Profit-Maximized User Acquisition Payback Windows

Russell Ovans, East Side Games (ESG)

Abstract How much should you spend on user acquisition (UA)? The effectiveness of paid UA campaigns is measured by their day n return on ad spend (Dn ROAS), defined as the ratio of average customer lifetime value by day n after install to the average cost per customer acquired; i.e., Dn ROAS $= \text{LTV}_n/\text{CPI}$. A campaign's profit per install is the terminal LTV of acquired users minus their CPI. However, the typical UA manager workflow is to ignore profit and instead target a specific break-even day N where DN ROAS $\geq 100\%$. N is also known as the payback window. Budget is constrained by this target, which effectively mandates that a campaign's average CPI remain at or below expected LTV_N. Additional spend is unlocked only by pushing the payback window further out and accepting an increase in average CPI. But how do we determine a profit-maximizing payback window (i.e., what value of N should be targeted) when CPI is correlated with daily ad spend? The diminishing marginal utility (profit) of each additional daily install becomes zero when its incremental CPI is equal to its LTV. Any expenditure beyond that point will increase top-line revenue, but diminish profit and cash deplete your company, even though average CPI remains below LTV. Profit is maximized when the daily ad spend results in a *marginal* CPI equal to terminal LTV, not when *average* CPI equals some arbitrary LTV_N.

The Problem Statement

How much should you spend on UA for a mature game with a well-understood monetization curve? One simple approach is to spend a maximum proportion (e.g., 20%) of a game's current revenue: if a game generated $10,000 in revenue yesterday, then spend $2,000 on UA today. A different approach is to divorce spend from current revenue, and instead operate under the constraint of attaining a 100% return on ad spend (ROAS) by a prescribed payback window of N days after expenditure (i.e., the install date of the users attributed to the campaign). In this case, there is no budget: UA managers increase daily ad spend so long as CPI $\leq \text{LTV}_N$, regardless of the actual CPI and total expenditure; spend is reduced if predicted days to break even exceed N. This common approach relies on install attribution and a predictive model of LTV based on early cohort behaviour; i.e., LTV_N is predicted based on the ongoing monetization of a cohort of installs whose average acquisition cost is CPI.

At ESG, we typically employ this latter approach: each game's UA budget is uncapped so long as campaigns are expected to back out (i.e., attain 100% ROAS) within $N = 90$ days. But is D90 the correct user acquisition payback window? Can the studio generate more revenue and more profit by extending that payback window to, say, D180? Barring any exceptional circumstances such as a new game launch, the growth team should be maximizing the *profit* of every title in their portfolio of games, where profit is defined as terminal net revenue minus UA expenditure. However, employing arbitrary industry-standard terminal days N like 90 and 180 does not guarantee you optimize long-term profit. In particular, extending the payback window can actually decrease your long-term profit.

Why? Well, spending until D90 ROAS is 100% results in a cohort average CPI of LTV_{90}. If you increase spend right up until CPI $= \text{LTV}_{180}$, then any profit must arise beyond D180. However, the jump from 90 to 180 days to break even can wreak havoc on your profitability. Refer to Figure 1.

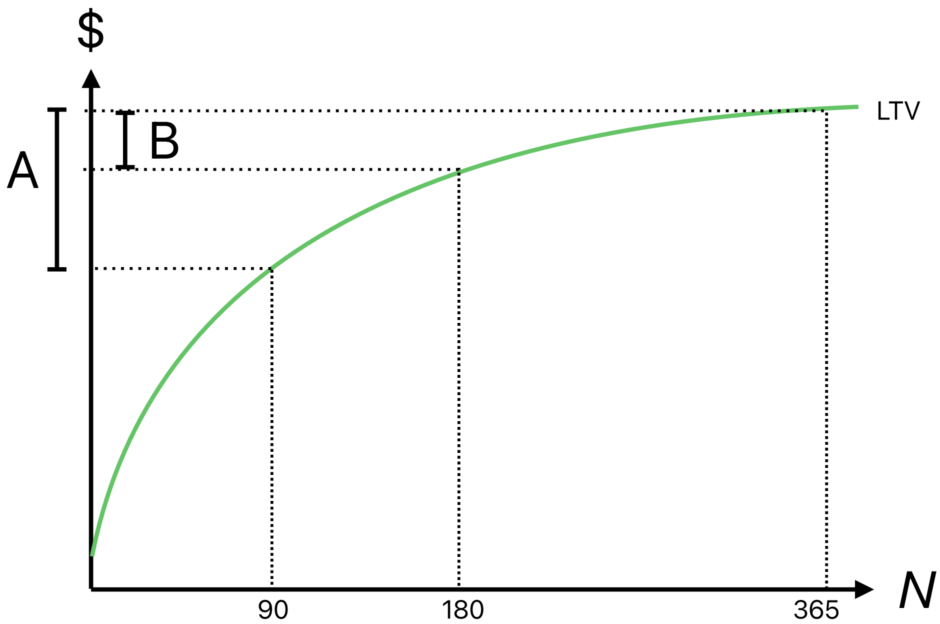

Figure 1. As the UA payback window *N* increases, the profit of each install diminishes.

With a power-function LTV curve (in green), the profit from each user acquired at a CPI of LTV180 (indicated by B) is much smaller than for users acquired at a cost equal to LTV90 (indicated by A). So while increasing the payback window from 90 to 180 days unlocks UA budget and enables the acquisition of additional installs, this will lead to less long-term profit unless you can buy `A/B` times more installs at an average cost of LTV180.

An example can help visualize what is going on here. Assume a game's LTV90 is \$5.00, LTV180 is \$7.50, and terminal LTV (which occurs at D365) is \$10.00. The profit A for users acquired at an average CPI of \$5.00 is \$5.00 each, whereas the profit B for users acquired at an average CPI of \$7.50 is only \$2.50 each. `A/B=2`, meaning a campaign must drive twice as many installs at an average cost of \$7.50 to generate the same profit as a campaign whose average CPI is only \$5.00. Refer to the following table.

Payback window *N*	daily spend	average CPI	daily installs	D365 revenue	D365 profit	profit margin
90	\$ 2,000	\$ 5.00	400	\$ 4,000	\$ 2,000	50%
180	\$ 6,000	\$ 7.50	800	\$ 8,000	\$ 2,000	25%

If a campaign with a payback window of 90 days is driving 400 installs/day, its D365 profit margin is a healthy 50%. To match its \$2,000 in D365 profit, a campaign with a 180 day payback window and an average CPI of \$7.50 must spend 3*X* per day. But if CPI increases as a function of ad spend, is it even possible to find 800 daily installs at an average CPI of \$7.50?

CPI Increases as a Function of Ad Spend

Mobile app-install (MAI) campaigns with a fixed cap on CPI are guaranteed to drive installs at a prescribed average cost (or less), but increasing daily budget while holding CPI steady does not necessarily result in more installs. To unlock the increase in budget, the maximum CPI the UA manager will tolerate also must be increased, leading to a longer payback window. However, MAI users are often low-quality (i.e., poor monetizers), and thus not guaranteed to back out within N days simply because their average CPI is below historical LTV_N.

For this reason, UA managers often employ app-event optimized (AEO) campaigns with a fixed daily budget, D*n* ROAS targets, and no cap on CPI. AEO campaigns that are algorithmically optimized by the networks automatically bid higher for known high-quality users. Whatever the AEO campaign budget, the network is incentivized to spend it, which entails bidding higher and higher in order to win more auctions over time. (Publishers pay for installs, not impressions.) Therefore, like its MAI counterparts, increasing campaign daily budgets results in more installs, but at an increasing CPI. See Figure 2, where each point represents a single day over an AEO campaign's history.

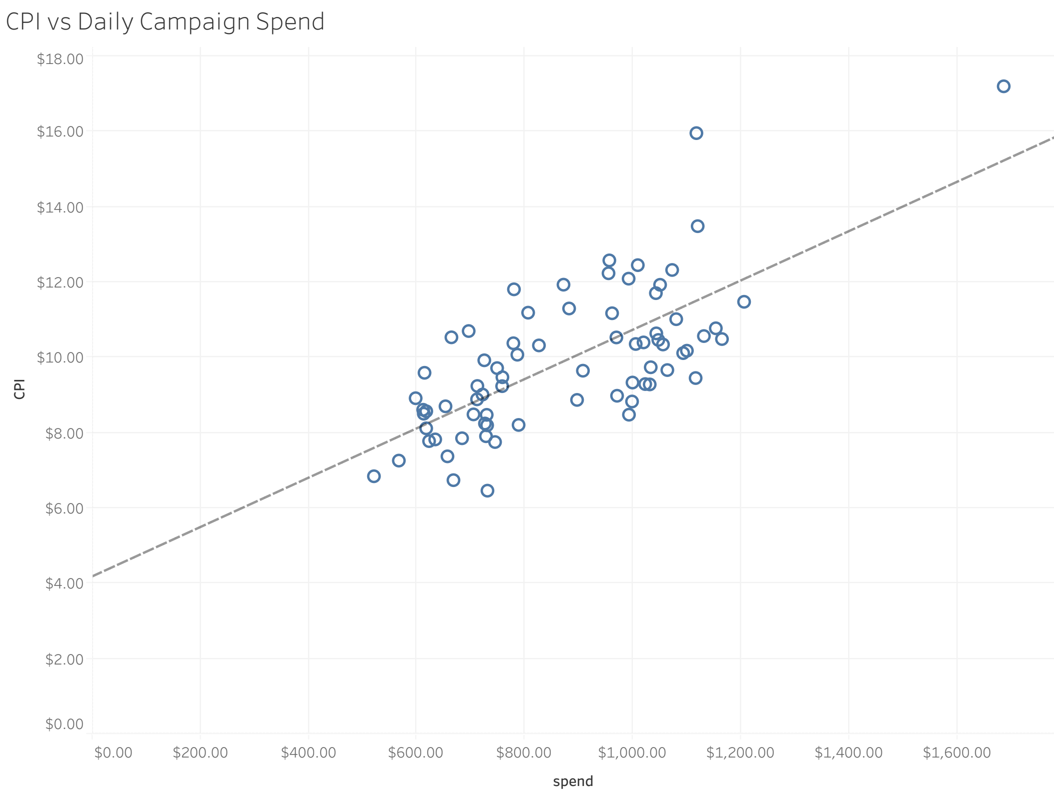

Figure 2. Historical daily spend versus CPI indicates that CPI is positively correlated to daily budget. The linear regression fit to this campaign data has a y-intercept of $4.00 and a slope of 0.0005.

This correlation works both ways, depending on the campaign type. For MAI, increasing the maximum CPI will increase daily spend. For AEO, increasing the daily spend will increase average CPI. In either case, any increase in spend must come at the expense of a longer payback window.

As indicated earlier, *modus operandi* at ESG is to increase daily ad spend on each active campaign that is expected to back out by D90. But is D90 the right target day to break even? Can the studio generate more installs and more revenue by extending that payback window beyond D90 and increasing daily ad spend? Yes, but if CPI increases with spend while expected LTV remains constant, what impact does this have on long-term profitability?

When is Profit Maximized?

Typically, there are diminishing returns on additional installs as daily ad spend budgets for campaigns are increased. The relationship between ad spend and CPI is like how tax brackets work; the cost of the first x installs per day is less than the cost of the next x installs, leading to a situation where the profit on each additional daily install decreases. The **diminishing marginal utility** (profit) of each additional install becomes zero when its incremental CPI is equal to its LTV. Any expenditure beyond that point will increase top-line revenue, but diminish profit, and cash deplete your company, *even though average CPI remains below LTV*.

Profit is maximized when the *marginal* CPI equals the LTV, which is not necessarily when *average* CPI = LTV_{90} or *average* CPI = LTV_{180}. The difference is subtle, but illustrated with the following example. Once again, assume a mature game's expected LTV90 is $5.00, LTV180 is $7.50, LTV365 is $10.00, and LTV beyond D365 is virtually non-existent; i.e., LTV = $10.00. Our objective is to solve for daily *spend* that maximizes long-term *profit* as defined by the following formula:

$$profit = revenue - spend$$

Total long-term revenue is dictated by the number of users acquired (*spend*/CPI) multiplied by their LTV:

$$revenue = \frac{spend}{\text{CPI}} * \text{LTV}$$

Therefore,

$$profit = \frac{spend}{\text{CPI}} * \text{LTV} - spend$$

Assuming a linear relationship between ad spend and CPI, like in Figure 2, then:

$$\text{CPI} = m * spend + b$$

Therefore,

$$profit = \frac{spend * \text{LTV}}{m * spend + b} - spend$$

Assume we have an active ad campaign that is returning an average CPI of $4.00 when daily spend is $500. Beyond that level of spend, assume CPI grows linearly according to Figure 2; namely, $0.0005 * spend + \$4.00$. Profit is therefore governed by the following function:

$$profit(spend) = \frac{spend * \$10}{0.0005(spend) + \$4} - spend$$

The maximum occurs at the inflection point on a graph of *profit* as a function of daily *spend*. Refer to Figure 3.

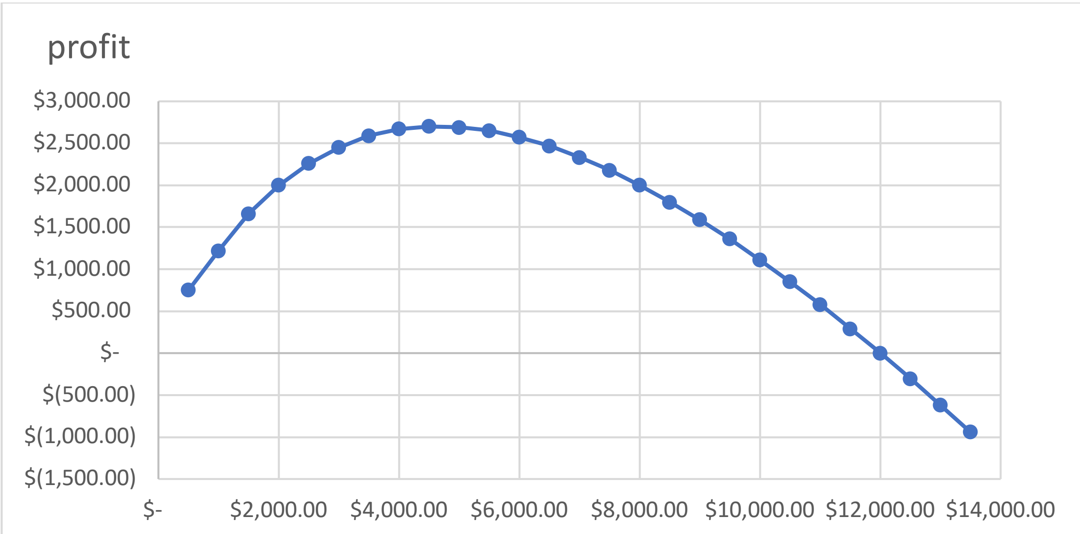

Figure 3. Profit (*y*-axis) as a function of daily ad spend (*x*-axis) for a campaign with a linear relationship between budget and CPI. Profit is maximized at the inflection point on the curve.

In this case, profit when $\text{LTV} = \$10$ and $\text{CPI} = 0.0005 * spend + \$4$ is maximized at an expenditure of \$4,500/day, where average CPI comes in at \$6.25. The optimal payback window is the day $N*$ when $\text{LTV}_{N^*} =$ \$6.25, which is easily determined from inspection of the game's historic LTV curve. Targeting a 100% ROAS for any terminal day $N, N > N^*$, results in less profit. Why? Because each of those incremental installs cost more than \$10.00!

Back to our original question: should we increase the payback window from 90 to 180 days? Since LTV90 is less than \$6.25, spending only up until average $\text{CPI} = \$5$ is not going to optimize profit. Ad spend should definitely be increased if campaigns are predicted to break even by D90.

The following table – and Figure 4 – summarizes the relative performance of the three payback windows. Also refer to Figure 4.

N	**daily spend**	**average CPI**	**daily installs**	**D365 revenue**	**D365 profit**	**profit margin**
90	\$ 2,000.00	\$ 5.00	400	\$ 4,000.00	\$ 2,000.00	50%
N*	\$ 4,500.00	\$ 6.25	720	\$ 7,200.00	\$ 2,700.00	37.5%
180	\$ 7,000.00	\$ 7.50	933	\$ 9,333.33	\$ 2,333.33	25%

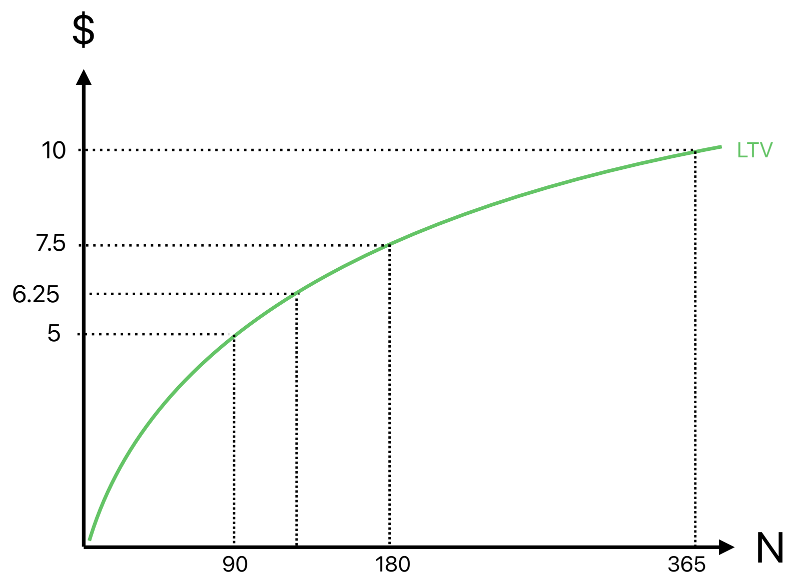

Figure 4. To increase the payback window from 90 to 180 days, average CPI increases from $5.00 to $7.50. The optimal payback window from a profit-maximization perspective is about 120 days, where average CPI is $6.25.

At a certain level of daily ad spend ($4,500.00), the last install costs about $10.00, which is expected LTV. This is the profit-maximizing budget. Even though the average CPI over all daily installs is $6.25, in order to win enough auctions to spend the daily budget, the last installs are significantly more expensive.

The situation becomes worse when D180 is the payback window: since LTV180 is $7.50, using that target results in significantly more spend ($7,000/day in total) and less D365 profit ($2,333 vs. $2,700) compared with a $6.25 CPI. Note that while an additional $333.33 of profit is expected if we extend the payback window of this campaign from 90 to 180 days, this requires an additional $5,000/day of ad spend. $333.33 profit on a $5,000 investment is *not* a great return, especially given the inherent risk of predictive models based on historical performance. Do you have the stomach for that?

Caveat

The optimum day N^* is campaign specific since it depends on how CPI behaves as a function of daily spend and the LTV of the specific users acquired by the campaign. Therefore, it is difficult to know while a campaign is running. Only in hindsight are we able to look at the data and say exactly what the optimal payback window should have been.