

Extracted from:

# Practical A/B Testing

## Creating Experimentation-Driven Products

This PDF file contains pages extracted from *Practical A/B Testing*, published by the Pragmatic Bookshelf. For more information or to purchase a paperback or PDF copy, please visit <http://www.pragprog.com>.

Note: This extract contains some colored text (particularly in code listing). This is available only in online versions of the books. The printed versions are black and white. Pagination might vary between the online and printed versions; the content is otherwise identical.

Copyright © 2023 The Pragmatic Programmers, LLC.

All rights reserved.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form, or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior consent of the publisher.

The Pragmatic Bookshelf

Raleigh, North Carolina

The  
Pragmatic  
Programmers

Pragmatic  
exp

# Practical A/B Testing

Creating Experimentation-Driven  
Products



Leemay Nassery  
*Edited by Vanya Wryter*



# Practical A/B Testing

Creating Experimentation-Driven Products

Leemay Nassery

The Pragmatic Bookshelf

Raleigh, North Carolina



Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this book, and The Pragmatic Programmers, LLC was aware of a trademark claim, the designations have been printed in initial capital letters or in all capitals. The Pragmatic Starter Kit, The Pragmatic Programmer, Pragmatic Programming, Pragmatic Bookshelf, PragProg and the linking *g* device are trademarks of The Pragmatic Programmers, LLC.

Every precaution was taken in the preparation of this book. However, the publisher assumes no responsibility for errors or omissions, or for damages that may result from the use of information (including program listings) contained herein.

For our complete catalog of hands-on, practical, and Pragmatic content for software developers, please visit <https://pragprog.com>.

The team that produced this book includes:

CEO: Dave Rankin

COO: Janet Furlow

Managing Editor: Tammy Coron

Development Editor: Vanya Wryter

Copy Editor: L. Sakhi MacMillan

Layout: Gilson Graphics

Founders: Andy Hunt and Dave Thomas

For sales, volume licensing, and support, please contact [support@pragprog.com](mailto:support@pragprog.com).

For international rights, please contact [rights@pragprog.com](mailto:rights@pragprog.com).

Copyright © 2023 The Pragmatic Programmers, LLC.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form, or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior consent of the publisher.

ISBN-13: 979-8-88865-008-0

Encoded using the finest acid-free high-entropy binary digits.

Book version: P1.0—May 2023

A/B testing is a key part of product development. On the surface, it may seem like a complex methodology that requires fancy tools and elaborate engineering systems. However, this isn't the case if you're just getting started. You don't need to know every detail of A/B testing to integrate this practice into your product development process, but understanding the fundamentals is necessary.

This book will expose you to concepts that can help kickstart incorporating A/B testing into your product. We'll lay out the anatomy of a test and the core components required to facilitate A/B testing. We'll explore practical concepts and strategies to incorporate A/B testing within your organization, starting with the many advantages of this experimentation methodology. Specifically, this chapter includes the following:

- How your team can benefit from A/B testing.
- How you can gain user insights you didn't even think to seek at the start of a project.
- How you can use A/B testing to catch engineering system vulnerabilities earlier rather than later.
- How you can leverage A/B testing as a risk management tool.

To further illustrate the advantages of A/B testing, you'll have the role of analyst for the personalization team at CableMax. Although CableMax is a fictitious company, it's loosely based on a real-life context that used A/B testing to improve the experience for a user-facing product. There will be analyst tasks related to your role at CableMax to get some hands-on experience.

So, without further ado, let's get started!

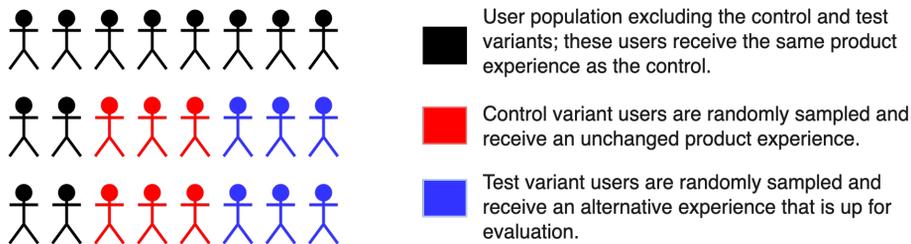
## What Is A/B Testing?

A/B testing is an online controlled experiment that measures the impact of a change on a subset of users. An effective A/B test is one where you feel confident in making decisions based on the results.

An A/B test, also referred to as an experiment, in its simplest form consists of two groups of users, as shown in the following image. The first group is the control variant which includes users who receive the unchanged functionality or product experience. The second group consists of users who receive a new feature or change that is up for evaluation, known as the test variant. Both the test and control variants need to be randomized to ensure that the

users assigned to each group are similar statistically, allowing for higher confidence when measuring the effect of the experiment. If the variants are not properly randomized, then you wouldn't know with high certainty that the changes evaluated in the experiment caused an increase or decrease in key metrics.

### User Population



Now to determine the effectiveness of the change, you'll use metrics to compare the engagement of the test variant to the control variant. Your metrics and variants are closely tied together. Without proper variant randomization, your test results may represent an inaccurate estimate of the effect of the change, as they could be biased toward a specific user population or other unknown factors. We'll be looking more closely at A/B testing terminology in [Learn the Fundamentals of an A/B Test](#), but for now, let's start with a basic understanding.

You're evaluating changes on real users and measuring the outcome using data. The change that is evaluated as part of the experiment could include:

- A new UX design to a user-facing product.
- A new feature or change to an existing feature.
- A new software architecture that, for example, improves the efficiency of the engineering system under the hood but, ideally, should be transparent to the user.

If you're building user-facing products, your goal is to improve the user experience and create business impact. When you leverage A/B testing to evaluate a change, the test results allow you to measure if you're meeting your intended goals.

## Who Should Run A/B Tests

Everyone and anyone should run A/B tests: product owners, engineers, engineering managers, designers, data scientists, user researchers, and really, anyone who has the desire to learn about the product they're building.

If you have a theory and the desire to prove your theory is true, then you should utilize A/B testing.

When you should run an A/B test is subjective. You may prefer all user-facing changes, whether big or small, to be A/B tested. Or you may choose a less conservative approach: A/B test more significant changes. For more minor changes, compare your metrics before versus after the change was rolled out to all your users. It really just depends on your use case and what type of questions you're looking to answer.

Here are a few points to consider when deciding on whether to evaluate a change with an A/B test:

- If your product is new or you're testing a brand-new domain, your metrics may be more volatile. In this case, it's better to start measuring changes early as you iterate on the product over time.
- If your product is mature, your metrics may not be as sensitive to changes, making it a bit harder to move the needle. That's not to say that you shouldn't leverage A/B testing. You may want to opt for an equivalence test, which will be discussed in detail in [Select the Right Type of Experiment](#).
- If you're keen on protecting the user from an unintentional degradation in the product experience, use A/B testing to gather data insights into the performance of change on a subset of users first.
- If you find it necessary to introduce the change at a lower scale first, to evaluate the risk or the potential for uncovering scale issues before it's available for all your users.

If you're still wondering whether you should evaluate a change using the A/B testing methodology, consider the following questions that data from an A/B test could answer:

- How did this change impact key business metrics?
- How are specific minority user groups affected by the new feature?
- How did this change influence engagement on other surfaces of the product?

This is all to say, if you have questions about the usage and the impact of changes made to the product, then you're the perfect candidate to run an A/B test.

## Beginning Your Analyst Role

You were recently hired at CableMax, a company with a large footprint in the United States, providing video services to millions of users.

Given the company's size, you'd assume there would be a sophisticated A/B testing platform that supports the innovation made to the video product. That's not the case. Even big companies need help getting started with A/B testing.

The engineering and product teams at CableMax want to better understand how new features influence user engagement and impact key business metrics. This is why they've asked you to help them push their A/B testing practices forward. You'll start with a team known to push changes straight to production, with little evaluation or measurement beforehand: the personalization team.

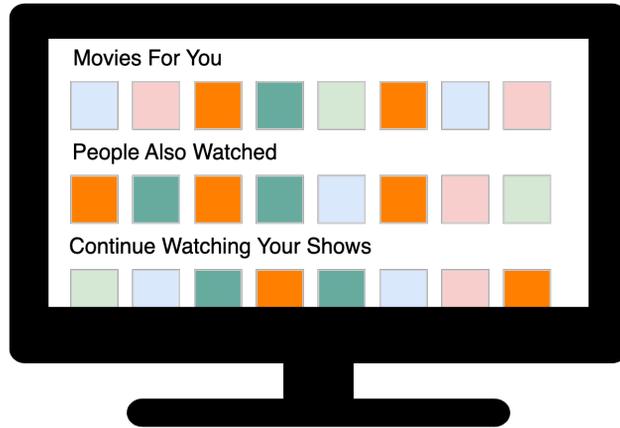
The personalization team is responsible for features that enable users to quickly access content to watch on the company's flagship video product. To get an idea of the user experience you're providing analyst support for, see the [image on page 5](#). The colored squares represent individual TV shows or movies that are personalized based a user's taste and watch history.

Since you're a new hire, let's set the stage a bit more.

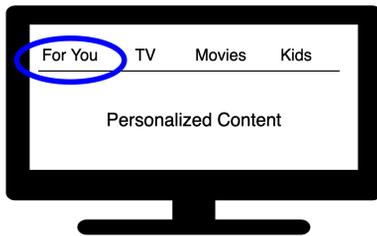
It's clear what is missing from the video product: a personalized For You homepage. CableMax users expect good value for what they pay for, so naturally, you want to make it as easy as possible to engage with. The For You homepage will be a single location where the user can find recommended content just for them, instead of the typical editorial experience that is "one-size-fits-all."

This is a great idea, but all ideas, whether great or not-so-great, should be A/B tested. And that's exactly what the team at CableMax decides to do: launch a For You A/B test on the video product.

To evaluate the For You homepage, users in the test variant will receive the For You experience (left). Users in the control variant will receive the unchanged, unpersonalized experience (right). See the following image.



Test Experience



Control Experience

