Aampe's no-nonsense guide to push notification timing and frequency, and its direct impact on your revenue $$$

IT'S ABOUT D@#& TIME

aampe

Hey! We're Aampe! 👋

We're a team of *seasoned data scientists* (plus one marketer), and we're sick and tired of reading eBook after eBook and guide after guide telling you that the best time to send push notifications is on Tuesday between 12pm and 2pm, as if all of your users magically fall into nice clean buckets.

Yeah, we know it's easier to pretend that the world is really that simple, but **this guide is for the people who want to see the actual data** we've collected from sending millions upon millions of push notifications, so you can draw your own conclusions about message timing and frequency and understand how much a "12pm to 2pm" mentality could be hurting your business.
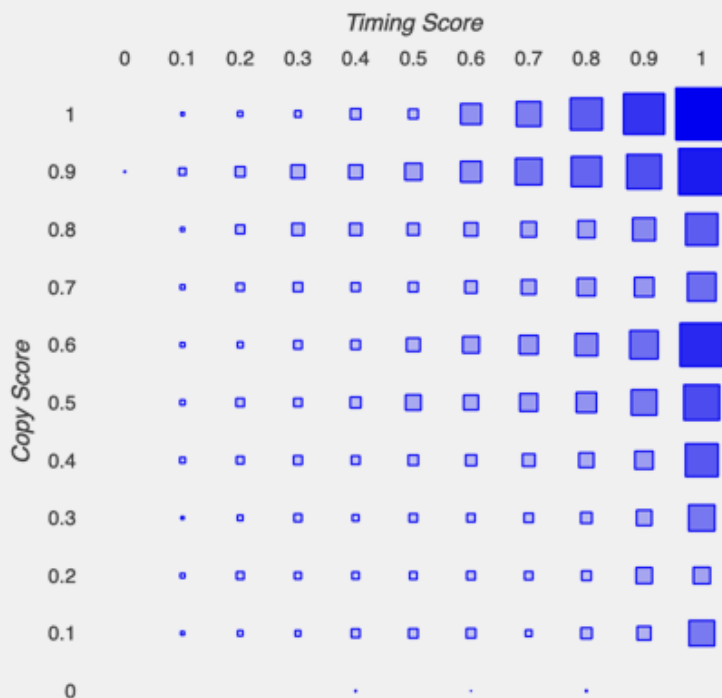
How important is push notification timing, really?

It's important enough to write a book about it, that's for sure.

In fact, we've found that **timing** can have the **single, largest effect** on the success of your push notifications.

Don't believe us?

Take a look at the evidence.

(Wondering how to read this graph? No problem, we'll break it down on the next page.)

Here's how to read this data—

Ok, before we can decipher these graphs, we first need to talk about **personalization scores**:

How Personalization Scores Work

When we message users, we do it through AI infrastructure that estimates the probability that they will respond.

The bottom line: The higher the score, the more likely this individual user will respond to a message at that specific time.

We use this model to generate two different estimates:

- The user's **baseline probability to act** (the probability that the user is going to engage in the desired behavior anyway, whether we message them or not), and

- The user's **target probability to respond** (the probability that the user is going to respond if we message them in a particular way (at a particular time, with particular copy, etc.).

Note: This means each user has only one baseline probability, but has as many target probabilities as there are many choices we can make about a message.

We then compare those two probabilities to calculate a **personalization score** – which represents the degree to which messaging the user in a particular way will increase the probability of response beyond what we'd expect from their individual baseline – and we calculate these scores for timing and several aspects of copy (tone, offer, value proposition, etc.).

So, for example, if you're planning on sending a message at 2pm, each user will get a personalization score that tells you the probability that they'll respond if you send them that message specifically at 2pm.

* If you want a more in-depth introduction to how personalization scores work, [see here](#).

For every timing choice, we start each user with a personalization score of 0.5 – which means they have a 50/50 chance of responding if we send them a message at that particular time – but, as we get more information by experimenting with lots of different times, the score could go as high as 0.99 (meaning the user is very likely to respond), or as low as 0.01 (meaning the user is very unlikely to respond).

Ok, back to our graph...

This graph is a breakdown of personalization scores where we're managing an app's messaging, and each box represents the click-through rate for messages sent on the basis of different personalization scores.

So, those big, blue boxes in the upper right corner mean, predictably, that users are very likely to respond when the personalization scores for both timing and copy are vanishingly close to a perfect '1.0'.

Also notice that there is no visible box in the lower left corner, where our copy and timing personalization scores are a poor '0.0'.
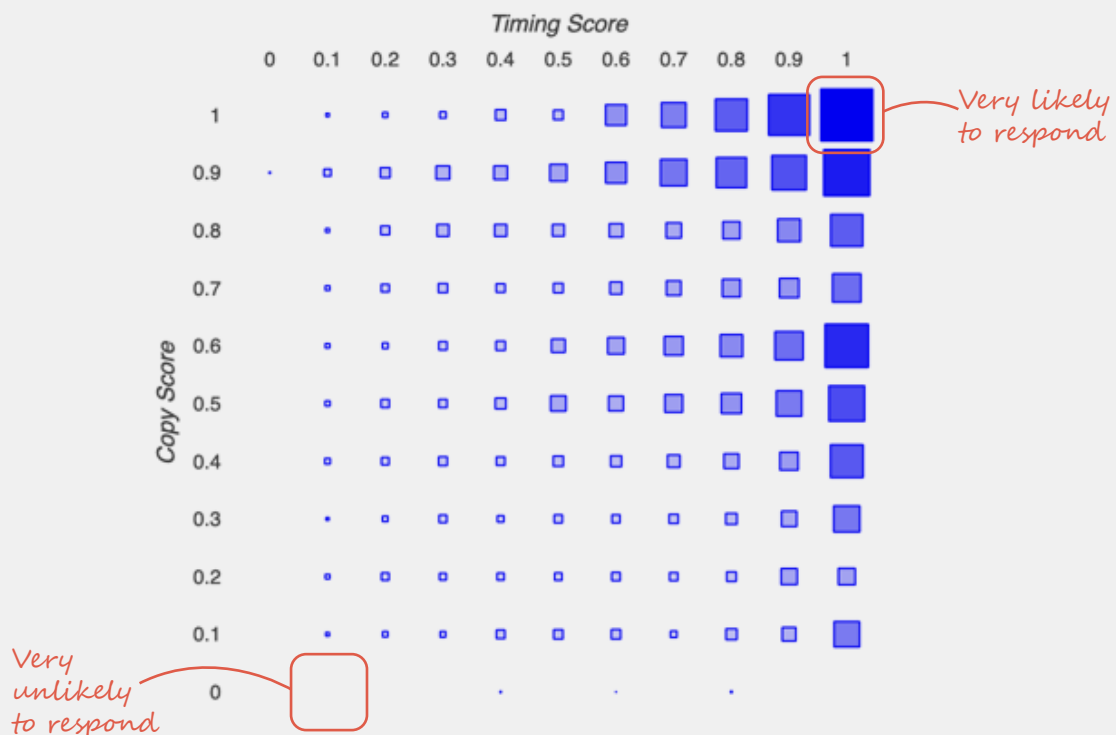
Now that's all to be expected.

What's more interesting is what happens when **timing is bad and copy is good** (the upper left portion of the graph), and when **copy is bad, but timing is good** (the lower right portion of the graph).

…and, in these cases, you can clearly see that there are much more prominent boxes (e.g. Higher Click-Thru Rates) in the lower right **(bad copy/good timing)** vs. the upper left (good copy/bad timing):
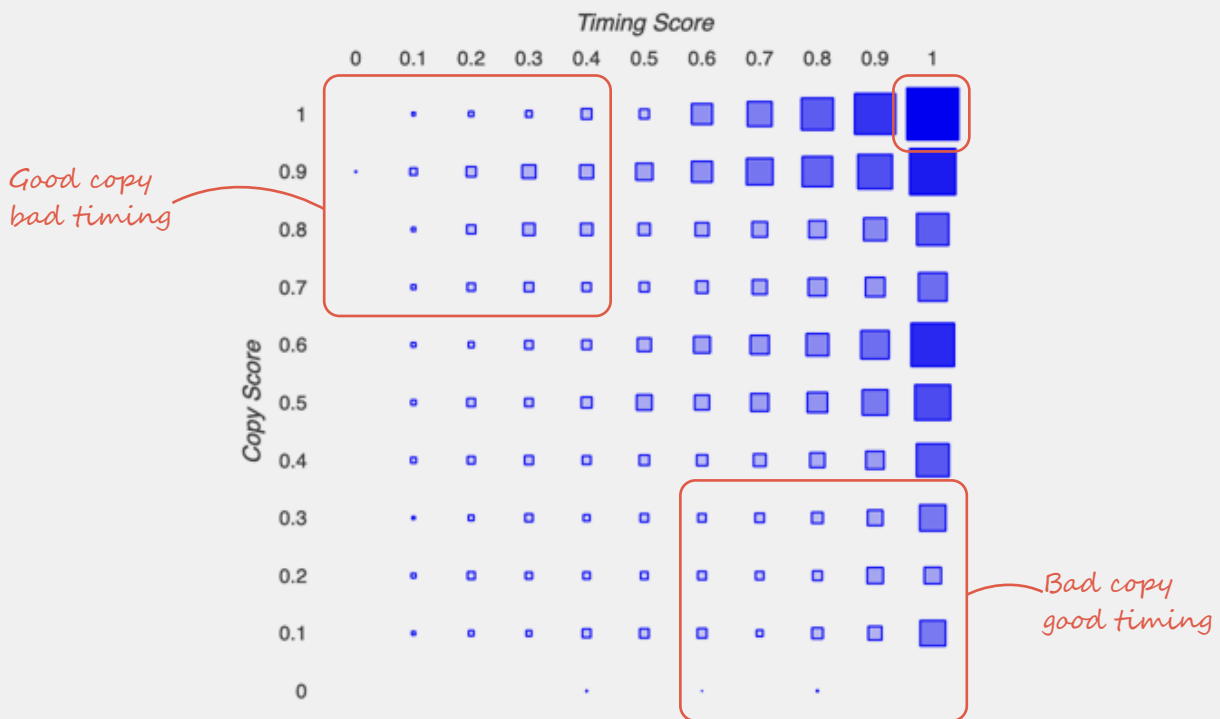
*Assuming both are decent, the impact of **timing** far exceeds the impact of the message copy itself.*

Wait...Why does bad copy/bad timing outperform good copy/bad timing?

If you study the chart closely enough, you might notice that 'bad copy/bad timing' (lower left) appears to outperform 'good copy/bad timing' (upper left).

Don't worry, it actually doesn't.

The reason the 'bad/bad' quadrant has bigger blue boxes than the 'good copy/bad timing' quadrant is simply because there are **more messages with bad timing and bad copy** than well crafted messages with good copy and bad timing.

And **higher message volume = more responses** (until all of your users unsubscribe that is!).

For the purposes of comparison, our point still stands:

Good timing is much more important than good copy alone!

Message timing: Why you shouldn't use the past to predict the future

Many conventional CRM delivery platforms claim to optimize message timing by sending messages when your users have traditionally been most active on your app.

Initially this seems to make sense, but, when you think about it, this is absolutely the wrong time to send your users push notifications.

The value of push notifications isn't to notify people when they're already on your app; the value is to pull them into your app at times when they otherwise wouldn't be on your app – this is what drives incremental revenue.

That's why **we use intelligent experimentation** (instead of just historical data) to determine optimal message timing.

You simply can't know the best message timing for each user unless you are regularly trying multiple options.
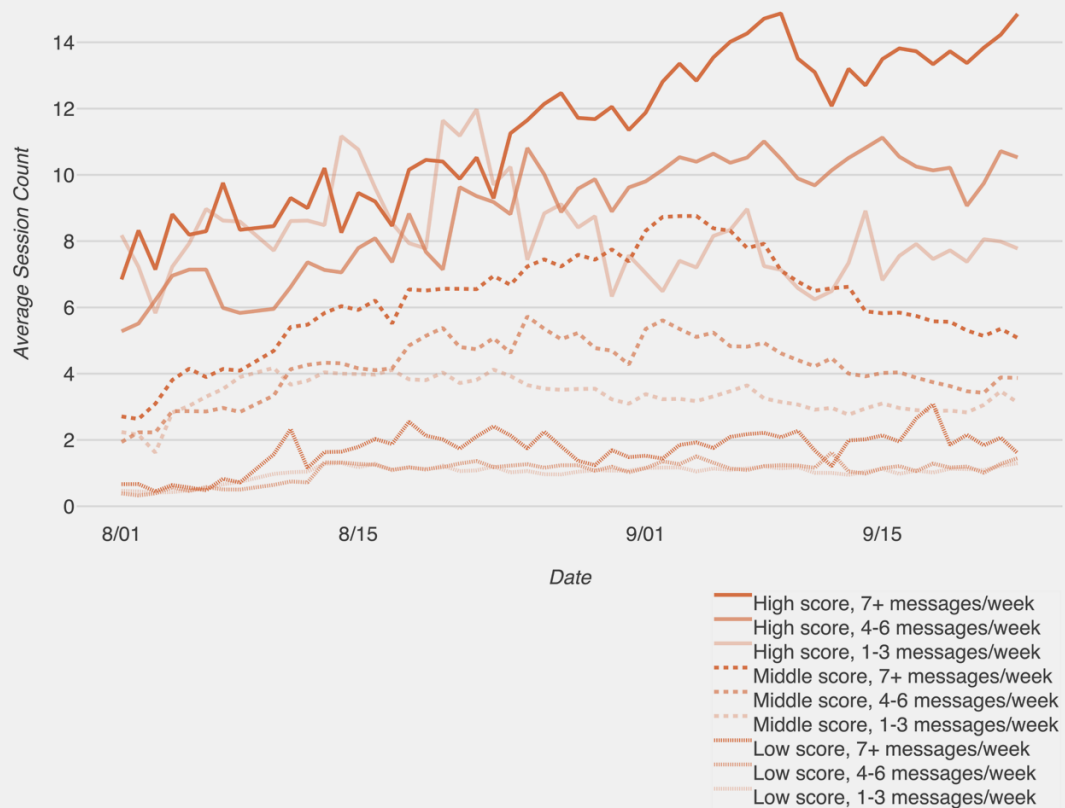
Keep users coming back, time after time

In addition to timing having the biggest impact on click-thru rate, we've also seen a positive correlation between increased user sessions (sessions being the period of time in which a user is active in the system) and optimized message timing.

(A higher average session count indicates that you're bringing users back over and over again, and bringing them back in over and over again means more chances to earn their trust, loyalty, and purchases.

If you don't get users to your app, neither you nor they get any value from the fact that they downloaded it in the first place. Increasing session counts is the first step towards increasing that value.)

Legend:
- High score, 7+ messages/week
- High score, 4-6 messages/week
- High score, 1-3 messages/week
- Middle score, 7+ messages/week
- Middle score, 4-6 messages/week
- Middle score, 1-3 messages/week
- Low score, 7+ messages/week
- Low score, 4-6 messages/week
- Low score, 1-3 messages/week

In the previous graph, solid lines meant our personalization score – for message timing alone – were high (meaning we had a lot of confidence in those choices).

For the dashed lines, we weren't as confident. The dotted lines meant we weren't confident at all. The darker the line, the more frequently we had messaged a user.

The more we expect users to engage in more sessions, the higher up they are on this list:

— High score, 7+ messages/week
— High score, 4-6 messages/week
— High score, 1-3 messages/week
---- Middle score, 7+ messages/week
---- Middle score, 4-6 messages/week
---- Middle score, 1-3 messages/week
||||||||| Low score, 7+ messages/week
||||||||| Low score, 4-6 messages/week
||||||||| Low score, 1-3 messages/week

...and, as you can see, that's exactly what happened.

The solid lines are higher than dashed and dotted lines, meaning that the more tailored the messages were to our individual users' timing preferences, the more they visited the app in the following week.

! Messages with confidently high personalization scores resulted in **1,300% to 1,800% higher average session count** than messages with low scores.
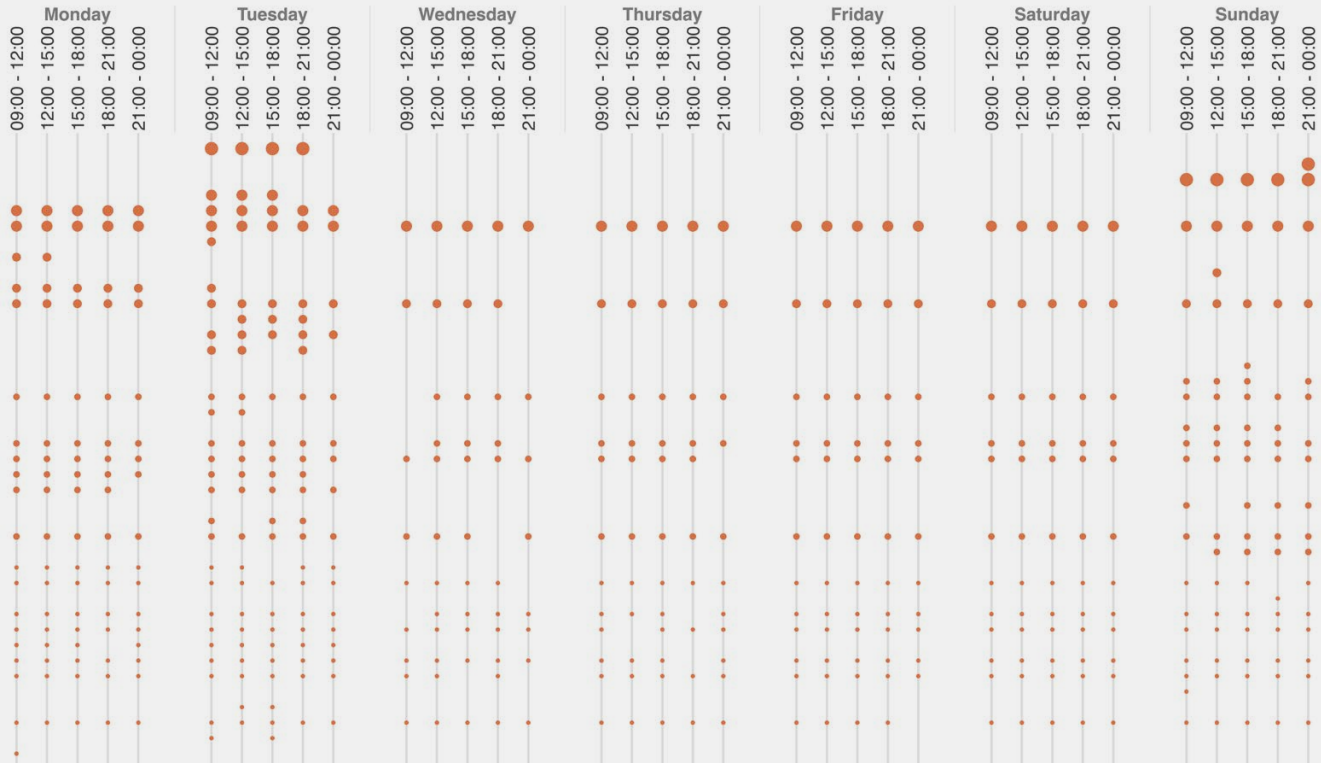
So when is the best time to send a push notification?

The quick answer is between 9am and noon on Tuesdays.

…but here's a peek behind the curtain at all the nuance that's hidden behind that simple number:

! *Between 9am and 12pm on Tuesdays.*

But there's a catch… yes this is fine print!

Top 40 Timing Patterns

487 other patterns, accounting for the remaining 28.2% of all users.

In the above graph, the columns represent different timing options per day, the rows represent different user preferences, and the intensity of the dots denotes the likelihood of a user to respond during those certain time slots.

So, for example,

- The very top row represents a relatively large number of users who only like to get messages on Tuesdays — but not late night on Tuesday.

- The next row is a big group of users who only want messages on Sundays...but only late at night.

- The row after that is a big group who is ok with messages pretty much any time on Sunday.

- ...and the next group is only ok with Tuesday again, but don't even think about messaging them in the evening!

In all, from this single group of users, our system was able to detect over 520 distinct timing preference patterns for when different users are receptive to receive messages.

So let's go back to our "9am to noon on Tuesdays"

So yeah, painting with a broad brush this time works for a good chunk of users, but, here's what really happens when you set your bulk send to "9am to noon on Tuesdays":

- You alienate your 2nd largest group of customers (who only want messages on Sunday evenings),

- You miss out on the opportunities to message multiple times per week for the users who are receptive to additional messages each week (more on message frequency in the next section),

- ...and, in general, you're still missing out on the majority of your users, who just aren't receptive to receiving messages during that narrow slice of time.

It's also worth noting that the '9am to noon on Tuesday' timeslot was the best time to message your users at the time that we calculated these results.

A few weeks later, with different patterns (say, a few new offerings on the app, or seasonal changes, or just shifts the population of users), we could easily see a completely different pattern emerge.

That's the point — Even when we find the "best" times to message users, these patterns aren't stable: They're a moving target that changes as your individual user's life situations and preferences change.

"Why are we messaging our users after 10pm?!"

Our team recently got this message from a major food delivery app in Asia, so we sprang into action to find the answer.

See, before using Aampe, this app had set its push notification cutoff at 9pm (so they wouldn't disturb anyone who was getting ready for bed or sleeping), but now that Aampe was optimizing their message timing, someone on the app's marketing team noticed **messages sending as late as 11pm**.

But how many people were getting these late messages and why? Were we driving users away?!

The answer turned out to be way more interesting than a timing bug:

When we dug in, we found out that around 5-7% of the messages sent during certain weekdays were sent between 9 and 11:30pm (the latest time a message was allowed to be sent, enforced by a system rule), and they were having a **strong rate of success**…but still —

# Why were these messages sending so late?
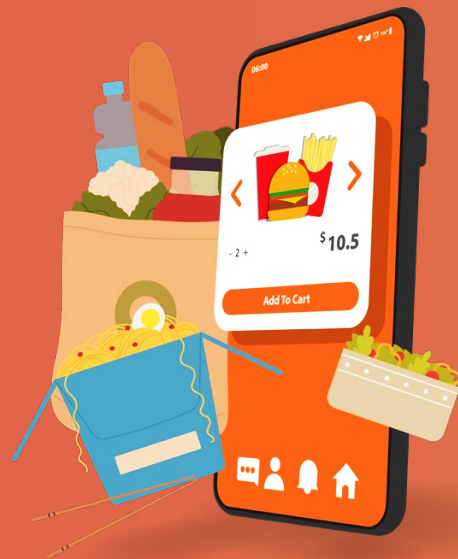
After looking into it, it turned out that a substantial population of users in this country worked for companies that operate on further-west timezones (e.g. Europe and the U.S.), which meant these were probably users overlapping with those European/American hours of business (In other words, **they were working into the late evenings**).

And, since they were working so late, one of their biggest struggles was finding "snacks" at night — When they didn't feel like cooking and most public transit was closed.

"One of their biggest struggles was finding snacks at night..."

The answer: **This food delivery app**.

Yep. Because of **Aampe's intelligent timing feature**, the app was able to discover a significant population of users who actually wanted to place orders well past the latest time the team was previously sending their batch campaigns.
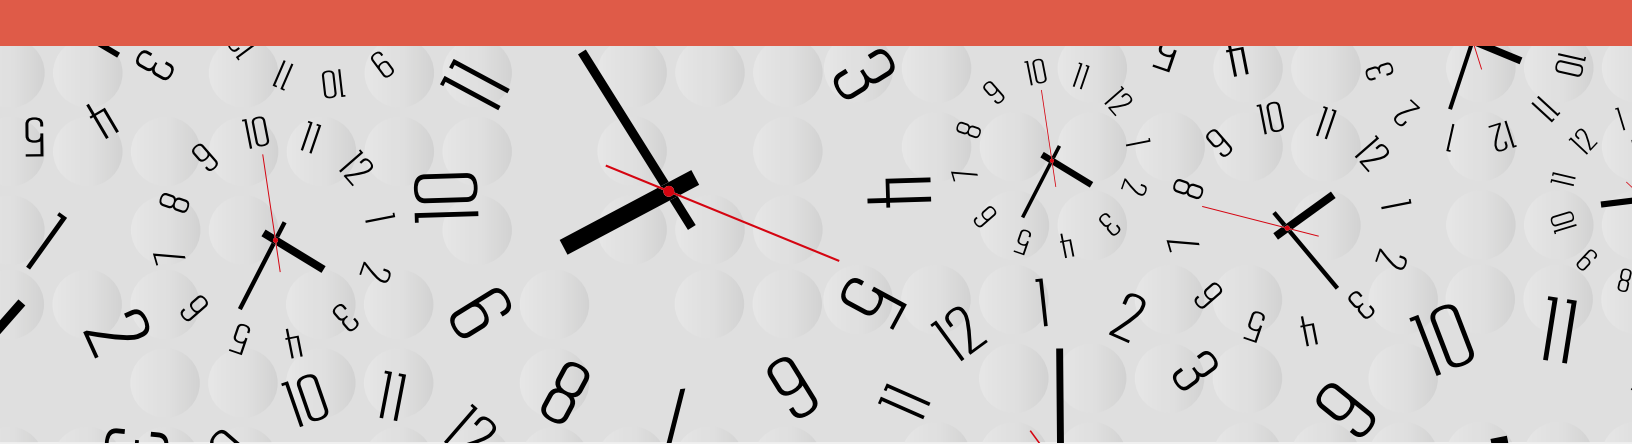
How much of a difference does the right timing make?

Well, here's just one example.

This is 2 months' worth of messaging and response data from one of our customers, across 3 different countries:

	Wrong Time	Right Time		% Difference
Country A	0.3%	1.1%	➡	**328.5%**
Country B	0.5%	1.8%	➡	**272.9%**
Country C	0.6%	2.1%	➡	**239.7%**

> *These percentages aren't just click-thru rates, they're the percentage of messages that resulted in a **checkout within 24 hours** of the message being sent. So, while the numbers by themselves look small, keep in mind that **every bit of that is money in our customer's pocket**.*

330%

increase in orders placed?

Stop the clock! Who wouldn't take that?!

So, as we've seen, **sending messages at the right time for each individual user is critical** for exceeding your key metrics and **increasing your ROI**.

But there's a different timing-related metric that affects ROI…and has an **even bigger effect** on user churn.
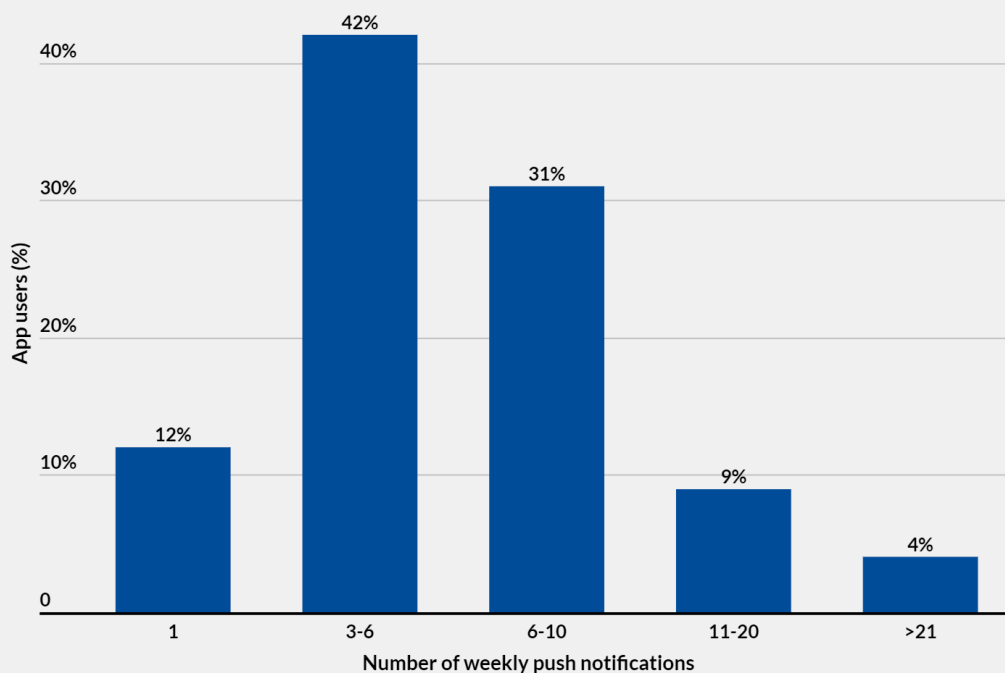
Bad frequency puts you on the wrong wavelength

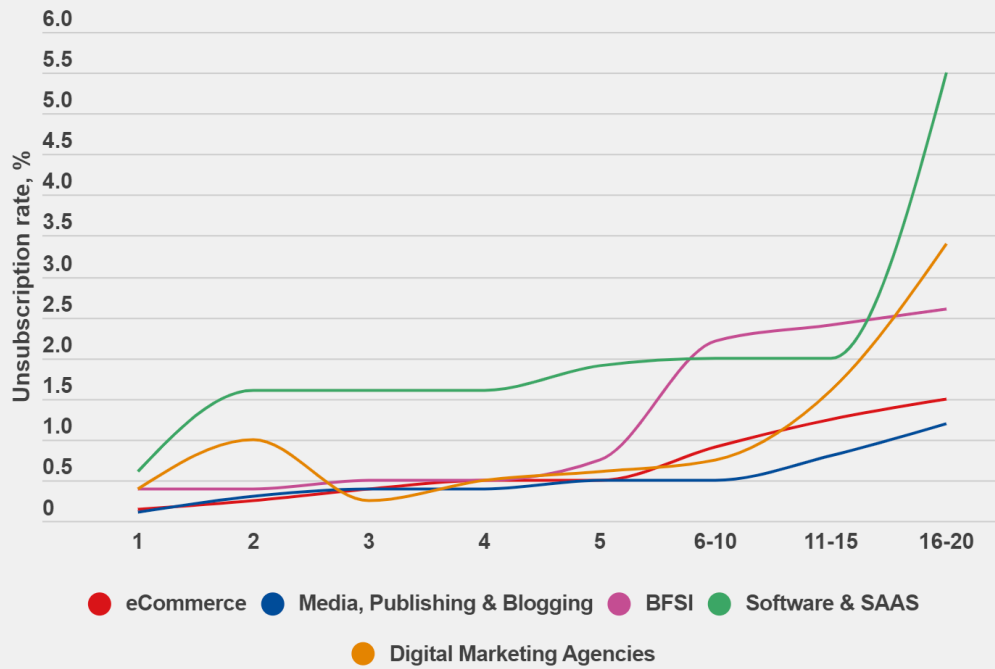Nothing drives a customer away faster than receiving too many messages too soon.

In fact, the latest research from Helplama shows that 42% of users will uninstall an app if they receive 3-6 push notifications per week, and VWO Engage found that your overall unsubscription rate generally increases the more messages you send.

Number of weekly push notifications that cause users to disable push notifications (%)

Daily push notification frequency vs. unsubscribe rate by industry (%)

Legend: ● eCommerce ● Media, Publishing & Blogging ● BFSI ● Software & SAAS ● Digital Marketing Agencies

Furthermore, setting your messaging frequency too high is the largest factor causing users to see your messages as spam:

What makes push notifications classify as spam? (%)

Category	Subscriber	User
Too many	62%	73%
Wrong timed	44.82%	12.24%
Irrelevant	55%	10.2%
Invalid / expired offers	48.18%	48.98%
Clickbait	55.18%	55.10%

Legend: ● Subscriber ● User
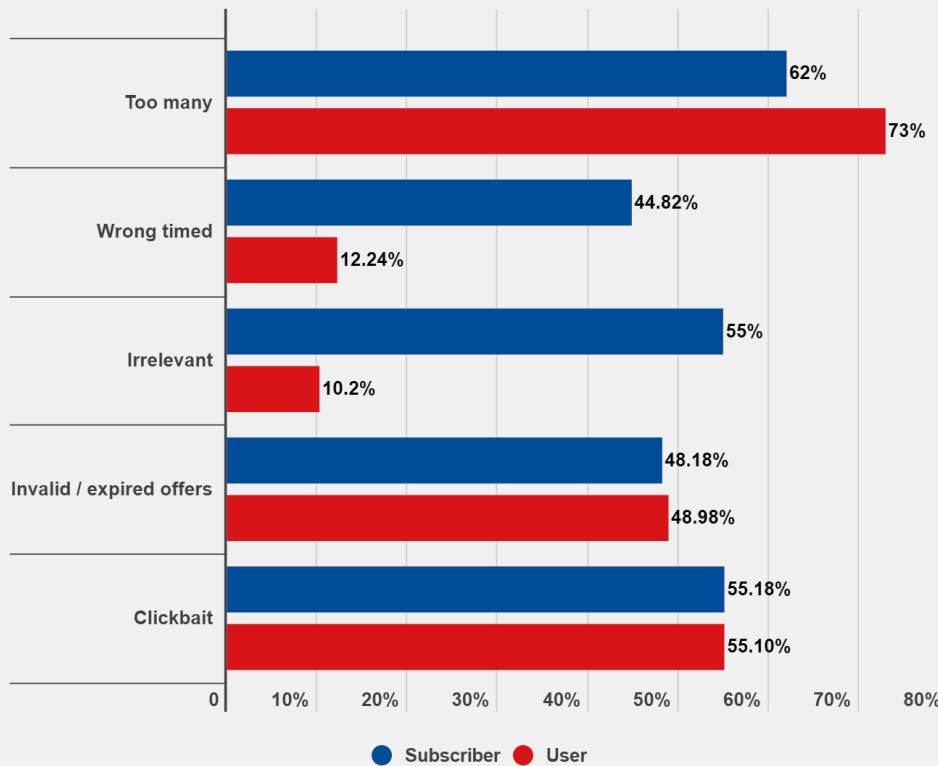
But not sending enough messages means missed sales

The truth is that **each person has their own preferences** for message frequency. Some are ok with daily notifications, but, for others, the same daily notification is way too much:

Jamie Perez

There's just something joyous about @slice daily notifications. Can't quite explain it.

Andrew Stella

love @slice but damn the push notifications trying to get me to eat pizza every night!

So how do you know the optimal messaging frequency for each individual user?

PushEngage says you can send 1-2 push notifications per day, but Sellbrite says not to send more than 3-5 push notifications per week.

As you can probably guess by now, the reality is much more messy (or beautiful, depending on your perspective):
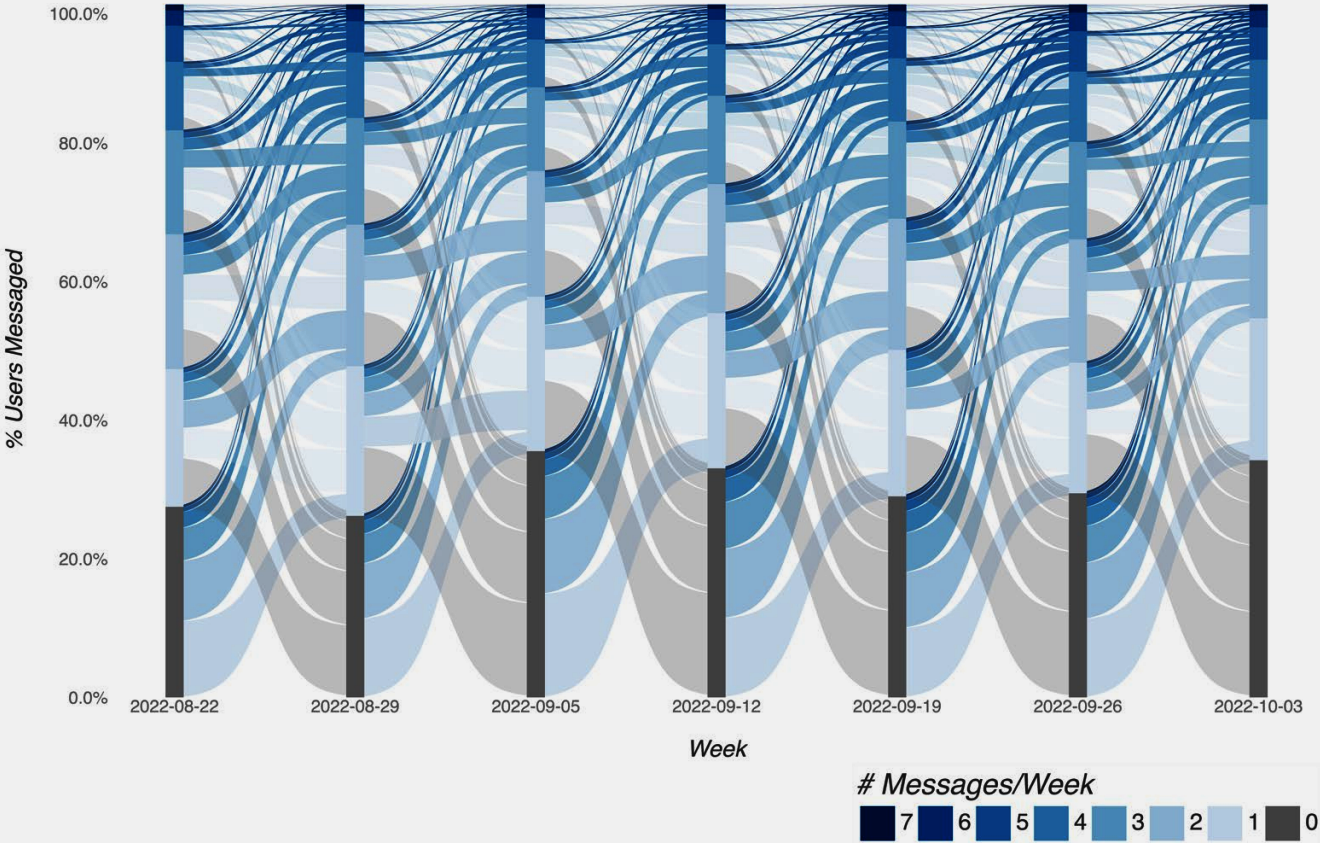
The graph above shows how many messages we're sending per week by percentage of users, and how this frequency changes week to week depending on user-initiated factors like activity level, type of activity, and even user inactivity.

Put another, far less beautiful way, the graph below shows what percentage of users are receiving which volumes of messages per week based on the results of our algorithm **determining, per user, how favorable each messaging opportunity is for them, specifically**.

(In other words, we don't just 'decide to send a user 2 messages per week' — We evaluate every opportunity in the entire opportunity set and decide whether or not to send a notification to a user during a specific time window.

…We don't think about frequency beforehand…**Frequency happens as a purposeful result**.)

Week	2022-05-22	2022-05-29	2022-06-05	2022-06-12	2022-06-19	2022-06-26	2022-07-03
# of Messages							
1	12%	14%	17%	14%	14%	15%	15%
2	17%	19%	22%	19%	19%	20%	19%
3	21%	22%	23%	21%	21%	21%	21%
4	19%	19%	16%	18%	18%	17%	17%
5	13%	12%	9%	12%	12%	11%	11%
6	7%	6%	5%	7%	7%	6%	7%
7	5%	4%	4%	4%	4%	4%	4%
8	3%	3%	2%	3%	3%	3%	3%
9	2%	1%	1%	2%	1%	2%	2%
10	<1%	<1%	<1%	<1%	<1%	<1%	<1%

So, in week 1, 12% of this app's users received a single message, 17% of users received 2 messages, etc., and this **messaging frequency continuously adjusts** (again, based on our assessment of each messaging opportunity).

How do we know we're not burning users out while we're learning?

Simple. We're **not messaging every single user every single time**.

The first thing Aampe's system does is cluster users based on app usage and behavior. By clustering users, we can then test messages efficiently, develop a model that incorporates a large set of features that represent users' app behavior, and then iterate, iterate, iterate.

This helps us **learn quickly** while not causing your users to hit unsubscribe.

Simple. We're not messaging every single user every single time.

So sure, **in this case**, 3 messages per week would be the most optimal (at **21% of the user population**)...

3

Messages per week

21%
Served

34%
Annoyed

45%
Underserved
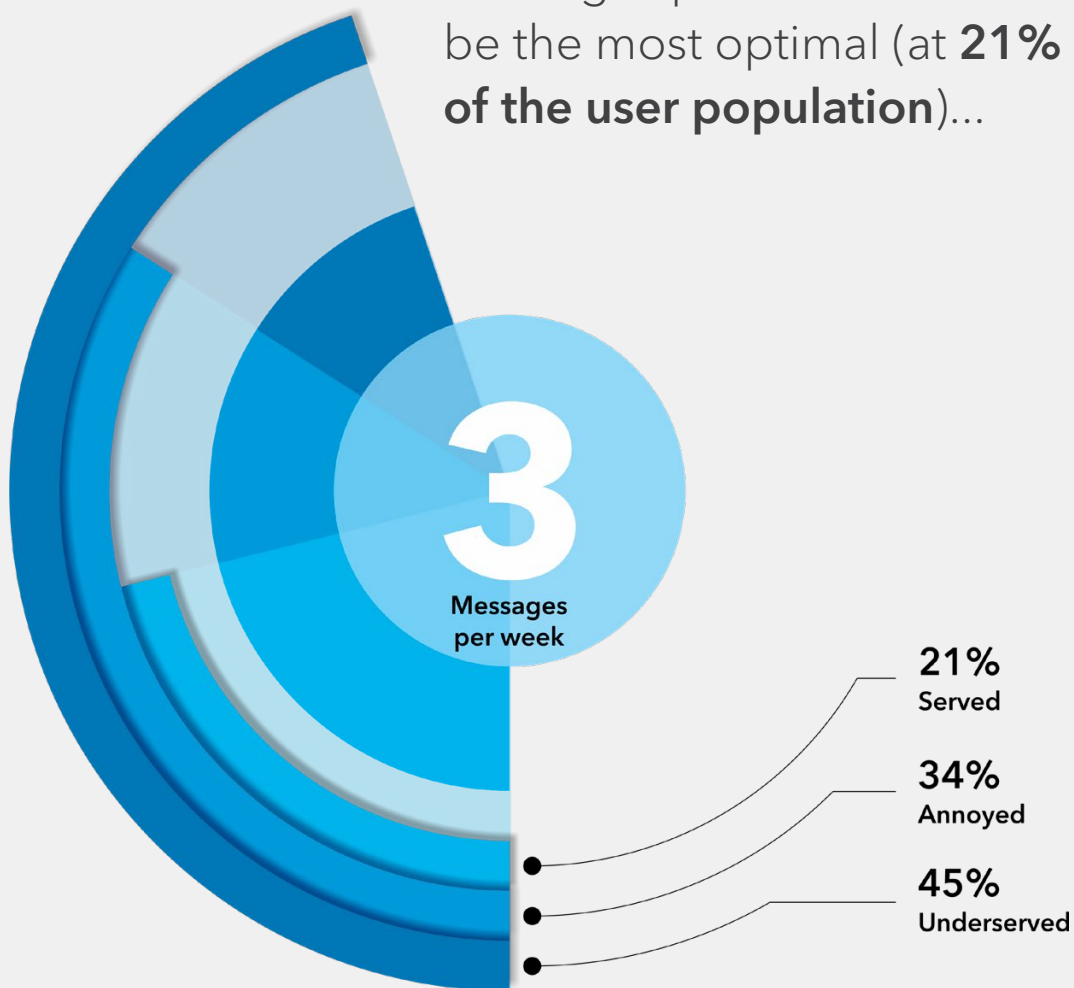
but there could be some big ramifications caused by setting this message frequency as a global rule:

- By choosing 3 messages per week, you're **potentially underserving the ~45%** of users who would be receptive to receiving more messages (and thus making more purchases) per week.

- You're also **potentially driving away a portion of the 34%** of users who prefer less than 3 notifications per week.

(…and of course this all assumes that your users preferences are static and their preferences for message frequency doesn't change…which we know they do.)

So, what is the best timing and frequency to send your users a push notification?

Did you really just ask that question? (...and, if you just skipped to this section hoping for the easy answer, you might just want to go back and read the whole eBook next time.)

By now you've learned that the best time to send push notifications is on Tuesday from 9am to noon, and the best messaging frequency is 3x week.

...and you've also learned that, if you follow those suggestions perfectly to a T, you'll:

- **Alienate** large groups of your customers

- **Miss out** on the majority of your users who aren't receptive to receiving messages during that narrow slice of time

- **Underserve ~46%** of your users who would be receptive to receive more than 3 messages per week

- **Annoy the 34%** of customers who don't want more than 2 messages per week

...and **miss out on a load of incremental revenue** in the process.

But wait! Is there a way to find these individual timing preferences yourself?

It's painful, but it's still possible, and it goes like this:

1. **Split your current users into like groups based on their app usage and behavior.**

   When did they install your app? What time do they typically use your app and for how long? When do people who behave similarly to the user in question also use your app?

2. **Conduct highly-parallelized experiments on these user groups to learn what timing works best.**

   Make sure to normalize your data across your audience, so you have an accurate baseline and control group. Also be sure you're **testing every single time window** – You never know if you have the best time window unless you test them all.
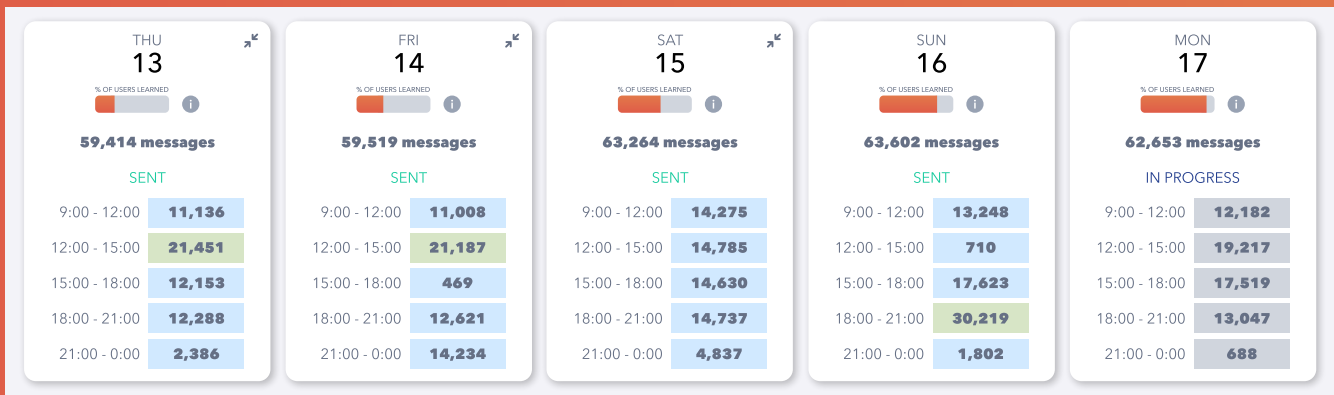
   Be sure to track more than clicks – Certain times are better at driving conversions than others. We need to find those times!

3. **Lather, rinse, and repeat!**

   People change, seasons change, and life events happen. The best time today, isn't likely to be the best time tomorrow, so keep iterating this model and testing new times to ensure you're still on target.

Yep, with a dedicated team of engineers and data scientists and a few months of effort, plus regular pipeline maintenance you can approach optimal timing for your users.

…or you can call Aampe and, with a few hours of set-up, you can start seeing effective timing personalization in just a couple of weeks.

THU 13	FRI 14	SAT 15	SUN 16	MON 17
% OF USERS LEARNED	% OF USERS LEARNED	% OF USERS LEARNED	% OF USERS LEARNED	% OF USERS LEARNED
59,414 messages	**59,519 messages**	**63,264 messages**	**63,602 messages**	**62,653 messages**
SENT	SENT	SENT	SENT	IN PROGRESS

	THU 13	FRI 14	SAT 15	SUN 16	MON 17
9:00 - 12:00	11,136	11,008	14,275	13,248	12,182
12:00 - 15:00	21,451	21,187	14,785	710	19,217
15:00 - 18:00	12,153	469	14,630	17,623	17,519
18:00 - 21:00	12,288	12,621	14,737	30,219	13,047
21:00 - 0:00	2,386	14,234	4,837	1,802	688

Want us to generate this data for your app?

We're Aampe, and using data science principles to intelligently optimize push notification messaging and timing is what we do.

Companies across the globe use Aampe to **help increase their conversions** and **drive customer retention** with push notifications that are *actually* personalized to each of their individual users – without having to replace their existing messaging deliverability platform.

To test Aampe out for yourself, **login here**, or drop us a line at **hello@Aampe.com**.

Ready to send better messages? It's about damn time.

aampe

The pioneers of Message-Led Personalization

Please visit **aampe.com** or email **get@aampe.com**