

Marketing Analytics and Privacy

Jaap Wieringa

University of Groningen

Thomas Reutterer

WU Vienna University of Economics and Business

René Laub

Goethe University Frankfurt

PK Kannan

Robert H. Smith School of Business, University of Maryland

Michael Platzer

MOSTLY AI

Gilian Ponte

Rijksuniversiteit Groningen

Cite as:

Wieringa Jaap, Reutterer Thomas, Laub René, Kannan PK, Platzer Michael, Ponte Gilian (2021), Marketing Analytics and Privacy. *Proceedings of the European Marketing Academy*, 50th, (102341)

Paper from the 50th Annual EMAC Conference, Madrid, May 25-28, 2021



Special Session: “Marketing Analytics and Privacy”

Session chairs: Jaap Wieringa and Thomas Reutterer

Jaap Wieringa is a Full Professor of Research Methods in Business at the Department of Marketing, Faculty of Economics and Business, University of Groningen, PO Box 800, 9700 AV Groningen, The Netherlands. Email: j.e.wieringa@rug.nl

Thomas Reutterer is a Full Professor of Marketing at the Department of Marketing, Vienna University of Economics and Business (WU Vienna), Austria. Email: thomas.reutterer@wu.ac.at

Paper 1

Title: “The Value of User Tracking and Behavioral Targeting for Publishers”

Authors: **Rene Laub**, Klaus M. Miller, Bernd Skiera, Goethe University Frankfurt, Germany.

Paper 2

Title: “A World Without Cookies: Challenges for Attribution and Media Mix Modeling”

Authors: **P.K. Kannan** and Zipei Lu, Robert H. Smith School of Business at the University of Maryland, USA.

Paper 3

Title: “AI-based re-identification exposes privacy risk of behavioral data. A case for synthetic data”

Authors: **Michael Platzer**, Thomas Reutterer, Stefan Vamosi, Vienna University of Economics and Business (WU Vienna), Austria.

Paper 4

Title: “Privacy-preserving Generative Adversarial Networks to Share Data and Derive Marketing Insights”

Authors: **Gilian Ponte**, Jaap Wieringa, University of Groningen, The Netherlands.

Abstract

In almost everything we do, we leave a trail of data that exposes our interests, traits, beliefs and intentions. In doing so, we provide information to firms and governmental institutions, which allows them to track individuals and collect personal information to predict customer behavior and deliver an unprecedented level of personalization. This offers great opportunities for marketing analytics. At the same time, increasing privacy concerns and stricter privacy regulations within the European Union and the United States limit firms in their ability to process individual-level data and to develop effective marketing analytics programs.

The goal of this special session is to address the marketing challenges in today's privacy environment. Our joint contribution is to strike a balance between data-utility and privacy protection. We cover the trade-off from data collection to deriving insights in a privacy-preserving way.

Laub, Miller and Skiera derive the monetary value that publishers are losing due to restrictions on online tracking. In their paper entitled "The Value of User Tracking and Behavioral Targeting for Publishers", they decompose the value of a third-party cookie into a privacy-intrusive part, that stems from behavioral targeting, and a privacy non-intrusive part, that stems from other features such as success measurement of online advertising. The authors further illustrate how the value of a cookie differs across publishers in the market.

Kannan and Lu discuss how with increased privacy concerns, tracking of consumers becomes more and more difficult. Consequently, marketing technology solutions may no longer be helpful for marketers. In their presentation "A World Without Cookies: Challenges for Attribution and Media Mix Modeling" they outline possible work-arounds using modeling approaches that preserve privacy but also enable attribution estimates and optimal allocation of media for effective outcomes.

Platzer, Reutterer and Vamosi investigate the utility-privacy trade-off in the context of anonymizing behavioral marketing data. In their presentation "AI-based re-identification exposes privacy risk of behavioral data. A case for synthetic data", they demonstrate that standard "anonymization" techniques fail to protect individual-level sequences of behavioral data. They show that data synthetization can effectively reduce the risk of privacy intrusion and help to conserve the value of the data for data-driven marketing.

Ponte and Wieringa develop generative adversarial networks (GANs) that generate individual-level artificial data that, when analyzed, deliver the same marketing insights as real consumer data. In their presentation “Privacy-preserving Generative Adversarial Networks to Share Data and Derive Marketing Insights”, they show that the artificial data estimations occasionally even outperform real data estimations in terms of predictive validity. Their approach allows for data sharing even under strict privacy regulations.

Paper 1 - The Value of User Tracking and Behavioral Targeting for Publishers,

Rene Laub, Klaus M. Miller, Bernd Skiera, Goethe University Frankfurt, Germany.

It is common practice in the online advertising industry to collect a record of a user’s online browsing history via various tracking technologies. Advertisers use this information to increase the performance of online advertising (e.g., Braun & Moe 2013). Advertisers draw on a user’s browsing history to infer a user’s interests for targeting or personalization of online advertising—a practice to which is often referred to as behavioral targeting (e.g., Goldfarb & Tucker 2011). The ability to observe CTRs and conversion rates, via user tracking technologies, allows advertisers to measure the success of online ads and thereby improve the performance of online advertising. Among others, advertisers also use this information for frequency capping of online ads (Sahni 2015) or cross-site attribution modeling (Berman 2018).

However, with the increasing discussion on the protection of consumer privacy, the tracking of a user’s browsing history is becoming more and more controversial (Beke et al. 2018). Some authors argue that the excessive usage of behavioral targeting led to this development (John 2018). Thus, policy makers have put forward regulations to restrict the collection and usage of a user’s online browsing history such as the recently introduced General Data Protection Regulation (GDPR) in the European Union or the California Consumer Privacy Act (CCPA) in the United States. As a result, online advertisers tend to lose the ability to target and personalize online ads and lose the ability to measure the success of online ads. This development might have an important economic consequence for publishers because publishers rely heavily on income from online advertising to finance their editorial content that is often offered free of charge. Naturally, this advertising income follows directly from advertisers’ willingness-to-pay (WTP) to serve certain users with an online ad. Thereby,

previous analytical studies suggest that the availability of consumer information (e.g., their browsing history) enhances advertisers' WTP (Chen & Stallert 2014). This higher WTP should result in higher prices for ads and thereby in higher advertising income for publishers. As a consequence, disabling tracking technologies is likely to decrease advertisers' WTP and, in consequence, endangers publishers' revenue streams.

So far, only few and diverting empirical evidence is available that uses observational data to quantify the potential loss of disabling online user tracking technologies, such as cookies, to the online advertising industry. Previous studies found a rather wide range of price differences (4% - 70%) between situations where user tracking was possible versus situations without user tracking (e.g. Johnson et al. 2020). As a results, policy makers, that need to balance the commercial interests of the online advertising industry and the privacy protection of users, are missing a clear empirical guideline for policy decisions. If there is no substantial value decrease for the online advertising market in a world without user tracking, the trade off, would only exist in theory.

Given the diverse findings of previous empirical studies, further research is needed to understand the value differences of a cookie and the underlying drivers. Thus, the aim of our paper is to contribute to this need by empirically quantifying the value of a cookie and by further decomposing the value in its drivers. We aim at answering the question (1) what the overall value of a cookie for a publisher is? We further aim at understanding differences in the value of a cookie by (2) examining the effect for each publisher, and provide empirical explanations for differences between publishers that could have led to the differences from previous findings. Lastly, we (3) disentangle the value of a cookie into the part that stems from behavioral targeting and the part that originates from other features that are enabled by cookies like success measurement of online advertising. If behavioral targeting is not the major value driver of a cookie, then policy makers could also design polices that restrict the usage of behavioral targeting but allow user tracking for the other purposes of user tracking, like ad success measurement, which might be more widely accepted by consumers.

To empirically investigate our research questions, we draw on a data set from a large European ad-exchange including ~ 1.4 billion users with cookies, ~100 unique publishers and in total more than 42 million ad-impressions observed over a period of two weeks in spring 2016.

We quantify the potential losses from disabling online-user tracking by comparing prices that were paid for ad-impressions from users with and without third party cookies. Initial

empirical analyses reveal significant differences between those prices. For example, we find an average paid ad-impression price of 0.69 EUR Cost-per-Mille (CPM) for ad-impressions with third-party cookies and 0.39 EUR CPM without third-party cookies. We show that the value of a cookie strongly differs between publishers and find that the size and reputation of a publisher drive these differences. We further show that behavioral targeting only creates a small share of the value of a cookie and much of the value of a cookie comes from other tracking-based features.

Paper 2 - A World Without Cookies: Challenges for Attribution and Media Mix Modeling

P.K. Kannan and Zipei Lu

PK Kannan is Dean's Chair in Marketing Science and Zipei Lu is a doctoral student in marketing at the Robert H. Smith School of Business, University of Maryland, College Park, MD 20742, USA.

With increased governmental privacy regulate, and self-regulation by firms in the wake of consumer privacy concerns, we are entering into an era where tracking of consumers in online and offline environments is going to become increasingly sparse. In recent years customer identity graph based on third-party cookies has fueled marketing technology solutions for advertisement targeting, fraud detection, attribution modeling and media mix personalization. However, with the change in the privacy environment, many such solutions may not be helpful for marketers. In this presentation, we describe the changing environment and the types of challenges that marketing solutions such as attribution and media mix modeling will face. We also outline possible work arounds using modeling approaches that preserve privacy but also enable attribution estimates and optimal allocation of media for effective outcomes. We describe our approach below.

Budget allocation and spending based on marketing mix modeling is generally implemented in a hierarchical fashion within a firm. The model may be run annually, semi-annually and budget allocated based on an optimization model (usually a variation of the knapsack formulation). In a typical organization, the budget is allocated to each silo/channel – for example, TV, social, display or search – and each silo (channel) tends to allocate it without much recourse to what other channels do. There may be adjustments in these allocations based on refinements and running the model more frequently. The planning cycle for each channel varies too from monthly, weekly to daily. The Multi-Touch Attribution model

captures the net impact of the implementations in each channel on conversions and allocates the credit to each touchpoint, which when aggregated over time (say a week to 2 weeks) to the channel level, provides a way to relate the marketing investment in each channel to the net conversions attributed to that channel in that time period.

- a. If conversions can be counted against a particular channel – say, conversions through display, conversions through search etc., (using, say, last touch attribution), then using MTA these conversions can be adjusted to arrive at the attributed conversion through each channel.
- b. If there are variations in the spending in each channel across such short time-periods, then a model (we call it a sub-model) relating these spends in each channel to the conversion outcomes can provide us with elasticity estimates for spends in each channel.
- c. These estimates can then be used to estimate the optimal re-allocation of budget across these channels. This is bottom-up approach to budget allocation but involves only a subset of the channels. This approach would work only if the spending in higher level channels – such as TV and print – are assumed constant during this short-time, so that the variation is lower level channel spends and the related attributed conversions are used to identify their elasticities.
- d. A number sub-models are estimated and these estimates are related to the metamodel, where changes in spends occur in the higher level channels – like TV, print, etc., - where spend decisions are made for longer periods and more infrequently as compared to, say, digital channels. The digital channel spends and outcomes may be aggregated at this level (that is, variation within these channels are not considered as the data is aggregated over the shorter-periods to make it consistent with time-periods where variations occur for the spends in higher level channels. These runs can be used to reallocate budgets across higher level channels and aggregated lower-level channels.

The problem with Media Mix Models is that it estimated at an aggregate level (say, monthly data) while variations in lower-level channels occur across weeks with a month and thus it may not provide as fine-grained estimates of elasticities as needed for good allocations. The hierarchical model proposed above would be able to provide such elasticities for fine-grained analysis across lower-level channels. At the same time, providing estimates from the submodels to the refinement of meta-model will enable changes in allocations in higher level

channels taking into account the interactions at the lower level and also incorporating the attribution estimates in the lower level sub-models.

The implementation would proceed from the top with an MMM, provide initial allocations based on monthly or quarterly level aggregate data. The lower-level budget spends will incorporate variations over shorter intervals. The MTA will provide attribution estimates (the cadence could be daily or weekly), which will then be incorporated into the estimation of weekly sub-models to provide estimates for elasticities for lower-level channels – budget reallocation will be done based on this if needed. When enough sub-models are done over a quarter or half-year, the higher-level model will be run to reallocate budgets based on the estimates. But with privacy preservation constraints, the MTA models will necessarily become MMM models at the lowest level. The estimation technique will be Bayesian as it will allow incorporation of priors and is also very consistent with the hierarchical nature of this integration. The project will focus on this, first using simulated data to understand how the approach would work, and then taking it to real data.

Paper 3 - AI-based re-identification exposes privacy risk of behavioural data. A case for synthetic data

Michael Platzer, Thomas Reutterer, and Stefan Vamosi

Michael Platzer is Founder and CSO at MOSTLY AI Solutions, Hegelgasse 21/3, A-1010 Vienna, Austria. Email: michael.platzer@mostly.ai

[Thomas](#) Reutterer is a Full Professor of Marketing and Stefan Vamosi is a Research and Teaching Associate and doctoral student at the Department of Marketing, Vienna University of Economics and Business (WU Vienna), Austria. Email: thomas.reutterer@wu.ac.at; stefan.vamosi@wu.ac.at

The steady rise of digital native business formats witnesses the tremendous opportunities offered by an exploding amount and variety of individual-level, behavioural micro-data accruing in a broad range of industries. For example, firms like Amazon, Netflix or Facebook track the behaviour of their customers to derive personalized recommendations and targeted marketing actions. Other companies realize that sharing customer information with other parties (e.g., linked with “Internet of Things” elements, such as mobile tracking meters, medical or fitness devices, etc.) can create synergies for both sides. Likewise, the non-profit

sector and research institutions increasingly rely on the availability or “shareability” of publicly available or open behavioural data (Beaulieu-Jones et al. 2019).

However, all these benefits are in strong contrast to the legitimate desire of individuals to protect their privacy and to refrain from sharing their personal data (Wieringa et al. 2021). In the vein of the Facebook-Cambridge Analytica scandal also firms have increasingly become sensitive to protect their customer data against re-identification attacks and their brands against a loss in customer trust (Schneider et al. 2017, 2018). All these concerns acuminate in modern privacy regulations (in particular EU’s GDPR and California’s CCPA) which impose very strict standards of data anonymization. Both the GDPR and the CCPA do not specify any specific process for anonymization, but they demand the outcome to be irreversible to prevent re-identification of individuals.

Against this background, the key challenge for many firms is to keep benefiting from data-driven marketing while maintaining the privacy of their customers’ data. Previous research already showed that conventional perturbation techniques (such as adding random noise, masking, or obfuscation) fail to do so in the presence of high dimensional, highly correlated data, which typically arise when observing individuals over an extended period of time, i.e. for sequential personal data. For example, Narayanan and Shmatikov (2008) document successful re-identification attacks in the context of Netflix user’s movie ratings, De Montjoye et al. (2013) for human mobility traces and De Montjoye et al. (2015) for credit card retail transaction data.

In this research, we extend this perspective and show that a powerful general-purpose, AI-based model recently proposed by Vamosi et al. (2020) is capable to re-identify behavioural data in a highly effective way. As we will demonstrate in detail, this makes standard “anonymization” techniques inapt to protect individual-level sequential data. We also show that data synthetization can help to manage the trade-off between preserving the useful information in the original data, while reducing the risk of violating privacy.

To this end, we use customers’ clickstream histories available to us via the 2018 ComScore Web Behavior Panel for conducting a series of re-identification experiments. In our anonymization experiments, we mimic the re-identification task of a potential intruder. In doing so, the intruder just observes a short sub-sequence or “snippet” of a non-private (i.e., presumably “anonymous”) clickstream sequence. The task is to make inferences about the user’s identity by comparing the non-private snippet with a private one from the “leaked”

original dataset. Our experimental setup varies (i) the length and location of the “browsing snippet” along the time horizon of the available flow of clickstream data, (ii) anonymization techniques used for privacy protection and (iii) levels of data obfuscation.

Our findings clearly demonstrate the inability of conventional anonymization techniques (more precisely, we tested sequence flipping and shuffling) to protect behavioural customer data from re-identification without completely distorting their usefulness in terms of preserving relevant information contained in the original data. For example, flipping 40 percent of a browsing sequence destroys most of the structural information reflected by the behavioural patterns, but still leaves more than 60 percent (!) of the individuals unprotected and therefore re-identifiable; and even adding 80 percent noise to the data does not protect all individuals. In contrast, using a deep-learning based generative model approach to convert the original data into synthetic data makes re-identification virtually impossible, but retains the behavioural characteristics of individuals and thus conserves the value of the data for data-driven marketing. We conclude that synthetic data is a viable way to cope with the privacy-preservation vs. data utility trade-off.

Paper 4 - Privacy-preserving Generative Adversarial Networks to Share Data and Derive Marketing Insights

Gilian Ponte and Jaap Wieringa

Gilian Ponte is a PhD Candidate and Jaap Wieringa is a Professor of Research Methods at the Department of Marketing, Faculty of Economics and Business, University of Groningen, PO Box 800, 9700 AV Groningen, The Netherlands. Email: g.r.ponte@rug.nl; j.e.wieringa@rug.nl

Privacy is a fundamental human right. Over the years, we observe an increase in privacy concerns due to the growing amount of data and the development of methodologies that pressure the fundamental right to privacy. The marketing literature has defined privacy as one of its biggest priorities and the importance of privacy in marketing continues to increase over the years (Wedel and Kannan 2016; Wieringa et al. 2021). With unprecedented access to data and marketing analytics as a top priority, firms annually spend around \$36 billion to benefit from customer data. Simultaneously, the vast amount of investment in leveraging customer data combined with the growth of data and possibilities to capture individual customer data lead to privacy concerns among individuals (Acquisti et al. 2015, 2016).

Ultimately, these two faces of marketing analytics pose a trade-off between data utility and privacy where the goal is to balance the ability to derive meaningful insights while preserving the privacy of the individual (Wedel & Kannan 2016; Wieringa et al. 2021). In this paper, we address the Marketing Science Institute (2020) research priority to fundamentally alter the rising concerns about data privacy. Specifically, we aim to resolve the trade-off between marketing insights and privacy, and show that it is possible to pursue both. We contribute to the marketing literature by showing that academics and firms are able to generate privacy-friendly artificial data, derive real marketing insights, improve the predictive validity of estimations, promote data sharing even under privacy regulations and accelerate scientific progress in and outside the field of marketing.

To deal with this delicate balance between privacy and data utility, we draw on the artificial intelligence (AI) literature on generative modeling. Goodfellow et al. (2014) show that generative adversarial networks (GANs) can approximate *any* data-generating process. GANs consist of two neural networks that compete. We can use maximum likelihood estimation to train two neural networks that are guaranteed to approximate any data-generating process to any degree of accuracy. GANs make a particularly interesting case to address the utility-privacy trade-off: We are able to generate artificial data that do not exist in reality and simultaneously maintain the utility of the data. During training, one player indirectly maximizes the likelihood of observing the real data through the other player, which allows the artificial data never to be identical to the real data, only in distribution. To possibly increase the utility of the data, we can use the capacity of GANs for feature extraction. GANs only learn to extract important factors of variation from the data. We are able to vary the capacity to only extract the important factors that are of interest to overcome the curse of dimensionality, which makes for sound statistical inference and can possibly increase generalization for prediction tasks.

Compared to data common in the AI literature, consumer data are often characterized by combinations of univariate marginals from a variety of distributional families (Danaher and Smith 2011). For example, a typical consumer data set may consist of a combination of unit sales, purchase decision and time until purchase following a log-normal, binomial and exponential distribution, respectively. These characteristics of consumer data present an additional challenge. We present an approach that overcome the main difficulties associated with Markov chains, scale to an arbitrary amount of dimensionality, and account for the complex joint distribution that shapes marketing data sets (Goodfellow et al. 2014).

However, GANs are in an early stage of development and notoriously difficult to train. The fact that we are able to *represent* any data generating process does not imply that we are guaranteed to *learn* the data generating process. For example, due to the non-convex nature of the optimization problem we might get stuck in a local minimum, a saddle point or not even arrive at a critical point. In this study, we introduce GANs to investigate the ability to generate privacy-preserving data. Subsequently, our research question is as follows: “*How are GANs able to generate privacy-preserving data that maintains the ability to derive meaningful insights?*”.

One of our empirical contributions is a simulation study. We investigate the bias in the parameter estimates on artificial data, identify potential performance measures to monitor during training, and investigate the large sample properties of GANs. Next, we apply GANs in the context of marketing applications to three consumer data sets, each with different characteristics. We are able to derive the same real marketing insights from artificial data estimations, and artificial data estimations occasionally outperform real data estimations in terms of predictive validity. Furthermore, we compare a variety of GAN architectures in terms of convergence speed and ability to learn the data generating process.

References

- Acquisti, A., Brandimarte, L., & Loewenstein, G. (2015). Privacy and human behavior in the age of information. *Science*, 347(6221), 509-514. doi:10.1126/science.aaa1465
- Acquisti, A., Taylor, C. R., & Wagman, L. (2016). The Economics of Privacy. *Journal of Economic Literature*, 52(2). doi:10.2139/ssrn.2580411
- Beaulieu-Jones, B. K., Wu, Z. S., Williams, C., Lee, R., Bhavnani, S. P., Byrd, J. B., & Greene, C. S. (2019). Privacy-preserving generative deep neural networks support clinical data sharing. *Circulation: Cardiovascular Quality and Outcomes*, 12(7), e005122.
- Beke, Frank T., Eggers, F., Verhoef, P. (2018). Consumer Informational Privacy: Current Knowledge and Research Directions. *Foundations and Trends in Marketing*, 11 (1), pp 1–71.
- Berman, R. (2018). Beyond the Last Touch: Attribution in Online Advertising. *Marketing Science*, 37 (5), 685-853.

- Braun, M., Moe, W. (2013). Online Display Advertising: Modeling the Effects of Multiple Creatives and Individual Impression Histories. *Marketing Science*, 32 (5), 753–767.
- Chen, J., Stallaert, J. (2014). An Economic Analysis of Online Advertising Using Behavioral Targeting. *MIS Quarterly*, 38 (2), 429-449.
- Danaher, P., & Smith, M. (2011). Modeling Multivariate Distributions Using Copulas: Applications in Marketing. *Marketing Science*, 30(1), 4-21.
<http://dx.doi.org/10.1287/mksc.1090.0491>
- De Montjoye, Y. A., Radaelli, L., & Singh, V. K. (2015). Unique in the shopping mall: On the reidentifiability of credit card metadata. *Science*, 347(6221), 536-539.
- De Montjoye, Y. A., Hidalgo, C. A., Verleysen, M., & Blondel, V. D. (2013). Unique in the crowd: The privacy bounds of human mobility. *Scientific Reports*, 3, 1376.
- Goldfarb, A., Tucker, C. (2011). Privacy Regulation and Online Advertising. *Management Science*, 57 (1), 57-71.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., . . . Bengio, Y. (2014). Generative adversarial networks. *Advances in Neural Information Processing Systems*, 27, 2672-2680. doi:<http://papers.nips.cc/paper/5423-generative-adversarial-nets.pdf>
- John, L. K., Kim, T., Barasz, K. (2018). Ads That Don't Overstep. *Harvard Business Review*. Available online: <https://hbr.org/2018/01/ads-that-dont-overstep>
- Johnson, G. A., Shriver, S. K., Du, S. (2020). Consumer Privacy Choice in Online Advertising: Who Opt's Out and at What Cost to Industry? *Marketing Science*, (forthcoming).
- Kannan, P. K., Reinartz, W., Verhoef, P. C. (2016). The Path to Purchase and Attribution Modeling: Introduction to Special Section. *International Journal of Research in Marketing*, 33 (3), 449-456.
- Marketing Science Institute. (2020). *RESEARCH PRIORITIES 2020-2022*. Cambridge, MA: Marketing Science Institute.

- Narayanan, A., & Shmatikov, V. (2006). How to break anonymity of the Netflix prize dataset. *arXiv preprint cs/0610105*.
- Schneider, M. J., Jagpal, S., Gupta, S., Li, S., & Yu, Y. (2017). Protecting customer privacy when marketing with second-party data. *International Journal of Research in Marketing*, 34(3), 593–603.
- Schneider, M. J., Jagpal, S., Gupta, S., Li, S., & Yu, Y. (2018). A flexible method for protecting marketing data: An application to point-of-sale data. *Marketing Science*, 37(1), 153–171.
- Vamosi, S., Reutterer, T., & Platzer, M. (2020). A Deep Recurrent Triplet Loss Model for Learning Sequence Similarities. *Working paper*, WU Vienna.
- Wedel, M., & Kannan, P. (2016). Marketing Analytics for Data-Rich Environments. *Journal of Marketing*, 80(6), 97-121. doi:10.1509/jm.15.0413
- Wieringa, J., Kannan, P., Ma, X., Reutterer, T., Risselada, H., & Skiera, B. (2021). Data analytics in a privacy-concerned world. *Journal of Business Research*, doi:10.1016/j.jbusres.2019.05.005