Show me the Money: Measuring Marketing Performance in F2P Games using Apple's App Tracking Transparency Framework

Frederick Ayala-Gómez*
fayala@bn-mobile.com
Bandai Namco
Barcelona, Spain

Ismo Horppu*
ihorppu@zynga.com
Zynga
Helsinki, Finland

Erlin Gülbenkoğlu
erlin.gulbenkoglu@rovio.com
Rovio Entertainment
Espoo, Finland

Vesa Siivola
vesa.siivola@rovio.com
Rovio Entertainment
Espoo, Finland

Balázs Pejó
pejo@crysys.hu
CrySyS Lab, HIT-BME & ELKH-BME
Budapest, Hungary

Abstract

Mobile app developers use paid advertising campaigns to acquire new users. Based on the campaigns' performance, marketing managers decide where and how much to spend. Apple's new privacy mechanisms profoundly impact how performance marketing is measured. Starting iOS 14.5, all apps must get system permission for tracking explicitly via the new App Tracking Transparency Framework. Instead of relying on individual identifiers, Apple proposed a new performance mechanism called conversion value, an integer set by the apps for each user. The conversion value follows a set of rules and a schema that defines the integers based on the user's in-app behavior. The developers can get the number of installs per conversion value for each campaign. For conversion values to be helpful, we need a method that translates them to revenue. This paper investigates the task of attributing revenue to advertising campaigns using their reported conversion values. Our contributions are to formalize the problem, find the theoretically optimal revenue attribution function for any conversion value schema and show empirical results on past data of a free-to-play mobile game using different conversion value schemas.

CCS Concepts

• **Applied computing** → **Marketing**; • **Information systems** → **Online advertising**; • **Security and privacy** → *Privacy protections*; *Data anonymization and sanitization.*

Keywords

conversion value, revenue attribution, mobile advertising, privacy

*The author's current institution for contacting purposes. The research was done while at Rovio.

AdKDD '22, August 15, 2022, Washington DC, USA

ACM ISBN 978-1-4503-XXXX-X/22/08…$15.00
https://doi.org/XXXXXXX.XXXXXXX

ACM Reference Format:
Frederick Ayala-Gómez, Ismo Horppu, Erlin Gülbenkoğlu, Vesa Siivola, and Balázs Pejó. 2022. Show me the Money: Measuring Marketing Performance in F2P Games using Apple's App Tracking Transparency Framework. In *Proceedings of (AdKDD '22)*. ACM, New York, NY, USA, 6 pages. https://doi.org/XXXXXXX.XXXXXXX

1 Introduction

Marketers look for online advertising channels that deliver the best return on investment (ROI). Measuring ROI requires calculating the revenue that the ad campaign brought compared to the money spent. Calculating return on investment requires attributing revenue to campaigns. This task is known as attribution, and there are different approaches to it [16]. The most common attribution model in online advertising is last-click attribution, which gives all the credit to the last ad that the user engaged with [8]. Online advertising companies started building user profiles in search of higher ROI. User profiles allow further advertising optimizations where companies may target users with particular standards. On the other hand, the more collected information in the user profiles, the more users become attentive to what the companies know about them. Several surveys show that people are concerned about the control that companies have over their data, and they disagree with the data collection and sharing practices of online services [7, 14]. Governments have taken action to rule how companies use personal data, which led to significant legislative changes.

As a consequence, technological giants such as Google and Microsoft started to utilize privacy-preserving techniques [9, 11]. Apple has previously introduced various privacy features [28], and in September 2020 they introduced new version of the ad network API (SKAdNetwork 2.0) with support for a new framework called App Tracking Transparency (ATT). Starting iOS 14.5, app developers cannot share any tracking identifier to advertising networks unless users allow it. The ATT framework allows showing a pop-up dialog asking the user if they want to enable the application to track or not. This privacy innovation has a profound impact on how ad campaigns' performance is measured. Inevitably, the effectiveness of mobile advertising is affected, as the lack of the identifier affects how performance is measured and what types of ad personalization are available.

At a glance, the conversion values are a privacy-preserving mechanism proposed by Apple to measure an advertising campaign's performance without disclosing the user's origin. At its core, the conversion value groups users by assigning them an integer. The application developers are free to determine the bucket for each user based on the available information about them, given a set of rules that we will explain later. Apple reports the count of users with the same conversion values per campaign via postbacks. Based on the currently available information — the conversion values provide an ad-hoc privacy protection in the hide-in-the-crowd sense; such as k-anonymity [26], instead of differential privacy [20] which has a formal privacy guarantee.

Apple released iOS 14.5 in April 2021, and its impact showed on advertising networks in a few months. In the earning reports of 2021 Q3, various advertising networks, e.g., Snapchat [25], Facebook [12], Twitter [29], and Alphabet [2] expressed that there was an impact on their revenue due to the changes brought by the SKAdNetwork 2.0. There is no doubt that marketing teams need a way to measure the performance of marketing campaigns under the new privacy-preserving mechanisms.

Contribution. This paper investigates various conversion value schemas in combination with revenue attribution functions. Our contributions shed light on using the conversion values for attributing the revenue to the advertising campaigns. More specifically, *i)* we formalize the problem of revenue attribution based on conversion values, *ii)* we find the revenue attribution function, which minimizes the attribution error for any conversion value schema, *iii)* we show the revenue attribution quality of different conversion value schemas via back-testing on historical data.

2 Background & Related Work

This section introduces the concepts and methods used in the rest of the paper, such as conversion value, Identifier for Advertisers (IDFA), last-click attribution, and user origin. Moreover, we survey the related literature concerning privacy and revenue attribution using conversion values. For a general view of the challenges of privacy-centric digital advertising, refer to [17]. An overview of the changes in `SKAdNetwork 2.0` is presented in [22]. In addition, `SKAdNetwork 3.0` was released on May 2021, for details see [18]. Most recently, `SKAdNetwork 4.0` was announced on June 2022 [5] with promising improvements to the conversion values rules and reporting.

2.1 Revenue Attribution

Knowing the advertising campaign that brought a player to the game helps assign the revenue generated from the player to the campaign, which is needed to measure ROI. Conversion values use last-click attribution, and it is the most commonly used method in online advertising [8].

User Origin. The origin of a user can be organic, paid advertisement, or cross-promotion depending on how they found an app. The organic users installed the app without engaging with an ad. Paid origin users are those that installed the app after seeing an ad. Cross-promotion users installed the app following an ad from the developer portfolio. This paper focuses on attributing revenue from the paid and organic origins.

Identifier for Advertisers. In the Apple ecosystem, the Identifier for Advertisers (IDFA) allows tracking without disclosing the user's identity. Starting iOS 14.5, users can set their preference for app tracking globally or per app. The IDFA will serve its purpose only for the apps to which the user gives system consent, the user must give tracking consent in the app where the ad is shown and in the app that is being promoted.

Conversion Value. The conversion values appeared in Apple's developer documentation starting `SKAdNetwork 2.0` [3] via the `update-PostbackConversionValue` [4]. The conversion value is an integer $v \in [0, 63]$ that developers can set. It is often modeled using a binary representation of six bits. Each bit may represent an action of the users as a logical condition (e.g., user passed tutorial, reached a certain level). The conversion value is assigned for the first time when a user opens the app (i.e., not when the user installs the app). Developers can arbitrarily increment the value within 24 hours of the last update. If there has been no update within 24 hours, the advertiser receives a postback of the install after a random time between 0 to 24 hours. Practically, a period of up to seven days seems to be a maximum delay that makes sense, with many ad networks recommending much shorter windows (e.g., 24 hours) [1, 24].

Revenue Attribution using Conversion Values. The task of attributing revenue using conversion values is very recent. To the best of our knowledge, our work is the first that formally investigates the task. However, related work can be found on the Web. Closer to our work is [30], where the authors present two approaches that rely on the user's conversion value empirical conditional probability of belonging to a campaign. The first approach is winner takes all, which assigns the user (and its revenue) to the campaign with the highest empirical conditional probability. The second approach is probabilistic attribution, which multiplies the user's revenue by the empirical conditional probability of coming from each campaign given the conversion value, and sums it at the campaign level. We use an equivalent approach in Equation 3, but instead of attributing revenue using the user-campaign probabilities, we use the expected revenue per conversion value multiplied by the count of conversion values per campaign and show that this method is optimal.

2.2 Related Privacy Literature

Many privacy-preserving techniques were introduced in the last quarter-century. One of the most famous is k-anonimity [26] which requires that any user contained in the dataset cannot be distinguished from at least $k-1$ other users. The method's main drawback is that they define anonymity as a property of the dataset. Another widespread privacy mechanism is Differential Privacy [10] where anonymity is defined as a property of the process, making it resilient to any privacy attack based on background knowledge. It was adapted to numerous scenarios, each requiring its own fine-tuning of the definition [20].

We are unaware of any official well-detailed documentation concerning how Apple applies privacy to the conversion values, only press releases and blog posts [15, 22, 30]. Also when Apple announced using differential privacy, they did it without telling crucial elements [27]. Thus, Apple is unwilling to reveal details about their privacy mechanisms. Although this can also be seen as an additional level of protection, it is well-known and widely believed that security and privacy by obscurity are not a good idea. The reason originates from cryptography, where it is always assumed that the enemy knows the system being used [23].

3 Formalizing & Analyzing Conversion Values

This section illustrates the problem, formalizes it, and analyzes it. Our goal is to capture the scene with all its details via a flexible mathematical model (e.g., it is adaptable for future changes concerning the conversion value schema), as it is not sure how the conversion value schema is enforced. Yet, based on our empirical observations (concerning the conversion value schema) and despite the intricate nature of the problem, our analysis lands itself on a simple solution relating to the optimal revenue attribution function, which minimizes the difference between attributing revenue using conversion values and last-click attribution with IDFA. The variables used in the paper are introduced individually in this section as well as summarized in Table 1.

3.1 Problem Illustration

Initially, the app developers could distinguish between paid and organic users. Thanks to IDFA and attribution methods like last-click attribution, they could map the users to their origin, network, and advertising campaign. A common practice for measuring the ROI is to group users in cohorts based on their registration date, origin, network, campaign, and country. At the cohort level, one can aggregate the cost of acquiring the users and their revenue, which helps monitor ROI.

Symbol	Type	Meaning
i	Cons.	User ID, in-between 1 and $\lvert U_t^d \rvert$.
d_i	Cons.	User's registration date: the first time a user opens the app.
t	Cons.	Number of days for the revenue to be accumulated (e.g., 3, 7, 14, 30, and 90).
d	Cons.	The date when the conversion values are reported (sufficiently later than any d_i).
α_i	Cons.	User combined network and campaign ID: $\alpha = 100 \cdot n + c$. Note that $0 \leq c \leq 99$.
β	Cons.	The upper limit on α (i.e., $100 \cdot n + c < \beta$). $\alpha = \beta$ corresponds to the organic users.
r_i^t	Cons.	Accumulated revenue of the corresponding user i for the first t days after d_i.
$\mathcal{U}_i$	Set	User features dataset (i.e., remaining information about the user).
u_i^d	Set	$= (d_i, r_i^t, \alpha_i, \mathcal{U}_i)$, user data at d where campaign IDs are known.
v_i^d	Cons.	Conversion value of user i at d. Without subscript we mark the different conversion values.
$\tilde{u}_i^d$	Set	$= (d_i, r_i^t, v_i^d, \mathcal{U}_i)$, user data when only conversion values are available instead of α_i.
$f(\cdot)$	Func.	Conversion value schema or conversion value model (e.g., $f(u_i^d \setminus \{\alpha_i\}) = v_i^d$).
$x_{v,\alpha}^d$	Cons.	$\in X^d$, the count of users in v bucket at d corresponding to α.
y_α^t	Cons.	Accumulated last-click attribution revenue for α based on the first t days of the users.
$\tilde{U}_v^d$	Set	Set of all users (i.e., independently of d_i) with conversion value v^d, i.e., $\forall \tilde{u}_i \in \tilde{U}_v^d : v_i = v$.
$\bar{r}_v^t$	Cons.	The average first t days revenue of users in $\tilde{U}_v^d$ at d, i.e., $\bar{r}_v^t = \sum_{\tilde{U}_v^d} r_i^t / \lvert \tilde{U}_v^d \rvert$.
$pr_p(\cdot)$	Func.	Privacy preserving method with privacy threshold p.
$\hat{x}_{v,\alpha}^d$	Cons.	$\in \hat{X}^d = pr_p(X^d)$, the conversion value counts after applying the privacy protection.
$g_\alpha(\cdot)$	Func.	Function to attribute the revenue of α at d. Input: $(\tilde{U}_v^d, \hat{x}_{v,\alpha})$.

Table 1: Summary of the variables used in the paper.

For simplicity, we note user IDs as $i \in \{1, 2, \dots\}$ and the registration date of a user i as d_i. We are interested in r_i^t, which is the accumulated revenue of user i from their registration d until time t, so we must restrict ourselves to users with $d_i \leq d - t$. This is necessary; otherwise, the revenue attribution would become a prediction problem (because we would not know r_i^t), which has been studied extensively [13, 21]. The users satisfying this condition are captured as $u_i^d \in U_t^d$. For convenience, we define the combination of network ID $n \in \mathbb{N}$ and ad campaign ID $c \in [0, 99]$ to be $\alpha = 100 \cdot n + c$ because Apple restricts the number of campaigns per network to 100. We denote with β the total number of different network and campaign combinations, hence $0 \leq \alpha < \beta$. Note that the organic users correspond neither to any networks nor campaigns. Therefore, we capture them by setting their combined network and campaign ID to β.

Table 2 presents the initial dataset when IDFA is available. User-wise data is shown in Table 2a, and Table 2b shows the cumulative revenue y_α^t of the first t days for each ad network & campaign, which simplifies calculating the ROI. Formally, this represents the data corresponding to u_i^d as a tuple $\{d_i, r_i^t, \alpha_i, \mathcal{U}_i\}$. The tuple includes the registration date, the first t day revenue generated by the user, the user's origin, and — for the sake of completeness — it also contains $\mathcal{U}_i$, which captures any other related information about the user i such as event-level data within the app useful for building conversion value models.

With the enforcement of ATT, the data presented in Table 2 will not be available for the vast majority of the users, as explicit tracking system consent must be given. Instead, the application developers have two tables available. The first contains the conversion values v_i and the revenues r_i^t for all the app users, as presented in Table 3a. The second contains the aggregate count of conversion values X^d at time d, as shown in Table 3b. Developers can build Table 3a by keeping track of the conversion values assigned to each user.

User ID i	Revenue r^t	Net. ID n	Cam. ID c	α_i
1	0	4	05	405
2	2.99	–	–	β
3	0	3	89	389
⋮	⋮	⋮	⋮	⋮
$\lvert U \rvert$	4.99	1	71	171

(a) Available user-wise data.

α	Net. ID n	Cam. ID c	Revenue y^t
000	0	00	458
001	0	01	927
⋮	⋮	⋮	⋮
099	0	99	811
100	1	00	373
⋮	⋮	⋮	⋮
199	1	99	731
⋮	⋮	⋮	⋮
β	–	–	1639

(b) Available campaign-wise data.

Table 2: Illustration of the user data available to developers before Apple's ATT came out.

User ID i	1	2	3	$\cdots$	$\|U\|$
r_i^t	0	2.99	0	$\cdots$	4.99
v_i	0	19	1	$\cdots$	60

(a) Available user-wise data.

$\alpha \rightarrow$	000	001	$\cdots$	099	100	$\cdots$	199	$\cdots$
$v = 0$	19	40	$\cdots$	88	0	$\cdots$	36	$\cdots$
$v = 1$	55	10	$\cdots$	31	79	$\cdots$	151	$\cdots$
$v = 2$	29	22	$\cdots$	34	58	$\cdots$	2	$\cdots$
$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\vdots$
$v = 63$	18	346	$\cdots$	54	7	$\cdots$	189	$\cdots$
Sum	458	927	$\cdots$	811	373	$\cdots$	731	$\cdots$

(b) Available campaign-wise data.

Table 3: Illustration of the user data available to developers after ATT came out.

Formally, when IDFA is not available, the user's tuple u_i^d contains the same data, but instead of α_i, the user i's conversion value v_i will be included. We encapsulate this with $\tilde{u}_i^d = \{d_i, r_i, v_i, \mathcal{U}_i\}$. The conversion value itself is computed from available user data via a conversion value schema f, i.e., $f(u_i^d \setminus \{\alpha_i\}) = v_i$.

Note that this model simplify the scenario, as we do neither consider the 24h delay from last conversion value update nor the additional 0-24h random delays, i.e., the conversion value count X^d includes all users with $d_i \leq d - t$. In contrast, our empirical analysis in Section 4 does consider these and further empirical observations obtained by interacting with the SKAdNetwork: any user belonging to X^d are disregarded for any future $X^{d'}$ where $d < d'$.

To further ease the presentation of this paper, we define three variables: $\tilde{U}_v^d$ is the set of users with the same conversion value on the day d, $\bar{r}_v^t$ is the average revenue from users in $\tilde{U}_v^d$ within the first t days, and g is the function to attribute the revenues based on the conversion value counts (i.e., g approximates y_α^t). In the rest of the paper, we slightly abuse the notations by leaving out superscript d when it does not play a significant role.

3.2 Privacy Protection

The count of conversion values provides privacy protection in hide-in-the-crowd, as the campaign information does not contain user identifiers. Individual users could still be connected with specific networks and campaigns if the size of conversion value buckets is small. For instance if only user i has a specific conversion value then Table 3 would indeed reveal user i's origin α_i. To overcome this problem, Apple proposed the privacy threshold p, a predefined (and currently unknown) value that provides further protection. On the other hand the documentation does not mention on which level (e.g., country, campaing) the p is enforced. In practice, Apple will not report the count of users in the conversion values where there are less than p users, and instead, those counts will be reported as `null`. Moreover, the users with such a conversion value are not discarded. Instead, the set of conversion values is extended with `null`, i.e., $v \in \{\texttt{null}, 0, 1, \ldots, 63\}$ which aggregates all the users from below the threshold conversion values.

Formally, our interpretation of the privacy mechanism (based on what we experience by interacting with the corresponding post-ATT ecosystem) is defined in Equation 1, where $\mathbb{1}$ is the indicator function. This mechanism is similar to k-anonymity [26], which requires all users to be indistinguishable from at least $k-1$ other users. On the other hand, it does not satisfy that because the condition is not enforced on $\hat{x}_{\texttt{null},\alpha}$.

$$pr_p(X) = \hat{X} = \begin{cases} \hat{x}_{v,\alpha} = \begin{cases} x_{v,\alpha} & \text{if } \sum_U \mathbb{1}(v_i = v) \geq p \\ \texttt{null} & \text{otherwise} \end{cases} \\ \hat{x}_{\texttt{null},\alpha} = \sum_v \mathbb{1}(\hat{x}_{v,\alpha} = \texttt{null}) \cdot x_{v,\alpha} \end{cases} \tag{1}$$

3.3 Revenue Attribution Functions

The revenue attribution function g plays a central role in our research, as we want to approximate the actual campaign-wise revenues y_α^t via the conversion values. This attribution error minimization problem is shown in Equation 2. Although f is not explicit in the formula to be minimized, it defines $\tilde{U}_v^d$ as it contains users with $f(u_i^d \setminus \{\alpha_i\}) = v_i^d$.

$$\min_f \left[\sum_\alpha \left(\sum_v g\left(\tilde{U}_v^d, \hat{x}_{v,\alpha}^d\right) - y_\alpha^t \right)^2 \right] \tag{2}$$

First, instead of focusing on f, we show the optimal g when there is no privacy threshold for conversion values (i.e. when $p < 2$). When $p = 0$ it is meaningless, and when $p = 1$ it only changes the 0 values to `null`, making no real difference between X and $\hat{X}$.

Theorem 1. *Only based on $\tilde{U}$ (i.e., without any prior background knowledge about the distribution of users corresponding to any α) and if $p < 2$ (i.e., we assume $\hat{x}_{v,\alpha}^d = x_{v,\alpha}^d$ so there is no additional privacy protection) then independently of f, the attribution function defined in Equation 3 minimizes Equation 2.*

$$g_\alpha\left(\tilde{U}_v^d, x_{v,\alpha}^d\right) = x_{v,\alpha}^d \cdot \bar{r}_v^t \tag{3}$$

Now we relax our initial condition about p and focus on the case when the privacy threshold is applied (i.e., when $p \geq 2$). The exact privacy-preserving mechanism used by Apple is unknown. Therefore we are using Equation 1 based on empirical available post-back data. The revenue attribution function defined in Equation 3 does not consider the `null` bucket. To account for the `null` bucket, we propose two attribution functions in the form of Equation 4, where $h(\{\hat{x}_{\texttt{null},\alpha}\}|_{\alpha=0}^{\beta})$ should be defined accordingly.

$$g_\alpha\left(\tilde{U}_v^d, x_{v,\alpha}\right) = \begin{cases} \bar{r}_v^t \cdot \hat{x}_{v,\alpha} & \text{if } \hat{x}_{v,\alpha} \neq \texttt{null} \\ \bar{r}_v^t \cdot h(\cdot) \cdot \sum_U \mathbb{1}(v_i = v) & \text{otherwise} \end{cases} \tag{4}$$

*Uniform Revenue Attribution (**U**).* Distributing the revenue uniformly across all possible networks and campaigns, i.e., $h_U(\cdot) = \frac{1}{\beta}$ should fill in Equation 4. This function is used as a pessimistic baseline because it does not use any information from $\hat{X}^d$.

*Null-based Revenue Attribution (**N**).* Distributing the revenue based on the empirical distribution defined by the `null` bucket, i.e., $h_N(\cdot) = \frac{\hat{x}_{\texttt{null},\alpha}}{\sum_\alpha \hat{x}_{\texttt{null},\alpha}}$ should fill in Equation 4. This function is based on the sum of the distribution corresponding to conversion values below the threshold p. Although we have no prior background information about the user distributions within the conversion values, we can still utilize `null` bucket for those below p.

THEOREM 2. *Only based on $\tilde{U}$ (e.g., without any prior background knowledge about the distribution of users corresponding to any α) for any f, the attribution function defined in Equation 4 minimizes Equation 2 where $h(\cdot)$ is a convex combination of $h_U(\cdot)$ and $h_N(\cdot)$.*

The proofs for these theorems can be found in the extended version [6].

4 Experiments

This section introduces several conversion value schemas and shows the corresponding empirical results combined with the introduced revenue attribution functions.

Setup. We experiment using data from a globally launched free-to-play mobile game developed by one of the biggest mobile game developers. To generate the ground truth dataset, we used six months of historical data from cohorts with revenue matured up to 90 days (i.e., players that have played at least 90 days). The dataset includes more than 500K paid players, which constitutes a significant share of players in the time window. The users include organic and 213 paid campaigns across seven networks.

Using historical data allows us to compare the attributed revenue with the actual data from last-click attribution. We calculate the conversion value for each user in the dataset according to the schema we want to evaluate. Because neither the exact privacy threshold nor the level is known, we use Equation 1 with different values of p and with country-level privacy protection. We build the matrix $\hat{X}^d$ eight times separately: six for countries with the largest user bases and two for the rest grouped randomly. Instead of using the registration date of the players, we use weekly cohorts, meaning that the matrix X^d contains the sum of counts of daily conversion values per week starting from Monday. We calculate the error per week and then aggregate the error from different weeks using a weighted average where the weight is the week's revenue, i.e., $e^t = \frac{\sum(r_w^t * e_w^t)}{\sum(r_w^t)}$, where r_w^t is the week's revenue and e_w^t is the error for that week $e_w^t = \frac{|r^t - \hat{r^t}|}{r^t}$. Then, the error is normalized by a hypothetical best-case baseline as explained in Section 4.2. The experiments were implemented in Python and ran on a single machine with 64 vCPU and 512 GB RAM.

4.1 Conversion value schemas

In Section 3.3 we described the revenue attribution function g and showed theoretically which is optimal. Concerning conversion value schema f, we do not pursue this direction. Rather, we assign meaning to each of the six bits of the conversion values by defining three types of bits: T bits used for time (i.e., day), V bits used for revenue, and C bits used for a logical condition (data captured within $\mathcal{U}$, for instance the device is tablet or smartphone, user passed tutorial, or the user reached a certain level). Using these bits, we specify various conversion value schemas.

It is worth mentioning that schemas using data beyond the registration day (i.e., day 0) are challenging in practice because they depend on the user coming back to play within 24 hours and updating the conversion value while the player is using the application. Next, we define five conversion value schemas.

*Day 0 event-based (**EV**).* Using data from $\mathcal{U}_i$, we encode actions taken by the user during the first day of gameplay (i.e., *CCCCCC*), each action mapped to one bit (e.g., finished tutorial as bit 0, reached a certain level as bit 1).

*Rolling Revenue & Rolling Purchase Count (**RR** & **RI**).* Both rolling schemas utilize some bits T for keeping track of the days that have passed from the first opening. The purchases define the remaining bits: **RR** uses bits V for bucketing the actual revenue. In contrast, **RI** uses bits C for bucketing the purchase counts of the user during the observation period, i.e., the first accumulates the total value of purchases while the latter counts how many purchases happened. Users without revenue are assigned to the zero buckets, and those with revenue are distributed uniformly based on their revenue. For example, **D7 RR** is defined as *TTTVVV*, where T bits capture day 0-7 and V bits are based on the current user's revenue. Plots on revenue schemas could not be provided not to disclose confidential business information.

*Uniform distribution (**UD**).* The users are distributed in conversion values at random. This schema is used as a pessimistic baseline because it does not use any information from the user.

*Perfect lifetime value (**PV**).* Using six V bits to bucket users based on the future cumulative revenue of the user. This is a hypothetical schema as it uses data that is not available in practice [19]. For example, **D30 PV** is defined as *VVVVVV*, where bits are based on the user's cumulative revenue until day 30. The schema serves as an optimistic baseline because it places users so that their revenue is close to the conversion value's expected revenue.

4.2 Results

We want to know how the introduced conversion value schemas presented in Section 4.1 perform in revenue attribution. We experimented by backtesting past data before ATT changes. Hence the ground truth is available via the data reported by our Mobile Measurement Partner. This allows us to measure the revenue attribution error of various conversion value schemas. Our results are presented in Table 4 where the prefix in the first column shows the number of bits used for a time, e.g., D1 corresponds to 1 bit (one day). The revenue attribution was calculated every week, so the results are compared using the average error for all the weeks. From Theorem 2 we know that the optimal revenue attribution function is a convex combination of uniform **U** and `null`-based empirical **N**. Because the exact combination is unknown, we consider both separately.

Table 4 shows the errors for attributing the cumulative revenue for 30 days. It shows that the best conversion value schemas use the observed users' revenues. The attribution errors are normalized with the hypothetical best-case **D30 PV** with **U** for every privacy parameter separately (also marked with a box). For example, in Table 4a, when $p = 2$, the conversion value schema **D7 RR** combined with **N** is 4% worse than the error of **D30 PV** with **U**.

As p increases, the **EV** and **UD** schemas' performance gets closer to the rest of the schemas because a high privacy threshold applied to the revenue-based schema sets most of the spending user's conversion value to `null`. As expected, the baseline schema **D30 PV** error is smaller than all others when there is no privacy threshold, and **UD** performs the worst.

Looking at the results for **RR** and **RI** with a low privacy threshold (i.e., $p \leq 2$) in Table 4 we see that using a more extended

	$p = 0$	$p = 2$		$p = 10$		$p = 100$	
f \ g	Eq. 3	**U**	**N**	**U**	**N**	**U**	**N**
D30 PV	$\boxed{0}$	$\boxed{0}$	1	$\boxed{0}$	-2	$\boxed{0}$	28
EV	-42	-25	-25	-2	-8	15	17
D1 RR	-32	-19	-19	-5	-5	18	24
D1 RI	-34	-20	-20	0	-3	17	26
D3 RR	-21	-9	-8	-15	-10	11	35
D3 RI	-25	-12	-11	-7	-17	6	33
D7 RR	-17	-4	-4	-14	-6	-7	25
D7 RI	-21	-8	-7	-7	-6	-6	32
UD	-62	-43	-43	-15	-15	10	10

(a) Attribution benchmark for cumulative *campaign revenue.*

	$p = 0$	$p = 2$		$p = 10$		$p = 100$	
f \ g	Eq. 3	**U**	**N**	**U**	**N**	**U**	**N**
D30 PV	$\boxed{0}$	$\boxed{0}$	2	$\boxed{0}$	7	$\boxed{0}$	46
EV	-32	-15	-15	7	7	23	34
D1 RR	-24	-11	-10	2	5	28	40
D1 RI	-26	-12	-11	8	14	26	43
D3 RR	-15	-2	-1	-11	-1	15	50
D3 RI	-17	-5	-4	-2	2	9	51
D7 RR	-11	1	2	-13	1	-6	44
D7 RI	-13	0	1	-5	6	-6	50
UD	-50	-31	-31	-3	-3	11	11

(b) Attribution benchmark for cumulative *network revenue.*

Table 4: Attribution benchmark for 30 days of cumulative revenue. The error metric is normalized with D30 PV combined with U (also noted with a box). A negative value means worse than baseline and a positive value means better than baseline.

period than the first 24 hours of gameplay reduces the attribution error. Intuitively, it takes some time for players to try the game, and they will start buying once they consider that it is worth it – which rarely happens on the first day of gameplay. **RR** and **RI** work well if enough players are spending during the observed period because it helps separate non-spenders from spenders and then places spenders in buckets based on their spending. However, higher privacy thresholds affect the quality of this schema because observing the players for a few days cannot correctly separate players into different conversion values. Most players will have a conversion value of zero, and those that do not will likely fall below the privacy threshold.

The results suggest that schemas that separate spenders and non-spenders and group users based on their spending are the most helpful for revenue attribution. That is why *rolling* schemas that include bits for carrying the count of days perform much better than the **EV** or **UD** schema.

5 Conclusion

This paper focuses on using Apple's new performance mechanism called conversion values to attribute revenue to advertising campaigns. To the best of our knowledge, our work is the first to formalize and investigate the conversion values for revenue attribution rigorously. We find the optimal revenue attribution function, and through various experiments, we shed light on how different conversion value schemas perform in revenue attribution. Based on empirical evaluation of real-world data, we postulate that the best conversion value schema is the one that relies on revenue and can separate players by their spending.

Limitations. Our work barely scratched the surface of Apple's conversion value schema and focused on the task of attributing revenue using conversion values. A major limitation is that the rules of the privacy threshold are not clear. We used data of free-to-play games where more than 95% of the users do not spend money.

Acknowledgments

The authors thank the various teams at Rovio for the great cross-team collaboration and support during the project. F. Ayala-Gómez was supported by Bandai Namco Mobile. B. Pejó was supported by the ITM of Hungary from the NKFI Fund, financed under TKP2021 and FK21 as part of project no. BME-NVA-02 and 138903, respectively.

References

[1] Adjust. Conversion value management, 2022.
[2] Alphabet. Q3 2021 earnings call, 2021.
[3] Apple. Skadnetwork. Apple Developer Documentation, 2020.
[4] Apple. updateconversionvalue(_:). Apple Developer Documentation, 2021.
[5] Apple. What's new with skadnetwork, 2022.
[6] Frederick Ayala-Gomez, Ismo Horppu, Erlin Gulbenkoglu, Vesa Siivola, and Balázs Pejó. Revenue attribution on ios 14 using conversion values in f2p games. *arXiv preprint arXiv:2102.08458*, 2021.
[7] European Commission. e-privacy. *Flash Eurobarometer*, 443, 2016.
[8] Brian Dalessandro, Claudia Perlich, Ori Stitelman, and Foster Provost. Causally motivated attribution for online advertising. In *Proceedings of the sixth international workshop on data mining for online advertising and internet economy*, 2012.
[9] Bolin Ding, Janardhan Kulkarni, and Sergey Yekhanin. Collecting telemetry data privately. In *Advances in Neural Information Processing Systems*, 2017.
[10] Cynthia Dwork. Differential privacy. In *Proceedings of the 33rd international conference on Automata, Languages and Programming*. ACM, 2006.
[11] Úlfar Erlingsson, Vasyl Pihur, and Aleksandra Korolova. RAPPOR: Randomized aggregatable privacy-preserving ordinal response. In *Proceedings of the 2014 ACM SIGSAC conference on computer and communications security*. ACM, 2014.
[12] Facebook. Third quarter 2021 results conference call, 2021.
[13] Peter S Fader, Bruce GS Hardie, and Ka Lok Lee. "counting your customers" the easy way: An alternative to the pareto/nbd model. *Marketing science*, 2005.
[14] Rafi Goldberg. Most americans continue to have privacy and security concerns, ntia survey finds, 2018.
[15] Vaibhav Gupta. Introduction to singular's skadnetwork solution, 2022.
[16] Gabriel Hughes and Damien Allison. Multiple attribution models with return on ad spend, 2014. US Patent 8,788,339.
[17] Garrett Johnson, Julian Runge, and Eric Seufert. Privacy-centric digital advertising: Implications for research. *SSRN 3947290*, 2021.
[18] Shumel Lais. Skadnetwork 3.0 and ios 15: Everything you need to know to prepare your user acquisition, 2021.
[19] Edward C Malthouse and Robert C Blattberg. Can we predict customer lifetime value? *Journal of interactive marketing*, 2005.
[20] Balázs Pejó and Damien Desfontaines. *Guide to Differential Privacy Modifications: A Taxonomy of Variants and Extensions*. Springer Nature, 2022.
[21] David C Schmittlein, Donald G Morrison, and Richard Colombo. Counting your customers: Who-are they and what will they do next? *Management science*, 1987.
[22] Eric Benjamin Seufert. Apple killed the idfa: A comprehensive guide to the future of mobile marketing, 2020.
[23] Claude E Shannon. Communication theory of secrecy systems. *The Bell system technical journal*, 1949.
[24] Singular. Skadnetwork model configuration faq, 2022.
[25] SNAP. Q3 2021 transcript, 2021.
[26] Latanya Sweeney. k-anonymity: A model for protecting privacy. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 2002.
[27] Jun Tang, Aleksandra Korolova, Xiaolong Bai, Xueqiang Wang, and Xiaofeng Wang. Privacy loss in apple's implementation of differential privacy on macos 10.12. *arXiv preprint arXiv:1709.02753*, 2017.
[28] Differential Privacy Team. Learning with privacy at scale. Apple, 2017.
[29] Twitter. Q3 2021 earnings report, 2021.
[30] Vungle. What is probabilistic attribution?, 2021.