

Beyond A/B testing: Automation within testing in Adobe Target

By Jason Hickey

Table of contents

The conundrum with A/B Testing

Manual allocation: When you have to know everything about every experience

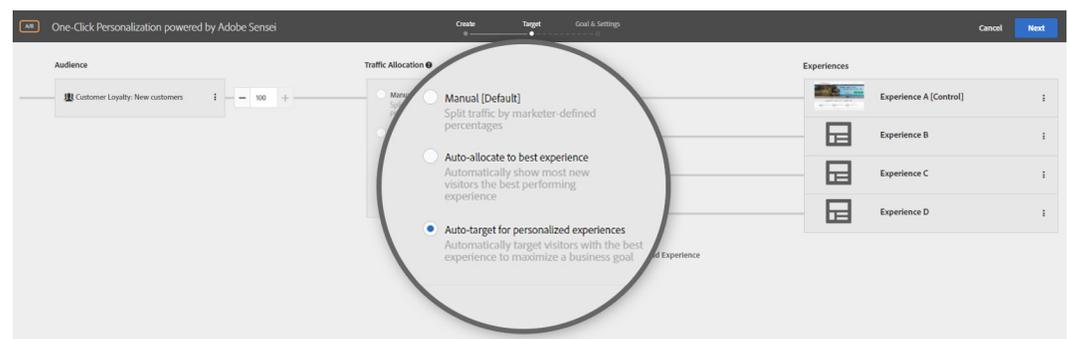
Auto-Allocate: When you want to find and exploit the winner faster

Auto-Target: When you want each visitor to get the experience that wins for them

The potential ROI and best uses for each traffic allocation method

Bridge the gap between testing and personalization with Adobe Target

You may have noticed a few new options in Adobe Target recently when creating an A/B test—specifically when selecting an audience on step two of the activity creation workflow. These new one-click options give you choices for how to allocate traffic between the control and test variations. You'll of course see the option you always had—Manual—but you now also see two new options that rely on the machine learning and artificial intelligence (AI) of Adobe Sensei, Auto-Allocate, and if you are a Target Premium user, Auto-Target. What's the difference between these three approaches to allocating your traffic, and when should you use one versus the other?



In Adobe Target, when setting up an A/B test activity, you can allocate traffic to a test using one of three different options—Manual, Auto-Allocate, and Auto-Target.

In this white paper, you'll learn about the three different ways of allocating traffic to your test experiences in Adobe Target, and when you might want to use one approach versus another. Ultimately, this paper helps clear the path toward using automation in personalization and optimization that's often been obscured by a lack of understanding of why you might want to use it, how it works, and the huge benefits it offers.

The conundrum with A/B testing

A/B testing, or split testing, allows you, whether you're a marketer, product manager, engineer or some other stakeholder, to use data to determine if one element or variation of your web page, mobile app, or other digital surface performs better than one or more alternatives. For example, if your CEO insists that different copy on a CTA on your home page will drive more visitors to click and explore a particular offer than the current copy, A/B testing can prove or disprove your CEO's hypothesis.

The most common approach to A/B testing is best described in statistical terms as fixed-horizon. That is to say, before running any test you should calculate the number of visitors—the traffic—that must participate in an A/B test before you can conclude that the test results are statistically meaningful. You've likely had one big objection to this approach: for the duration of the test, Adobe Target is serving visitors your losing alternatives, and they're losing money for you. Although you may have waited for the test to conclude before serving the winning variation to all your visitors, you've likely stopped at least one test early, serving the variation that appeared to be a winner to all visitors in your test audience.

But this still leaves you with some uncomfortable questions:

- If you concluded a test early because you are fairly certain you know the winner, was the winning experience really the winner, or was it a false positive?
- How do you avoid the opportunity cost from serving an underperforming experience to members of your test audience?
- Did the winning experience win with everyone in your test audience, or just a few select audiences?
- Did you call a winner too early—before the test concluded?

These very reasonable questions led to the development of two brand new ways to allocate traffic to your test audience in Adobe Target beyond the traditional manual approach—Auto-Allocate and Auto-Target. Both of these traffic allocation options are powered by the machine learning and artificial intelligence of Adobe Sensei.

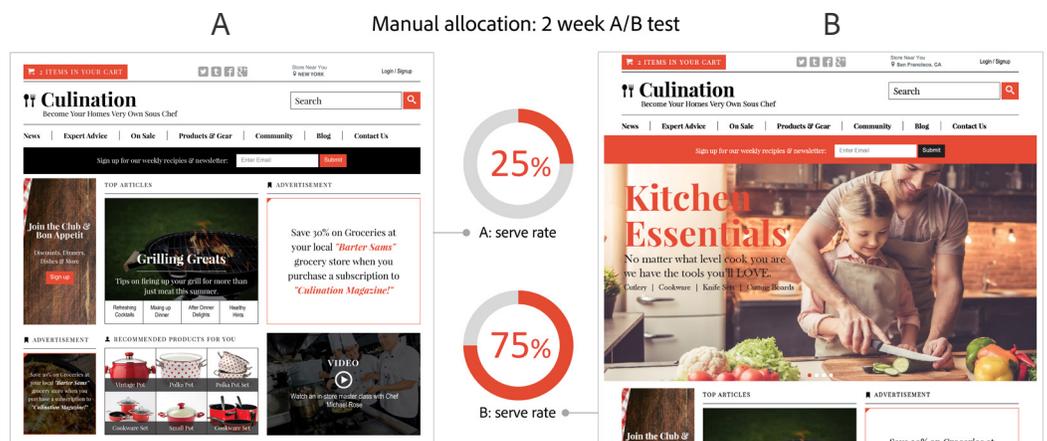
Manual allocation: When you have to know everything about every experience

In some cases, it isn't enough to know that B is better than A—you have to know exactly by how much and under what circumstances. This may be because a change has to exceed certain thresholds to be considered viable long term, or because the change is a major business or technology decision for the brand. When you need additional insight via deep analytics reporting, a consistent set of parameters for measurement, and a predictable number of visitors in every experience, you should use Manual traffic allocation.

Here's why:

Manual allocation of traffic to a test is a time tested and statistically robust way to determine the winner. Using the Adobe Target confidence calculator, you can control exactly how much statistical power and confidence to factor into the decision-making process. This translates into being able to control and calculate your risk of false positives (type I error) and false negatives (type II error). However, getting to that point of certainty about the winner takes time, and you have to be willing to take that time—otherwise you lack that confidence. In addition, it's easy to fall prey to several common issues, which are described in the very understandable terms in the Adobe Target documentation topic: [9 Pitfalls of A/B Testing](#).*

That said, you do have a few options to consider when setting up your test. For example, if you're super confident in your hypothesis and just want to prove that a specific alternative is the best performing, you can divert most of your audience traffic to that alternative for the duration of the test. For example, you could divert 25% to the current variation, but 75% to the alternative. Conversely, if you're not at all certain that any of your test variants will win over the control, then you might do just the opposite, pushing more traffic to your control. Note that allocating traffic in a non-uniform way will prolong the test and contains the risk that if your idea of which experience will be the winner is incorrect, it could cost more than if you had allocated the same traffic to every experience.



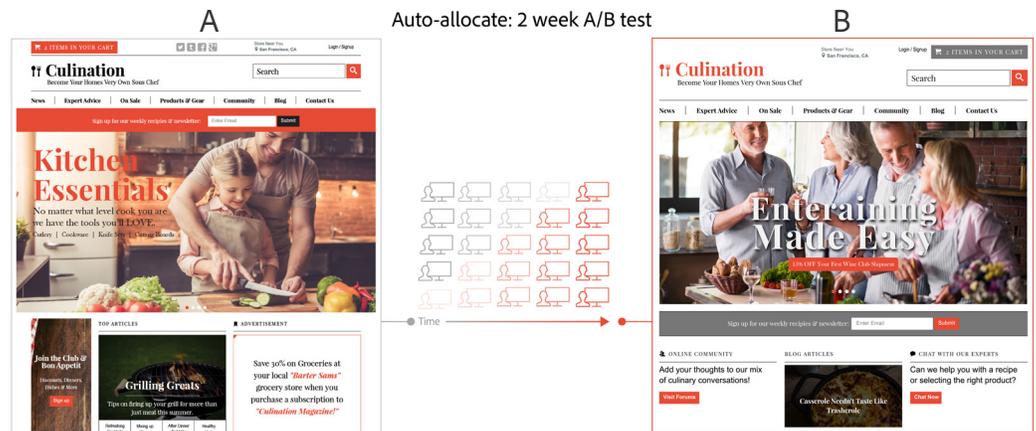
When you choose to use Manual traffic allocation in an A/B test, you specify how much traffic to divert to each experience for the duration of the test.

Regardless of how you manually allocate your audience traffic, you'll still have to serve the test to the number of visitors you originally specified. In other words, you run the test until it's done.

Auto-Allocate: When you want to find and exploit the winner faster

As described above, manual A/B tests are statistically rigorous, but they also come with an inherent cost—you have to spend traffic to measure the performance of each experience, and that traffic distribution remains fixed even after you recognize that some experiences are outperforming others. It can be complicated to figure out the sample size, and the activity must run its entire course before you can act on a winner. After doing all of this, you still have the chance of a false positive if you skipped any critical steps along the way, or if you or your marketer or analyst running the test lacks a foundational understanding of statistics and statistical inferencing.

For many of you, this inherent cost simply gets in the way of finding the answer to your question. What is the best experience to show my customers? Is B or C or D better than A? You need the option to serve the winning experience more often and earlier in the test while simultaneously removing or reducing the setup and calculation cost of determining sample sizes, confidence levels, and other statistical concepts. This is where the next method of allocating your audience traffic comes in, Auto-Allocate. If you're a marketer, it's also an easy first step into the realm of using automation based on algorithms, probability, and machine learning.



With Auto-Allocate, as a winning experience emerges, Adobe Target diverts more of the traffic to that winner early on in the test.

So how does Auto-Allocate work?

Auto-Allocate uses the principle of the multi-armed bandit. In case the term is unfamiliar, a one-armed bandit is a colloquial term for a slot machine (think Las Vegas). Visualize auto-allocation of traffic as having multiple slot machines, in this case, test variations, and starting off by pulling all handles equally. Over time, one or more machines, or test variations, might pay out more than others. When this happens, you would naturally start pulling the handles of the ones that win more often. That's what Auto-Allocate does; in traffic allocation terms, Adobe Target will serve more of your visitors the experience or experiences that are winning more.

The upside is pretty clear: more of your visitors see the variations that perform best. And as a single variation pulls ahead, even more visitors get diverted to that winning experience. This overcomes the issue with manual A/B testing of being unable to capitalize on the winning experience while the test is still running. This can be especially critical you're running an A/B test during a core business moment, such as a holiday, product launch, or a weather or world news event. For example, if your business thrives on a "back-to-school" period, those months are immensely more valuable than the dead of winter. Auto-Allocate both finds the winner faster than a manual A/B split, but also allows you to exploit that winner immediately—capturing upside revenue that would have been lost in the traditional manual approach.

Auto-Allocate also differs from manual A/B testing in the statistical methodologies it uses to help you find a winner. The net effect of these differences is that with Auto-Allocate, you don't have to calculate a sample size up front—it's useful to set expectations, but not mandatory because the system proactively determines the winner instead of relying on your input as to when to end a test. Auto-Allocate intelligently allocates new visitors to your various experiences until it has identified a winning experience. When it stops, there is a 95% chance that the true response of the experience it returns is within N% of the true response of the best experience. Note that auto-allocate may return any experience that is close enough to being a winner if there are multiple winners whose responses are within its relative sensitivity.

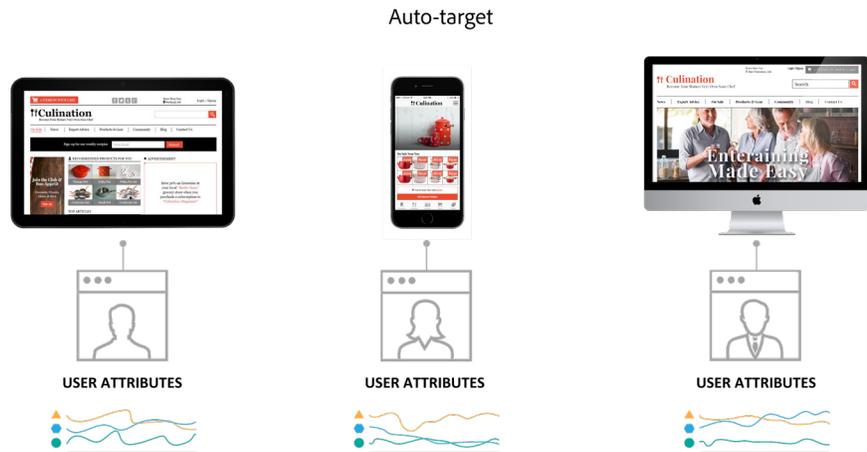
You do have to make a tradeoff when using this auto-allocated approach—although it does find and exploit your top performing experiences, it does not guarantee that you'll see notable differences between your lower performing experiences. That's because the relative differences between your top performers are amplified due to the higher traffic levels they receive. The converse is true for your lower-performing experiences—the differences between them will be so much less pronounced that you can't really detect differences reliably. If it's important to rank the performance of all your experiences, then manual A/B testing is the appropriate test approach.

Auto-Target: When you want each visitor to get the experience that wins for them

Each time you serve a visitor the experience that's not the winner for them, you know you're sacrificing the opportunity to convert them. In most cases that means lost revenue. You probably always have the feeling when you run an A/B test that while it proves one experience as the winner, that experience couldn't really be the winner for every visitor every time. Maybe some visitors actually prefer the variation that came in second place. Others may actually prefer the original control.

Obviously the best approach to this dilemma is to give each of your visitors their own personal winner from among the available experiences. In this case all experiences have the potential of being a winner with someone. Although segmentation can solve part of this issue, it can only take you so far. What if you failed to consider or didn't even know about certain segments or combinations of segments? Consider the manual and tedious process of finding these segments and assigning them experiences. There must be a better way.

One-click personalization with Auto-Target provides you with the ideal solution to this problem. This approach to A/B testing uses every bit of information provided by your visitor in the visitor profile, such as geo-location, device information, new or returning visitor, and past visit patterns or behavior, to determine what's most predictive about your visitor at that moment and deliver the best variation for that visitor for that visit. It's essentially a perpetual testing and personalization machine. Auto-Target leverages what it knows about your visitor to deliver the right experience in the moment, and unlike the Auto-Allocate and Manual traffic allocation options, these experiences can change over time. In a no-single-size-fits-all world, Auto-Target finds the best fit for each of your visitors.

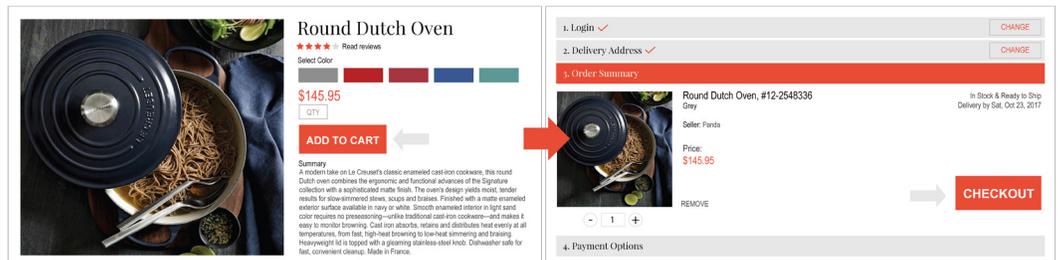


Adobe-Target delivers the experience to each visitor that wins with that visitor for that specific point in time using all the visitor profile data it has for that visitor.

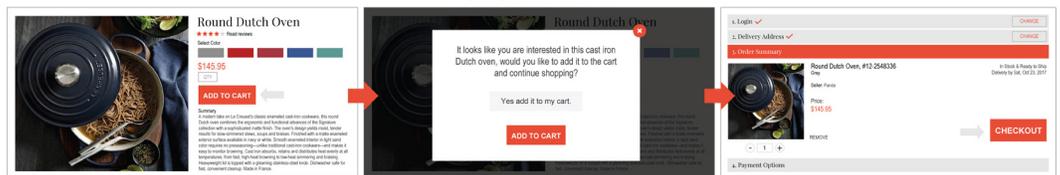
Auto-Target lets you leverage powerful artificial intelligence to make these decisions—powered by a patented and proprietary ensemble algorithm approach from Adobe that takes advantage of the latest in machine learning capabilities. The algorithm factors the entire visitor profile into its decision trees, including any data shared from other Adobe solutions like Adobe Analytics or Adobe Audience Manager. It also employs multi-armed bandit testing as a "backup policy" to act as a form of lift-insurance, reducing the likelihood of negative lift. As the newest capability in Adobe Target, Auto-Target offers you sophisticated personalization that with a single click turns any A/B test into a data-driven personalization activity.

While Auto-Target sits next to Manual and Auto-Allocate options in the test setup workflow, testing tells you only half of the story of this feature. Both Manual and Auto-Allocate seek to find a single best winner that fits the majority of your visitors overall. Auto-Target, by contrast, seeks to find a winner for each of your visitors on every visit so that every experience is the winner for some of your site population. However, all three traffic allocation approaches—Manual, Auto-Allocate, and Auto-Target—allow you to personalize entire customer experiences, not just a banner or a button. For example, experience testing can involve testing conversions for the entire experience of your visitors adding an item to a shopping cart. This experience can involve your visitor simply clicking an Add to Cart button from one of your product landing pages, but it could also involve providing your visitor a popup offer while they're on your product page that entices them to add the item to the cart and having them click to add it from the offer.

Checkout: A



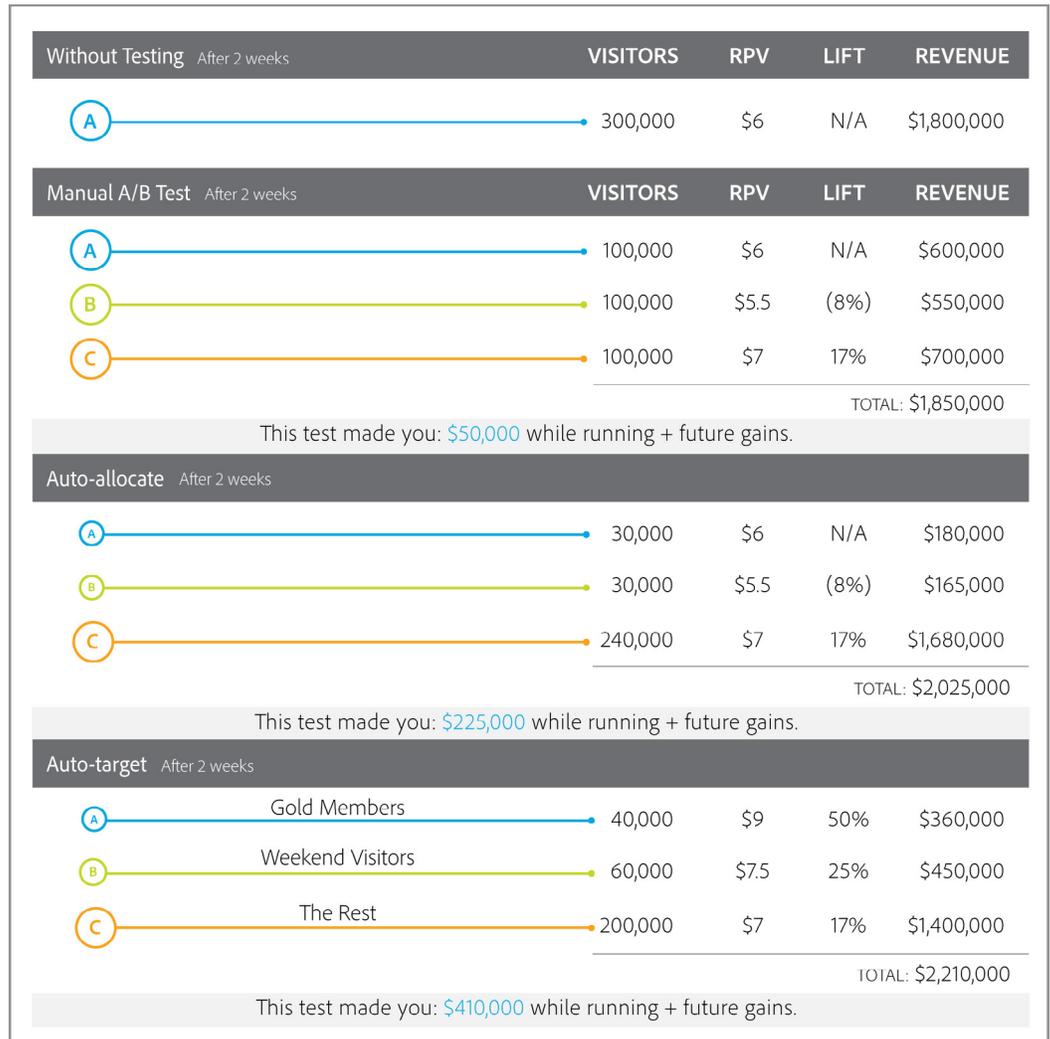
Checkout: B



Regardless of the traffic allocation method you choose, Adobe Target lets you test entire experiences, not just a single page or element of a website page, mobile app page, kiosk, and so on.

The potential ROI and best uses for each traffic allocation method

It might seem clear by now that Auto-Target stands to provide you the greatest ROI—after all, if you deliver the experience that wins to each visitor, those visitors are far more likely to convert. That means you're far more likely to see greater ROI from the test activity. Similarly, you'd expect Auto-Allocate to deliver greater ROI than manual A/B testing, but less than you'd gain from Auto-Target. The image below gives a hypothetical, but realistic example of the relative levels of ROI you could expect from each method. But before you jump in to use Auto-Target, you need to understand when it's best to use one method over another and the tradeoffs with each.



A hypothetical, but plausible example of the potential ROI from using each traffic allocation method.

You've learned about a few of those tradeoffs already—with a manual A/B test, you must determine the sample size before you launch your test and run that test for the time it takes to reach that sample size. You also have to determine the statistical confidence and power that you want for your test. Most importantly, you may be losing conversions because you can't deliver the winning experience to your entire audience until the test completes. But when you need complete control over your traffic allocation and risk tolerances, Manual is the method to use. In addition, Manual allocation allows you to dig more deeply into your test results to identify valuable segments and metrics in Adobe Analytics.

But what about the uses for and tradeoffs with the two options powered by Adobe Sensei—Auto-Allocate and Auto-Target?

As discussed earlier, if you need to know how each experience performed relative to the others, Auto-Allocate can't do that. In this case, Manual allocation is your better option. But if you want to expose more traffic to the winning experience earlier, then Auto-Allocate is a great choice.

Auto-Target is ideal when you don't plan to change your content and experiences too frequently—that fails to give the algorithm sufficient time to refine its understanding of your visitors and their preferences so that it can deliver the best experience. Compared to Manual and Auto-Allocate, in which each test has a definitive beginning and end, Auto-Target can be considered to be "always on." With Auto-Target, you have no one best winner, so you don't need to stop the test while it is optimizing personalized experiences. For this reason, Auto-Target may not be the best fit if you know that operating costs or brand considerations may necessitate your stopping a test activity. Barring these two situations, though, Auto-Target can deliver incredible results and is worth trying out.

Bridge the gap between testing and personalization with Adobe Target

Adobe Target lets you choose how much or little you want to let data and automation drive the content and experiences that visitors see and have on your site. Use Manual traffic allocation for complete control of how much of your visitor traffic sees each test variant and for customization of the statistical thresholds that are relevant for your business. Take a first step into automation from Adobe Sensei with Auto-Allocate, letting Adobe Target dynamically decide how much traffic to serve each experience; you'll get more conversions and increase ROI as your tests run. Finally, with Sensei-driven Auto-Target, let Adobe Target analyze all the data you have about each visitor and use it to make AI-driven decisions about what experience to serve each visitor for the best results possible.

About the author

Jason Hickey is a Senior Product Marketing Manager for Adobe Target and is highly passionate about data driven decision making and conversion rate optimization. A user of Adobe Target for over ten years, Jason is now sharing all things personalization and optimization with the world.

For more information

[http://www.adobe.com/
marketing-cloud/target.html](http://www.adobe.com/marketing-cloud/target.html)



Adobe

Adobe
345 Park Avenue
San Jose, CA 95110-2704
USA
www.adobe.com

Adobe and the Adobe logo are either registered trademarks or trademarks of Adobe in the United States and/or other countries. All other trademarks are the property of their respective owners.

© 2019 Adobe . All rights reserved. Printed in the USA.