
A/B Testing at Vungle

Andrew Kritzer and Hammond Guerin stared at the screen and then at each other. It was June 30, 2014—six weeks since they had graduated from the Darden School of Business. The ad-serving algorithm Kritzer and Guerin had spent six months developing for Vungle, a mobile advertising company, seemed to be outperforming the company's current algorithm. But they did not want to start celebrating too soon. Could their algorithm really deliver the type of improvement they had promised Vungle's CEO? Would install rates of advertised apps really increase? Would Vungle see an increase in ad-serving efficiency as a result?

Neither Kritzer nor Guerin could afford for the algorithm to disappoint. Now that he had graduated, Kritzer was headed to LinkedIn, having left a legend among MBA students for his appreciation of data science, tech, and media and raising expectations for what Darden students knew and could learn about data science, analytics, and the ever-growing world of big data. His work on the Vungle project during his second year had received a lot of attention, and he was looking forward to having the results support the effort.

Guerin's data science capabilities were also legendary among his MBA peers. He won every school forecasting competition, and his data mining algorithms even beat those of the professional consultants who did classroom visits. Late in his second year, Guerin decided to turn down a generous offer from a well-known consulting firm in favor of an offer from Vungle for an annual salary of \$100,000 and stock options to serve as the head of Vungle's brand new data science team out in the company's San Francisco headquarters. The job was a dream for the computer scientist turned MBA. He and his wife were already house hunting in the Bay Area, looking for the right place to raise their baby daughter.

Company Overview

Vungle was an in-app video advertising company. With 70 employees and \$25.5 million from three rounds of investments, Vungle was routinely listed as one of the most promising start-ups operating in Silicon Valley.¹ The three-year-old company offered a platform that embedded video ads in mobile apps to encourage users to download and install additional apps. It was estimated that more than 100 million people saw an advertisement enabled by Vungle each month.²

¹ Anthony Ha, "In-App Video Ad Startup Vungle Raises \$17M More," *TechCrunch*, February 6, 2014, <http://techcrunch.com/2014/02/06/vungle-series-b> (accessed Aug. 23, 2014).

² Steven Loch, "In-App Video Advertising Platform Vungle Raises \$17M," *VatorNews*, February 6, 2014, <http://vator.tv/news/2014-02-06-in-app-video-advertising-platform-vungle-raises-17m> (accessed Aug. 23, 2014).

Vungle was founded in 2011 by two young entrepreneurs from the United Kingdom, Zain Jaffer and Jack Smith, during their graduate studies at University College London. Initially a video ad production firm, Vungle's expenses in its first year were running too high and revenue was not reaching the founders' expectations. Late in 2011, Jaffer borrowed funds from his then girlfriend (and future wife) and his business professor, Bert De Reyck. Each invested \$15,000 and the company remained afloat.³

The turning point for Vungle came in 2012, when the two founders creatively used their own video production technology to get the attention of the San Francisco-based start-up incubator AngelPad. In doing so, they beat 2,000 applicants for the final slot in the incubator program. This opportunity provided Vungle \$120,000 in seed funding. Jaffer moved to San Francisco to serve as the firm's CEO and remained in that position. He was profiled in a "35 Under 35" list by *Inc.* magazine in 2014.⁴

The Mobile Advertising Ecosystem: Market, Operations, and Pricing

In 2013, the average U.S. consumer spent two hours and 42 minutes on mobile devices per day; 86% of that was spent in apps, the clear dominant form of mobile usage.⁵ The growth in the mobile market and the extensive time spent in apps introduced a new advertising channel. According to the Mobile Marketing Association, 75% of ads served to mobile consumers in 2013 were served while they were using apps.⁶ Mobile in-app ads experienced a 60% annual growth in 2013 and were expected to surpass PC online ad revenues by 2017.⁷

By 2014, in-app video advertising was replacing mobile banner ads—the latter offered a lower-quality user experience and were typically clicked on accidentally. The in-app video ads were typically 15 seconds long and promoted a new app or product. Apple's iOS system accounted for 80% of ads being served. Video ads peaked during prime-time TV hours.⁸

Four parties participated in the in-app mobile advertisement channel. The user of the mobile device (*user*), the owner of the app being used (*publisher*), the sponsor of the video ad the user was exposed to (*advertiser*), and the platform that matched the choice of ad to a specific user (e.g., *Vungle*). In the mobile advertising domain, supply was considered to be the slots available for showing ads, and demand consisted of the advertisers willing to buy the supply by placing ads.

When the user launched an app, his or her device would send a request to Vungle for an ad. For instance, suppose user Chris was playing Sonic Dash by the publisher Sega. Vungle's platform would then determine the best ad to serve to Chris while he played Sonic Dash. Assume Vungle decided to serve Chris an ad for the game Hay Day (the advertiser; see **Exhibit 1** for a schematic of this process). Assuming Chris was still playing Sonic Dash when the ad was served, then Chris would see the video for Hay Day. If Chris was interested in learning more about Hay Day, he would click on the ad and be redirected to the app store. Chris might then decide to install Hay Day.

³ Laura Montini, "Creating Ads That Blend In to Stand Out," *Inc.*, June 24, 2014, <http://www.inc.com/laura-montini/35-under-35-using-ads-that-blend-in-to-stand-out.html> (accessed Jan. 15, 2015).

⁴ Donna Fenn, "Generation Why Not: Meet the 35 Under 35, Class of 2014," *Inc.*, July/August 2014, <http://www.inc.com/donna-fenn/35-under-35-2014.html?cid=readmore> (accessed Jan. 15, 2015).

⁵ Ewan Spence, "The Mobile Browser Is Dead, Long Live the App," *Forbes*, April 2, 2014, <http://www.forbes.com/sites/ewanspence/2014/04/02/the-mobile-browser-is-dead-long-live-the-app> (accessed Aug. 23, 2014).

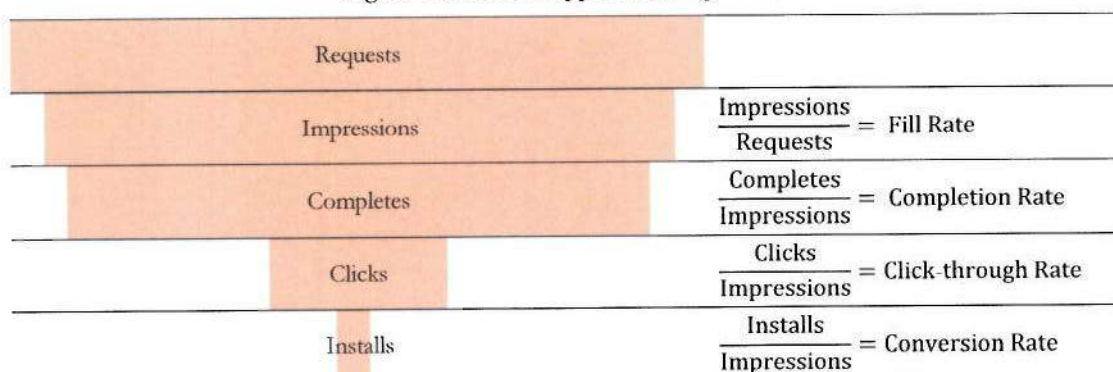
⁶ Spence.

⁷ Dean Takahashi, "Mobile In-App Ad Revenues Will Surpass PC Online Display Advertising by 2017," *VentureBeat*, March 26, 2014, <http://venturebeat.com/2014/03/26/mobile-in-app-ad-revenues-will-surpass-pc-online-display-advertising-by-2017> (accessed Aug. 23, 2014).

⁸ Christopher Heine, "75% of Mobile Video Ads Happen In-App," *AdWeek*, April 24, 2014, <http://www.adweek.com/news/technology/75-mobile-video-ads-happen-app-157217> (accessed Jan. 15, 2015).

In most cases, payment was made by the advertiser upon installation. Publishers typically received 60% of the revenues and the ad provider the remaining 40%. See **Figure 1** for the conversion funnel depicting how an install is achieved. Of all ad *requests*, most were served and became *impressions*. When at least 80% of a video ad was watched, it was considered *complete*. When the user clicked on the ad to get more information, it was counted as a *click*. The process could then result in an *install*.

Figure 1. Mobile in-app advertising funnel.



Source: Created by case writer.

Ads were monetized at all different points along the funnel—whether CPI (cost per install), CPC (cost per click), CPCV (cost per completed view), or CPM (cost per 1,000 views). The vast majority of ads were CPI.

On a typical day, using its current ad-serving algorithm, Vungle experienced a 98% fill rate, 88% completion rate, 5% click-through rate, and 0.5% conversion rate. The funnel for Vungle narrowed substantially at the end. Small improvements in the click-through or conversion rates could have a large effect on Vungle's revenue. The effectiveness of an app-promotion campaign and the success of the serving platform were typically measured by eRPM, or effective revenue (for both publisher and Vungle) per 1,000 impressions,⁹ which could vary from \$2 to as high as \$7 per campaign.

A/B Testing and the Data Science Project

Kritzer and Guerin were tasked with developing an ad-serving learning algorithm. Their data science approach would use historical information about users, publishers, and install rates to determine which ad campaign to serve in order to increase the chance of a conversion and, more specifically, eRPM. If the system proved successful, implementing it would require regular updates to the model by a data scientist, most likely Guerin himself.

Jaffer consulted with Vungle's chief technology officer, Wayne Chan, on how best to test the developed algorithm. Chan planned to test the developed method in parallel with the existing Vungle algorithm. As was typical in such experiments, the two conditions, A (Vungle's existing algorithm) and B (the data science approach) would be evaluated in parallel on randomly assigned users. Since Kritzer and Guerin's algorithm was new and unproven, Chan's team thought it would make sense to direct only 1/16th of the users to the B

⁹ See <http://www.amobcc.com/dictionary/#e>.

condition. The other randomly assigned 15/16ths of users would receive an ad based on the existing algorithm (i.e., the A condition).

Users were assigned to the A or B algorithm using a process called MD5 hashing. An MD5 hash transforms each user ID into a unique 32-character hexadecimal string. Each character of the hexadecimal string could be 0–9, A, B, C, D, E, or F—16 options in total. Each character occurred with equal likelihood, making it simple for Vungle to direct traffic in 1/16th increments using a logic statement (assuming that the original string was random).

The parallel run of the two algorithms began on June 1, 2014. Jaffer was excited to see if B would outperform A and, if so, what the financial benefits would be. He also wondered how long Chan would have to wait to declare a winner. Would a few days be enough time? Or would he need to wait longer? After two weeks, B was looking pretty good. Its daily eRPM was on average \$0.131 higher than A's. Would this translate into annual revenues worthy of the necessary data science investment? **Exhibits 2 and 3** provide the daily results of the A/B test.¹⁰

Thinking about his new role at Vungle, Guerin was curious to see how the superior condition would be chosen. How would one conclude that B was better than A? If he could be confident about such a conclusion, he would be able to develop a robust testing platform for many future experiments.

¹⁰ All numbers in the exhibits are disguised and serve illustrative purposes only.

Exhibit 1
A/B Testing at Vungle
Schematic of the Vungle Platform Role



Source: Company document; used with permission.

Exhibit 2

A/B Testing at Vungle

Data from Vungle A Test Condition

Date	Impressions	Completes	Clicks	Installs	eRPM
1-Jun-14	6,777,407	5,978,434	345,309	31,119	3.327
2-Jun-14	6,004,310	5,331,727	299,732	24,601	2.943
3-Jun-14	5,832,627	5,193,549	291,384	24,220	3.025
4-Jun-14	5,875,702	5,227,917	295,099	23,382	2.985
5-Jun-14	6,843,405	6,111,378	339,529	27,725	3.076
6-Jun-14	7,790,350	6,981,471	392,987	31,820	3.137
7-Jun-14	8,643,430	7,733,750	444,682	38,119	3.322
8-Jun-14	8,929,848	7,993,169	449,680	38,260	3.269
9-Jun-14	8,075,571	7,259,148	392,829	32,825	3.153
10-Jun-14	7,726,694	6,941,293	382,769	31,609	3.237
11-Jun-14	7,781,497	6,999,630	389,369	31,683	3.199
12-Jun-14	7,770,595	6,984,082	391,254	30,985	3.206
13-Jun-14	7,916,282	7,091,841	407,582	31,679	3.246
14-Jun-14	8,724,061	7,782,877	459,952	36,773	3.482
15-Jun-14	9,027,910	8,075,018	465,869	37,701	3.467
16-Jun-14	7,957,999	7,149,399	395,612	31,098	3.245
17-Jun-14	8,102,155	7,283,722	404,716	31,359	3.315
18-Jun-14	8,043,855	7,229,427	407,014	32,414	3.460
19-Jun-14	8,073,992	7,226,473	403,193	31,665	3.583
20-Jun-14	8,085,480	7,224,975	406,766	30,473	3.479
21-Jun-14	8,760,745	7,825,166	454,646	33,178	3.475
22-Jun-14	8,884,803	7,937,481	453,647	33,543	3.459
23-Jun-14	8,040,402	7,182,500	401,226	28,864	3.337
24-Jun-14	7,882,136	7,013,876	389,975	30,302	3.326
25-Jun-14	7,782,617	6,932,529	385,477	30,369	3.367
26-Jun-14	7,734,447	6,887,125	388,935	30,920	3.530
27-Jun-14	7,891,063	7,025,318	409,449	31,689	3.672
28-Jun-14	8,460,726	7,487,623	457,487	34,664	3.830
29-Jun-14	8,849,803	7,785,905	478,901	36,467	3.777
30-Jun-14	8,189,490	7,233,880	411,884	32,160	3.484

Source: Created by case writer.

Exhibit 3

A/B Testing at Vungle

Data from Vungle B Test Condition

Date	Impressions	Completes	Clicks	Installs	eRPM
1-Jun-14	569,044	499,235	28,035	2,111	2.953
2-Jun-14	505,963	447,695	24,621	1,713	2.587
3-Jun-14	492,804	437,495	24,070	1,705	2.755
4-Jun-14	498,772	442,791	25,023	1,801	3.004
5-Jun-14	491,463	436,858	24,337	1,875	3.243
6-Jun-14	509,657	454,702	25,223	1,932	3.430
7-Jun-14	564,247	502,016	28,127	2,221	3.438
8-Jun-14	575,302	512,228	28,200	2,203	3.455
9-Jun-14	523,689	469,082	25,075	1,950	3.272
10-Jun-14	504,636	452,753	24,414	1,914	3.394
11-Jun-14	506,060	454,773	24,637	1,839	3.366
12-Jun-14	505,083	452,687	24,879	1,812	3.321
13-Jun-14	513,106	458,354	26,018	1,893	3.488
14-Jun-14	562,772	499,196	29,088	2,076	3.525
15-Jun-14	586,702	522,522	29,163	2,097	3.341
16-Jun-14	516,148	462,646	24,635	1,805	3.297
17-Jun-14	526,671	471,763	25,325	1,786	3.333
18-Jun-14	526,713	471,137	25,761	1,912	3.604
19-Jun-14	531,452	472,466	25,361	1,740	3.847
20-Jun-14	420,187	373,085	20,629	1,360	3.887
21-Jun-14	548,116	485,150	27,480	1,668	3.694
22-Jun-14	581,785	515,575	28,701	1,816	3.636
23-Jun-14	525,631	466,427	25,462	1,618	3.602
24-Jun-14	517,748	455,814	24,808	1,715	3.418
25-Jun-14	511,505	451,388	24,894	1,725	3.408
26-Jun-14	508,097	448,333	25,111	1,773	3.722
27-Jun-14	518,004	457,335	25,832	1,852	3.939
28-Jun-14	562,854	494,686	28,491	2,041	4.073
29-Jun-14	583,732	510,194	29,483	2,168	4.051
30-Jun-14	537,433	470,054	26,669	1,910	3.687

Source: Created by case writer.