

FAQ

Getting Started

What are the benefits of Dynamic Testing experiences?

Dynamic Testing experiences have the following benefits compared to A/B testing:

- **Find higher ROI faster:** There's no waiting for the testing to reach significance, with incremental changes to the winning variant over time
- **View real-time results:** Continuously monitor traffic and data as it rolls in
- **Reduce risk:** Automatically downplay low-performing experience variants, and minimize risk quickly when you try out a new idea
- **Understand results more easily:** View a simple probability-based stats interface to answer the question, "Which is better?"
- **React to changing behavior:** A Dynamic Testing experience tracks which content is the statistical leader and observes all other variants as well
- **Use fewer resources with automation:** A Dynamic Testing experience automatically adjusts variant distribution on the fly, so you don't have to worry about which variant is performing better

When do I use Dynamic Testing experiences, and when should I use standard A/B testing?

Dynamic Testing experiences are great for most testing scenarios. They help you find results faster, exploit your goal metric, offer real-time analytics, and reduce the need for a strong testing methodology.

However, in some situations you may want to consider standard A/B testing instead:

- **When you need to differentiate clearly between low performers:** A/B testing helps define the performance of all variants before you optimize your site. Dynamic Testing experiences, on the other hand, focus on differentiating between best performers, downplaying underperforming variants. Standard A/B testing helps you clearly understand statistically significant results for all variants.
- **When you have a dramatically different sense of urgency across variants:** The Dynamic Testing experience algorithm is sensitive to its recent observations. If one variant drives urgency when another does not, this could influence results. In this case, a standard A/B test may be your best option.

Are Dynamic Testing experiences as statistically valid as standard A/B testing?

Yes, Dynamic Testing experiences uses sequential Bayesian updating to learn as the test moves forward. It's important to note that the question that Dynamic Testing experiences ask is much different than standard A/B testing.

Standard A/B testing, or traditional statistics, begins with a concept called the null hypothesis. When you set up

an experiment, "all variants have equal performance" is an example of the null hypothesis. A standard A/B test then gathers evidence to make a judgment about whether you can reject this null hypothesis. In other words, it makes the judgement that there is no evidence that all variants have equal performance. If you can reject the null hypothesis, you've found a statistically significant result.

A Dynamic Testing experience, on the other hand, asks a fundamentally different question through its mathematical procedure: "Given what I know at this moment, what is the probability that this is the best performing variant?"

Standard A/B/n testing asks a very different question: "What is the probability that I see a specific outcome if all the variants perform equally?"

Both questions are valid and sound, but the question a Dynamic Testing experience asks is a bit more intuitive and easier for most people to grasp. It's basically asking, "Which is better?"

Standard A/B testing requires careful methodology to ensure that you're not susceptible to risky decisions or that you're not susceptible to statistical error. You must wait until you've seen a specific number of sessions before you can ensure your test is powerful enough to avoid being susceptible to false positives.

In standard A/B testing, statistical significance prevents you from accidentally picking a variant that isn't better for your goal metric than any other variant. This is known as a false negative. This is useful but not the most important observation in the context of website or online optimization. Instead, a type 2 error is the most costly error. This error occurs when you might fail to choose a better variant and is managed by reaching statistical power. This is considered a false positive. In this case, you're missing out on conversions or, perhaps even worse, negatively influencing performance.

Dynamic Testing experiences naturally manage these errors in its statistical model and make use of data as it becomes available to exploit your goal metric.

How Dynamic Testing Experiences Work

Are Dynamic Testing experiences always shorter than a standard A/B test?

While Dynamic Testing experiences yield results much faster than a standard A/B test and with the same statistical validity, there may be occasional experiments that take longer. This can occur because Dynamic Testing experiences are balancing learning and optimizing towards your goal metric.

Do Dynamic Testing experiences work in low-traffic conditions?

Yes, a Dynamic Testing experience works immediately. It begins with no knowledge of the performance of the variants and learns over time. As the probability of a winner rises, the Dynamic Testing experience makes incremental adjustments.

How does the Dynamic Testing experience model react to outliers?

The algorithm learns from a naïve state and continues to learn based on a cumulative understanding of all collected data. Specific outliers may have small effects on a specific day, but the model will analyze the changes your variants have on traffic cumulatively.

How do Dynamic Testing experiences handle two variants that perform equally well?

When two variants perform exactly the same, a Dynamic Testing experience balances between the two by alternating, exploring, and exploiting. Here, the statistical engine behaves this way on average, but at any given time, one variant may accumulate more sessions than another due to chance.

Remember that clients run Dynamic Testing experiences because they think they can improve the site experience for their audiences by avoiding overexamining the worst case scenario assumed by standard A/B testing.

Do site visitors see multiple variants?

No, visitors who arrive on your site and begin their analytics session do not see more than one variant. However, if a visitor leaves your site and returns after the previous session ended, they may be reallocated to a higher-performing variant.

When should I end a Dynamic Testing experience?

If possible, you should never end a Dynamic Testing experience. Monetate's decision engine always observes variant performance and dynamically promotes the best-performing variant. This activity reduces the risk of making an incorrect assumption or decision by ending a Dynamic Testing experience.

When should I promote the current leader within a Dynamic Testing experience to 100%?

Monetate's decision engine promotes the best variant for you automatically. If you stop a Dynamic Testing experience and promote a winning variant to 100%, you introduce risk because the variant you promoted to your entire audience may not remain the best experience in the future. If you continue to run a Dynamic Testing experience, you reduce this risk and allow Monetate to dynamically ensure that the best experience is displayed at the right time all the time.

What if I can't let a Dynamic Testing experience continue to run and need to promote a winner to align with an internal business initiative?

If you absolutely need to stop a Dynamic Testing experience and promote the best variant to 100% of your audience, you should let the test run until the point in time when you absolutely need to select a winner and promote that variant to your entire audience.

You should consider the risks and costs associated with stopping a Dynamic Testing experience before you do it.

When you determine which variant is the best, you should focus on the current chance to win percentage and lift for the best variant. You also should understand the value in continued learning or how often the non-winning variants perform better than the best variant by continuing to run as a Dynamic Testing experience.

If the value in continued learning is acceptable and the current chance to win is high for a specific variant, you can feel comfortable promoting that variant to 100% of your audience if you absolutely have to.

You can schedule a Dynamic Testing experience to end by adding an end date to the WHEN settings of any Dynamic Testing experience, or you can pause the test at any time.

What if a Dynamic Testing experience has found multiple high-performing variants and multiple low-performing variants?

Let the Dynamic Testing experience continue to run because doing so provides you the best chance to automatically capitalize on performance changes that drive ROI over time.

If you absolutely need to end the Dynamic Testing experience in this instance, you can do so a few different ways:

- Take the two best variants that seem to perform evenly and introduce an additional variant in a new Dynamic Testing experience.
- If you need to pick a winning variant at this time, review the best-performing variants and their current chance to win along with the value in letting the Dynamic Testing experience continue to run. If you are comfortable with selecting a variant to promote to 100%, select one of the best-performing variants and promote to 100% in a new experience.

Can I promote the leader in an experience with a 60/20/20 variant?

The option to promote the current leader is available to you at any point during the life of a Dynamic Testing experience. When evaluating the current performance of a Dynamic Testing experience, you should understand specifically what the current chance to win represents in this case:

- There's a 60% chance that variant A is the best experience to optimize your goal metric.
- There's a 40% chance that variant B and variant C are the best experiences to optimize your goal metric.

If you were to promote the current leader, which has a 60% current chance to be the winner, there's a 40% chance that you are making the wrong decision to show that variant to 100% of your traffic.

You should consider the value provided by the non-leading variants and the cost to maintain the Dynamic Testing experience before making a decision to promote or let the Dynamic Testing experience continue to run.

To minimize risk and continue to maximize the ROI from the Dynamic Testing experience, always let the test continue to run.

How does a visitor qualify for a Dynamic Testing experience?

Visitors must qualify for all splits in Dynamic Testing and Automated Personalization experiences. Therefore, all conditions and all recommendation strategies must have results if recommendations are used. The reason why is that the engine decides between all splits and then places visitors in the best one, so they must be eligible to see all of them.

Does a visitor remain in the same variant during a Dynamic Testing experience?

Not always. A visitor can be reassigned if the variant they saw in a prior session is not the current leader and is now receiving a smaller percentage of traffic. Therefore, every visitor session has the opportunity to experience the best variant to maximize as much return as possible during a Dynamic Testing experience. Dynamic Testing experiences do not change a visitor's variant within a single session.

How do Monetate's cross-device capabilities work with Dynamic Testing experiences?

Dynamic Testing experiences leverage cross-device behavioral targeting for Dynamic Testing experiences to qualify a person with behavior exhibited across all devices on which they're recognized. Dynamic Testing experiences do not leverage cross-device testing by keeping an individual in the same variant across devices. Given that the purpose of Dynamic Testing experiences is to maximize ROI by showing the best variant as often as possible, Monetate believes keeping a person in a variant that may actually be the worst performer does not provide value. Each visitor session has the opportunity to see the best experience based on the current traffic distribution for the Dynamic Testing experience.

Understanding Dynamic Testing Experience Analytics

When can I start viewing results?

With standard A/B testing, you have to wait until the test has enough sessions to be powerful and meet a fixed horizon before you can view results. However, with Dynamic Testing experiences you can continuously monitor analytics and inspect progress at any time.

As sessions close, they are included in Dynamic Testing experience analytics. This means that the data reflects site activity as recent as 30 minutes ago.

How often are results updated?

Monetate operates on a real-time streaming infrastructure. As sessions stream across the platform and close, analytics update and so too the decisioning engine that powers Dynamic Testing experiences. Dynamic Testing experience analytics update every 60 to 90 seconds with the latest closed session data.

What metrics can I use?

All eCommerce metrics are available to you, depending on the data your organization passes to Monetate. Check with your account manager if you aren't sure about which metrics you have available. In addition, you can use any page event that you build in Event Builder as a goal metric for a Dynamic Testing experience.

Probability went down from yesterday. How can I trust this test?

Traffic and performance changes occur all the time and are influenced by a number of different factors, such as seasonality, pay cycles, promotions or sales, and new product releases.

Dynamic Testing experiences automatically recognize these changes and begins to adjust the traffic, if necessary, to capitalize on the best-performing variant at that moment in time. Not only can you trust the test when you see the current chance to win fluctuate, but you can also trust that Monetate makes these changes to drive the highest ROI possible based on performance observations over time.

Is there a way to make Dynamic Testing experiences more conservative?

The purpose of Dynamic Testing experiences is to recognize and maximize ROI as quickly as possible. You may see traffic and performance fluctuate over time as you learn whether there is a meaningful difference between the variants configured in the Dynamic Testing experience.

This means that Monetate recognizes performance changes and explores whether a variant is outperforming another.

If you are concerned with the fluctuation of traffic distribution and you want more control over the exposure of specific variants, you should consider leveraging the Standard Test experience with a fixed traffic distribution.

I launched a new test on my homepage and other Dynamic Testing experiences on the same page changed their leaders. Have I introduced a statistical error?

Absolutely not. Dynamic Testing experiences automatically handle the statistical error and reacts to changes over time. When you run simultaneous experiences on the same page that start at different times, Dynamic Testing experiences recognize if those changes have an impact on goal performance and handle traffic redistribution to maximize as much ROI as possible.

Remember that the purpose of a Dynamic Testing experience is to drive ROI, regardless of which variant is the best at any point in time. It's not about finding a single statistically significant winner.

Should I be concerned about the impact a Dynamic Testing experience may have on my non-goal metrics?

Today, Dynamic Testing experiences operate off the performance of a single goal metric. It's important to be thoughtful when you think through the purpose and goal of your test before you decide to use standard A/B/n testing or Dynamic Testing experiences. If you want to continually monitor the performance of all metrics and have control over the exposure of those variants throughout the test, you should consider using standard A/B/n testing.

When the purpose of your test is to drive ROI for a single goal metric, a Dynamic Testing experience is the best solution. It's also important to consider the appropriate goal metric to impact as a result of a Dynamic Testing experience. Dynamic Testing experiences optimize based on the performance of the selected goal metric only and do not take into account secondary metrics.

Within Dynamic Control experience analytics, you can monitor the performance of secondary metrics for the current leader and the overall performance of secondary metrics across all variants.

Monetate only displays the performance of metrics for the current leader because this variant receives the bulk of the traffic included in a Dynamic Testing experience. If you are concerned about the performance of secondary metrics as a result of the dynamic optimization, you could consider running additional Dynamic Testing experiences with different goal metrics or consider using standard A/B testing.

Can I edit a Dynamic Testing experience's action conditions after I activate it?

You cannot edit Dynamic Testing experiences once you activate them. For example, you want to debug a Dynamic Testing experience that was configured to target a product detail page, but that action condition was removed for debugging purposes. However, that condition still exists in the activated state of the Dynamic Testing experience, but you can still modify the action conditions and other aspects of the experience by duplicating the experience and properly configuring all conditions before you activate it.

Refer to [Build a Dynamic Testing Experience](#) for more information.