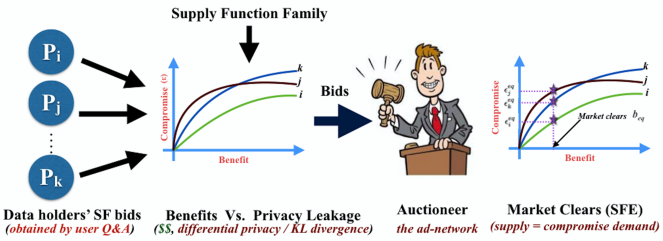

Figure 2: Illustrating Privacy Trading per Ad-Network

The Basic Idea - This auction-based economic framework, initially introduced by Klemperer and Meyer [14], demands apps provide as bids to their auctioneers (ad-networks) their own heterogeneous "supply functions" that mathematically characterize a tradeoff function between the amount of desired client group privacy compromise of an app versus the per-unit compromise benefit (monetary or otherwise) that is to be handed over to the app by the ad-network (see Figure 2). In practice, the supply function per app can be approximated from individual consumer Q&A, about their privacy preferences. Units of compromise are captured via the popular notion of ϵ - differential privacy (ϵ-DP). More specifically, for each value of ϵ on the y-axis of its supply function, the app expects a benefit/incentive of a certain value, on the x-axis of the function. The individual auctioneer, i.e., ad-network, has a demand of ϵ_d units of privacy compromise from its upstream entities, i.e., the advertisers/marketers. ϵ_d can again be approximated via Q&A dealings with these entities. (as a function of the composition-induced differential privacy metric) The individual auctioneers then work in the hope of clearing their own market, i.e., matches the compromise supply with advertiser demand and results in [10], a socially efficient competitive market (e.g., oligopoly, duopoly, perfect competition) equilibrium drives the optimal amount of compromise by each individual app at market equilibrium (i.e., ϵ_i^{eq} for each app i at market equilibrium such that $\sum_i \epsilon_i^{eq} = \epsilon_d$)[11] in return for a market equilibrium benefit value, b^{eq} (see Figure 2 for a conceptual illustration). In the case of multiple auctioneers (ad-networks) in a system, they work to form a parallel market resulting in heterogeneous market equilibrium parameters (see Figure 3). Note that a major advantage of the supply function framework is that no private consumer information such as the cost function for apps (part of app utility function) for compromised client group privacy, is shared with the auctioneer. The (parallel) markets in operation converge fast to market equilibrium using efficient market convergence algorithms that fall in the class of distributed gradient algorithms proposed in [4]. A detailed market analysis of the proposed ecosystem is provided in [20].

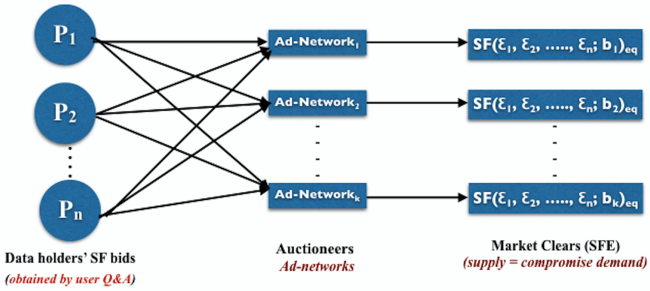

Figure 3: Privacy Trading with Multiple Ad-Networks

The Win-Win Implication - Thus far, we have commented on the socially efficient privacy trading market possibility as a result of deploying supply function economics. In terms of a win-win privacy trading solution, the advertisers/marketers get benefitted from the privacy leaks, induced by the ϵ_i^{eq}-DP values at market equilibrium. The ad-network intermediates between the data-holders/publishers and the advertisers, and balances it stance of preserving group privacy and satisfying advertisers through socially efficient ϵ_i^{eq}-DP values at market equilibrium. The apps intermediate between their

[7] Here, regulation follows the widely popular neoclassical microeconomic and Keynesian macroeconomic school of thought.
[8] The works in [1] and [16] only pitch the idea of privacy trading markets qualitatively.
[9] In practice, it is infeasible to preserve the individual privacy requirements of potentially hundreds of thousands of clients in the trading process.
[10] Current regulation now implicitly acknowledge that personal data is a commodity, tradeable, and subject to the laws of supply and demand [9].
[11] This summation relationship is in accordance with the composition property in differential privacy [8].

clients and ad-networks - on one hand they get commercially benefited by leaking a certain portion, i.e., the compromised portion, of personal client information to ad-networks. On the other hand, they protect the group privacy upto ϵ_i^{eq}-DP levels at market equilibrium. At the individual consumer level, things are not that straightforward - the consumers do compromise a certain acceptable amount of their privacy in return for targeted advertising benefits, but there might be a set of consumers whose privacy is exceedingly/unacceptably compromised with benefits, because privacy is preserved at the group level compared to an individual level.

Why SF Economics over Other Approaches? - The rationale behind the choice of auction-based models like the supply function mechanism, over the traditional economic Bertrand(B)[12] and Cournot(C)[13] oligopoly market competition structures is the resulting proven high market inefficiencies at equilibria under the (B /C) structures [13]. The rationale behind using the aforementioned supply function auction (SFA) mechanism over the well known Vickrey-Clarkes-Groves (VCG) multi-unit auction mechanism[14] is multifold [13]: (i) SFA adapts very well to variable compromise demands, and converges to market equilibria on existence, without the need to run the auction mechanism again, (b) as already mentioned above, like VCG, the SFA mechanism is private, i.e., private variables of apps are not shared with the auctioneers, (c) like VCG, the SFA mechanism is incentive compatible for the apps, i.e., the apps find it utility optimal to correctly report their compromise preferences to the auctioneers and not lie, (d) unlike multi-unit VCG, the SFA mechanism is fair in the sense that at market equilibrium, every app will get the same per-unit compromise benefit, and (e) in the case of market inefficiencies at equilibria, i.e., the case when social regulatory objective is not maximized at market equilibrium, the SFA mechanism results in bounded market inefficiency compared to the VCG mechanism.

Privacy Compromise Use-Cases in Practice - Among the possibly many scenarios apt for privacy compromise include settings related to fitness trackers, insurance industry, online bookstore, online music streaming industry, and the energy sector. Customers might be willing to compromise their personal, but yet 'not-so-very-personal' data like number of steps walked, number of steps climbed, quality of sleep etc. with fitness apps (e.g. *Strava*, *MapMyRun*, *Nike + RunClub*) in return for free recommendations like diet plan, exercise plan for weight reduction, to maintain a healthy lifestyle. In the insurance sector, consumers might compromise a bit of their privacy by sharing personal data like number of steps they walked, amount of time they spent on exercises everyday, with insurance company apps (e.g. *Oscar*, *UnitedHealthcare*) to receive reduced premiums. Customers might be willing to compromise their reading tastes with online bookstores (e.g., *Amazon Kindle*, *Booktopia*), which can potentially reveal their political/religious/social beliefs, in return for book recommendations [15]. In online music streaming sector, customers might be willing to compromise their musical tastes with streaming apps (e.g., *Sportify*, *Google Play Music*), in return of benefits like subscription discounts, or free services [10][11]. Similarly, in the energy sector, customers might be willing to share their (fine-grained) energy consumption data collected by smart meters, with utility apps (e.g. *First Utility*) for better quality of service even though this data can indicate approximate times householders tend to leave their home, or approximate times when they sleep. The common theme to all these use-cases is societal value creation of consumer data from applications. We should embrace this concept of societal value creation, at the same time appropriately respecting privacy constraints of individuals.

[12]In the Bertrand market competition, firms compete to arrive at a per-unit price (in our case benefit) equilibrium, where the price is a function of the quantity of resource (e.g., privacy compromise ϵ units in our case) produced [18].

[13]In the Cournot market competition, firms compete to arrive at a quantity (in our case, compromise) equilibrium, where the compromise amount is a function of the homogeneous price (benefit in our case) per unit of compromise [18].

[14]In the VCG mechanism, also popularly known as the second-price auction mechanism, participants (in our case, the apps) announce their bids (supply functions in our case) per unit of the item (compromise in our case), and the auctioneer selects as winner the participant with the highest bid, who needs to pay the second highest valuation.

4 FUTURE CHALLENGES TO TRADING

An immediate future challenge presents itself on the implementation front of our proposed idea of privacy trading - here, an important direction going forward is to approximate the supply function for various data holders. In practice, the supply functions derived via experiments might not fit the functional assumptions (e.g., differentiability) required for the theory to work. In such cases, the design of numerical approaches to reaching market equilibria need to be investigated. As other future challenges to effective real-world realization of our proposed privacy trading ecosystem, we must first ensure that privacy commodification should not contribute to unwanted additional privacy intrusions as mentioned by [6, 22]. As mentioned before, differential privacy does ensure group privacy leak checks at the ad-network, however, a proper regulatory supervision through proper contracts is necessary to ensure that obeying DP constraints by the ad-network is economically incentive compatible. Second, property rights are a challenge for apps to exercise despite incentive compatible contracts, when the personal data held by apps is often mixed with other data belonging to the app firm. This lack of boundary of data flow might make property rights for consumers too much of a challenge to implement and enforce, even in the presence of differential privacy tools, leading to higher transaction costs to be shouldered by regulators. More specifically, the optimal differential privacy ensuring statistical noise might be high enough to satisfy privacy constraints of the apps, but may be too high to make utility of the "other data" attached with the personal data. This calls for the design of new technical privacy preserving solutions that balance the privacy-utility tradeoffs. In addition, personal data contracts cannot specify all states of nature or all future actions, use, and shelf value of the data (perishable or non-perishable) in advance. When there are states or actions that cannot be verified and determined *ex post* by third parties, they are therefore not possible to be contractible *ex ante*, and results in incomplete contracts [7]. Incompleteness of contracts matter in terms of who has the power to take ownership related action, and the presumption is always that the economic actors (entities within an ad-network, in our case) will do so according to their interests. Aptly deciding who should have the ownership power to take certain actions is therefore a matter of foreseeing which actors will be most likely to act in the desired way. Given an appropriate decision framework, the transaction costs for realizing a privacy trading market would be low and an efficient market could exist [5]. In the case of non-low transaction costs, the need of the hour is the design of appropriate economic mechanisms. An interesting research direction is the static and dynamic valuation of consumer data. It is commonly known from the Metcalfe's law regarding data that static aggregated data is non-linearly more valuable to

analysts than its static individual counterpart. The effects are more pronounced when aggregate data is multi-dimensional, i.e., having multiple attributes. An important question to study here is how the ecosystem would function when privacy cost and utility functions would reflect such valuation metrics. With respect to dynamic valuation of consumer data, an important challenge towards trading efficacy lies in predicting/estimating the value of (multi) dimensional data over time. The importance of resolving this challenge has implications to apps or data collection firms to strategize the allocation of a limited budget towards collecting quality data on relevant dimensions, obeying privacy preservation constraints. Apart from the data valuation aspects, one important subject of concern is the veracity of quality data. As the latter is traded for money or incentives, there is most likely to be fraud and consequently the ecosystem must ensure via the design of proper crowdfunding mechanisms that statistically there are enough number of samples for a given data type before it decides to pay apps. Finally, there is the need for research on alternative trading structures to capture the many-many interactions between multiple data sellers and multiple data buyers, compared to the many sellers, one buyer trading structure illustrated in this article.

5 CONCLUSION

The current personal data ecosystem is mostly de-regulated, fragmented, and socially inefficient. As steps to improve social welfare, we commented on the the existence and design of efficient and regulated consumer-data releasing ecosystems aimed at achieving a maximum social welfare state amongst competing data holders, and creating societal value of consumer application data. In view of (a) the behavioral assumption that humans are 'compromising' beings, (b) privacy not being a well-boundaried good, and (c) the practical inevitability of inappropriate data leakage by data holders upstream in the supply-chain, we showcased the idea of a regulated (in the neoclassical microeconomic and Keynesian macroeconomic sense) and radical privacy trading mechanism that preserves the *heterogeneous* privacy preservation constraints (at an aggregate consumer, i.e., app, level) upto certain compromise levels, and at the same time satisfying commercial requirements of agencies (e.g., advertising organizations) that collect and trade client data for the purpose of behavioral advertising. Our radical idea merged *supply function economics*, introduced by Klemperer and Meyer, with differential privacy, that, together with their powerful theoretical properties, leads to a stable and efficient, i.e., a maximum social welfare, state, and that too in an algorithmically scalable manner. We also discussed interesting additional techno-economic challenges related to realizing effective privacy trading ecosystems.

ACKNOWLEDGEMENTS

We thank Prof. Edgar J. Whitley of the London School of Economics and Political Science for his detailed and constructive comments on our paper. We would also like to thank Prof. Swades De of the Indian Institute of Technology Delhi, and Professors Pan Hui and Sasu Tarkoma of the University of Helsinki for their useful feedback on the manuscript.

REFERENCES

[1] Alessandro Acquisti, Curtis Taylor, and Liad Wagman. 2016. The economics of privacy. *Journal of Economic Literature* 54, 2 (2016), 442–92.
[2] Susan B Barnes. 2006. A privacy paradox: Social networking in the United States. *First Monday* 11, 9 (2006).
[3] Volker Benndorf and Hans-Theo Normann. 2018. The willingness to sell personal data. *The Scandinavian Journal of Economics* 120, 4 (2018), 1260–1278.
[4] Dimitri P Bertsekas and John N Tsitsiklis. 1989. *Parallel and distributed computation: numerical methods*. Vol. 23. Prentice hall Englewood Cliffs, NJ.
[5] Ronald H Coase. 1960. The problem of social cost. In *Classic papers in natural resource economics*. Springer, 87–137.
[6] Julie E Cohen. 2012. What privacy is for. *Harv. L. Rev.* 126 (2012), 1904.
[7] Mathias Dewatripont and Jean Tirole. 1994. A theory of debt and equity: Diversity of securities and manager-shareholder congruence. *The quarterly journal of economics* 109, 4 (1994), 1027–1054.
[8] Cynthia Dwork, Aaron Roth, et al. 2014. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science* 9, 3–4 (2014), 211–407.
[9] Moritz Godel, Annabel Litchfield, and Iris Mantovani. 2012. The value of personal information: Evidence from empirical economic studies. *Communications & Strategies* 88 (2012), 41–60.
[10] David Greenberg. [n. d.]. The Open Book: What Your Reading Choices Say About You. *The Independent* ([n. d.]).
[11] David M. Greenberg, Simon Baron-Cohen, David J. Stillwell, Michal Kosinski, and Peter J. Rentfrow. 2015. Musical Preferences are Linked to Cognitive Styles. *PLOS ONE* 10, 7 (07 2015), 1–22. https://doi.org/10.1371/journal.pone.0131151
[12] Saikat Guha, Bin Cheng, and Paul Francis. 2011. Privad: Practical privacy in online advertising. In *USENIX conference on Networked systems design and implementation*. 169–182.
[13] Ramesh Johari and John N Tsitsiklis. 2011. Parameterized supply function bidding: Equilibrium and efficiency. *Operations research* 59, 5 (2011), 1079–1089.
[14] Paul D Klemperer and Margaret A Meyer. 1989. Supply Function Equilibria in Oligopoly Under Uncertainty. *Econometrica* 57, 6 (1989), 1243–1277.
[15] Wendy L. Patrick. [n. d.]. The Open Book: What Your Reading Choices Say About You. *Psychology Today* ([n. d.]).
[16] Kenneth C. Laudon. 1996. Markets and Privacy. *Commun. ACM* 39, 9 (Sept. 1996), 92–104. https://doi.org/10.1145/234215.234476
[17] Natasha Lomas. 2019. The case against behavioral advertising is stacking up. *TechCrunch* (Jan 2019). https://techcrunch.com/2019/01/20/dont-be-creepy/
[18] Andreu Mas-Colell, Michael Dennis Whinston, Jerry R Green, et al. 1995. *Microeconomic theory*. Vol. 1. Oxford university press New York.
[19] Mallesh M Pai and Aaron Roth. 2013. Privacy and mechanism design. *ACM SIGecom Exchanges* 12, 1 (2013), 8–29.
[20] Ranjan Pal, Jon Crowcroft, Abhishek Kumar, Pan Hui, Hamed Haddadi, Swades De, Irene Ng, Sasu Tarkoma, and Richard Mortier. 2018. *Privacy markets in the Apps and IoT age*. Technical Report. University of Cambridge, Computer Laboratory.
[21] Alvin E Roth. 1991. Game theory as a part of empirical economics. *The Economic Journal* 101, 404 (1991), 107–114.
[22] Julie Tuan. 2000. US West v. FCC. *Berkeley Tech. LJ* 15 (2000), 354.
[23] Wikipedia. 2018. Facebook-Cambridge Analytica data scandal. (2018).
[24] Wikipedia. 2018. General Data Protection Regulation. (2018).
[25] Shoshana Zuboff. 2019. *The age of surveillance capitalism: the fight for the future at the new frontier of power*. Profile Books.